TEXTBOOK · A-LEVEL MUSIC TECHNOLOGY · COMPONENT 4 · 1.3

The Story of Synthesis

A Century of Electronic Sound, 1920 to Now

BY MIKE LEHNERT

THE STORY OF SYNTHESIS

Contents

Introduction

This book tells one story in five parts: a century of arguing about what a sound is allowed to be. It runs from a wooden cabinet played in 1920 without ever being touched, to a rack of modules a musician chooses to patch from empty, on purpose, this year.

Every synthesiser in these pages is a position in that argument. How a performance already in the air got captured, blended and rebuilt once a microphone had heard it is told in The Story of the Studio, Volume One on this shelf — a story driven by constraint on capture. This book asks an earlier question: where does the note come from before any microphone is involved, and what is a musician allowed to do to it once it exists? That is a story driven by constraint on the instrument itself — what one machine will let you play, save, copy or touch.

It is written for Section 1.3 of Component 4 of the Pearson Edexcel Level 3 Advanced GCE in Music Technology: Synthesis. That section asks you to programme a synthesiser with intention and then explain, precisely, why your choices produced the sound they did. Both halves of that task rest on the same underlying skill — hearing a sound and naming the mechanism that made it — and that is what this book exists to teach.

NOTE

This is a vocabulary book as much as a history book. Its history explains why oscillators, filters, envelopes and modulation exist in the order they do; its Hearing chapters and its one framework chapter give you the exact words the specification expects you to use for them.

So as you read, keep asking two questions: what would I hear, and what is the correct name for it?

How This Book Is Arranged

Each of the five parts opens with a drawing of the instrument that defines its era, then tells that era's story — what changed, who built it, and why it mattered enough to argue about. A Hearing chapter follows, turning the era's sound into exam-ready description. Part One carries one chapter more than the rest: a framework chapter that names, once and for all, the four building blocks — Source, Filter, Amplifier, and Control over all three — that every synthesiser since has only ever rearranged. Learn that chapter and the remaining four parts read as variations on it, not new material.

Five Eras, Named by Their Constraints

constraint line you will meet again on that part's opening page.

ELECTRICITY FINDS A VOICE · 1920–1963	The instrument is a laboratory: one new sound at a time, and almost nobody can play it.
THE PATCH-CORD YEARS · 1964–1977	One voice, patched from scratch. Nothing can be saved.
PRESETS AND CHIPS · 1978–1995	Sound becomes storable — then becomes numbers.
SOFTWARE AND THE INFINITE INSTANCE · 1996–2009	The instrument becomes a file: free to duplicate, endless by default.
THE MODULAR RETURN · 2010–NOW	Musicians choose scarcity back on purpose.

What Changes Between Eras

- **Electricity Finds a Voice → The Patch-Cord Years:** the laboratory becomes an instrument you can tour with.
- **The Patch-Cord Years → Presets and Chips:** a patch becomes a memory.
- **Presets and Chips → Software and the Infinite Instance:** hardware becomes a file.
- **Software and the Infinite Instance → The Modular Return:** unlimited gets traded back for scarcity, deliberately.

Place a synthesised sound by what its instrument would let you do, not by the year it was made. Each era's limitation is what you can hear.

The five instruments on the part-opening pages tell the story on their own: a theremin played without ever touching it, a patch bay dense with cable, a keyboard that remembers, a laptop with an unreasonable number of windows open, and a modular case somebody chose to buy anyway.

A Note on What the Specification Actually Requires

Section 1.3 names no particular synthesiser and no historical instrument for any of these five eras. You need the categories: oscillator waveforms, filter types, envelope stages, modulation sources. Everything else in this book — the named machines, the named musicians, the arguments they were having — is enrichment. But it is the enrichment that turns a labelled diagram into a justified decision, which is where the marks are.

identification drill, a timeline, and full plate notes.

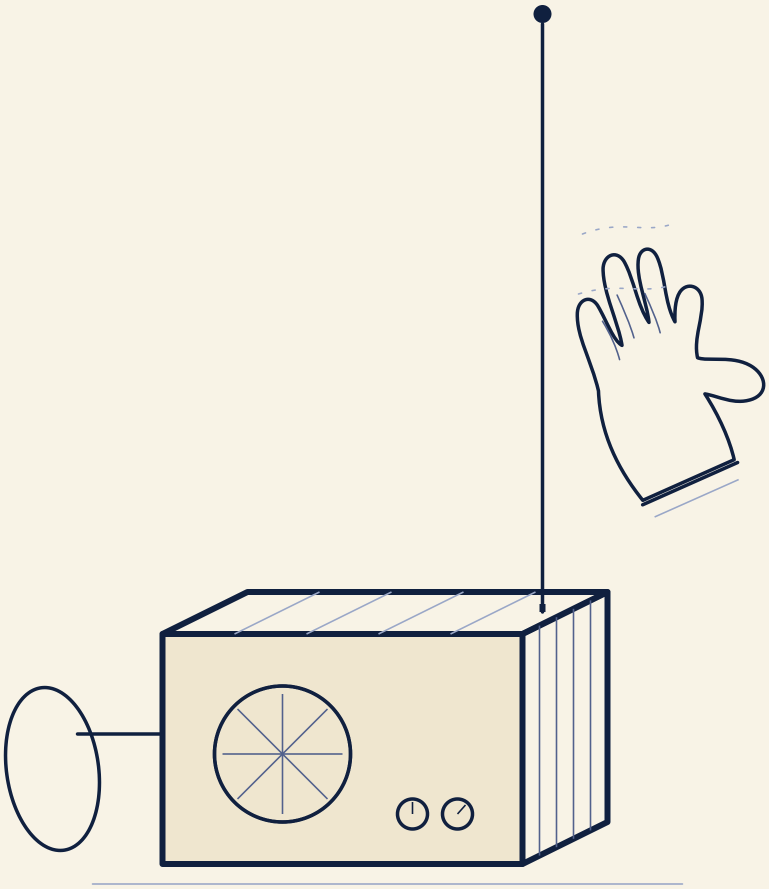

PART ONE · 1920–1963

Electricity Finds a Voice

The instrument is a laboratory: one new sound at a time, and almost nobody can play it.

CHAPTER ONE

The Laboratory Instrument

NOTE

The specification does not name any of the machines in this chapter, and 1.3 Synthesis is not tested by dating a sound to a decade. This history earns its place for a different reason: it is the evidence a "how synthesis developed" extended-writing answer needs, and it explains why the technical vocabulary this part builds towards had to be invented at all.

Every machine in this chapter was built by a physicist, a radio engineer or a broadcasting technician — not one of them by an instrument maker working for a musician. That single fact decides everything else about them. It decides who could learn to play one, whether a working musician could ever own one, and why, by 1963, not a single electronic instrument had found its way into ordinary musical life.

A Voice Made from Dynamos

The idea that electricity could generate a musical tone, rather than merely amplify one, is older than most people expect. In 1897 the American inventor Thaddeus Cahill patented a machine that generated pitched tones from banks of spinning, precisely tuned dynamos, and in doing so used a word that would outlive the machine itself: his patent describes a device built to "synthesize" the vibrations of music electrically. Cahill built the instrument, the Telharmonium, in New York in 1906, and it worked — its tones travelled down ordinary telephone lines into hotel dining rooms and restaurants as a novel form of piped music. But the Telharmonium was never anything a musician could own. It filled a building, its running costs were enormous, and its powerful signals bled into telephone conversations across the network it depended on. The company behind it collapsed within a decade. The Telharmonium proved that electricity could make music; it did nothing to prove that a musician could carry that music out of the building.

Played Without Touch

In 1920, the Russian physicist Leon Theremin built an instrument that solved the size problem completely and created an entirely new one. Developed out of research into proximity sensing, the Theremin generates its pitch from the interference between two radio-frequency oscillators, one of which no string, no fixed point of reference of any kind — only open air and the performer's memory of where a pitch happens to sit in space.

That absence of physical reference is precisely why the Theremin never became a working musician's instrument. Its wavering, disembodied tone made it a gift to film composers, but playing it accurately in tune demanded a kind of trained muscle memory that almost nobody had the years to acquire. Clara Rockmore, a violinist whose own playing career had ended early, became the instrument's great virtuoso and performed it with major orchestras — a demonstration of what was possible, achieved by one of a handful of people alive who could do it. A laboratory curiosity does not stop being one just because a genius eventually masters it.

The One That Reached the Orchestra

Maurice Martenot took a different route in 1928. His Ondes Martenot generated pitch by the same heterodyne principle as the Theremin, but he gave it a piano-style keyboard for the right hand, so that any pianist could sit down and find a note by sight rather than by ear alone. The keys could even rock slightly from side to side under the finger to add vibrato, and a ring worn on that same hand, threaded to a sliding wire behind the keyboard, offered a theremin-like continuous glide as an alternative to the keys. A separate control under the left hand, housed in a sliding drawer, shaped intensity and timbre in real time. That keyboard is the whole difference. The Ondes Martenot is the one instrument in this chapter that entered serious orchestral repertoire and stayed there: Olivier Messiaen scored it into his vast *Turangalîla-Symphonie*, and specialist performers, known as ondists, still play the instrument in orchestras that programme the piece today.

Yet the Ondes Martenot's success only sharpens the chapter's point rather than contradicting it. Only a few hundred instruments were ever built, learning one properly took years of dedicated study, and the whole tradition survived as a rare specialism inside classical music rather than as something a touring musician could buy from a shop and carry to a gig. The Ondes Martenot got further than any other machine in this chapter. It still did not get all the way to an ordinary working musician.

Composed on a Paper Roll

Built by RCA engineers and delivered to Columbia University in 1957, the RCA Mark II Sound Synthesizer removed the performer from the process altogether. There was no keyboard to play in real time and no antenna to wave a hand near. Instead, a composer specified every pitch, duration, timbre and dynamic level as a sequence of punched holes on a roll of paper tape — much like the roll that drives a player piano — and fed that roll into the machine, which read the instructions and generated the sound accordingly. Composers including Milton Babbitt worked this way, programming a piece over days or weeks and only hearing the result once the roll had run through. Nobody played the RCA Mark II the way a violinist plays a violin. It filled a room at a university electronic-music studio, one machine built for the purpose, operated by composers with research access to that room — not by musicians, and not anywhere a musician could reach.

The BBC Radiophonic Workshop, founded at Maida Vale in London in 1958 with Daphne Oram among its co-founders, existed to build original electronic sound for radio and television drama. It had no synthesiser at all, in 1958 or for a decade afterwards. What it had was a workshop: banks of test oscillators borrowed from the BBC's own engineering department — equipment built to check that transmitters were working, never to make music — alongside tape recorders, filters and a great deal of patient, manual razor-blade splicing. Every sound the Workshop produced was assembled one recorded event at a time, by hand.

THE DISTINCTION THAT MATTERS

In 1963, Delia Derbyshire built the Radiophonic Workshop's most famous piece of work: a realisation of Ron Grainer's new theme tune for the BBC's new science-fiction series, *Doctor Who*. She recorded each note of the melody separately from the Workshop's bank of test-tone oscillators onto quarter-inch tape, then measured and cut each note to its exact length with a razor blade, and spliced the pieces back together in rhythm — building the bass line, the melody and its swooping glissando as three separate tape loops before mixing them down together. There was no synthesiser in the room. There was not yet one anywhere in the building. What became one of the most recognisable pieces of electronic music of the century was built entirely from individually recorded oscillator tones and a roll of splicing tape, note by physically cut note — a landmark of tape technique, not of synthesiser performance. "Sounds electronic" and "made on a synthesiser" are not the same claim, and this piece is the cleanest possible proof that they are not.

Why None of Them Became Instruments for Musicians

THE ONE FACT THAT EXPLAINS THE CHAPTER

Every machine on these pages made a new sound. Not one of them put an instrument into a musician's hands that could be learned like a guitar and carried to a gig — because playing one required a trained virtuoso's body, a university research post, or a room full of borrowed test equipment and a razor blade.

Line the five machines up and the pattern holds exactly. The Telharmonium was too vast and too costly for anyone to own. The Theremin could be owned, but its total lack of any physical reference point put accurate playing beyond all but a handful of trained virtuosi. The Ondes Martenot solved that problem with a keyboard and reached the concert hall for real, but remained a rare specialism, a few hundred instruments deep, learned over years inside classical music rather than sold to working musicians. The RCA Mark II was never played by anyone in real time at all; it was programmed, like a player piano roll, sound was not performed on an instrument in any sense — it was built, tape splice by tape splice, from raw test-oscillator tone.

The common thread running under all five is the thing this part is named for: each produced one new sound at a time, at the cost of specialist physical skill, laboratory access, or hours of manual construction that an ordinary musician simply did not have. None of them put a single object into a musician's hands that could be learned, carried to a rehearsal, and played the same way twice. The instrument for musicians was still to be built. It would take a different kind of engineer, working in a different discipline entirely, to build it.

CHAPTER TWO

Hearing the Laboratory

This is your listening toolkit: place an unfamiliar recording by the technology you can hear.

What to Listen For

- A pure or near-sine tone — almost no harmonic buzz or edge, nothing like the bite of a bowed string or a brass instrument; a smooth, rounded pitch closer to a wordless human voice than to any conventional instrument
- No attack transient — the note swells into existence rather than starting; there is no pluck, strike or reed-buzz onset, because no envelope generator exists yet to shape one. Whatever “attack” you hear is supplied live, by a hand
- Stepless, continuous pitch — no frets and no keyboard divisions to lock the player to fixed notes, so portamento and glissando are the instrument’s default behaviour, not an effect switched on for colour
- Vibrato from a body, not a circuit — a small, uneven waver in the hand near an antenna or in finger pressure on a ribbon, irregular in rate and depth in a way no later, circuit-driven modulation will ever quite reproduce
- Mechanical exactness where a machine did the playing instead — pitches arriving with a precision no wrist could match, because a wrist was never involved; the tell for anything built from punched paper tape rather than a player’s gesture
- Tape artefacts around material assembled by splicing — a soft click or a hairline pitch and tempo hiccup where two lengths of tape were joined, a faint hiss underneath, and pitches confined to a short discrete scale because each note began life as a separately recorded tone

Writing About What You Hear

A strong answer is a chain: technique → sonic outcome → judgement. In this era the technique is usually a gesture, not a circuit, so name the gesture.

MODEL SENTENCE

Hollywood composers found a use for the theremin that its inventor never anticipated: not a concert instrument in its own right, but a single, uncanny colour dropped into an orchestral score. Miklós Rózsa scored the dream sequences of Spellbound (1945) with a theremin part played by Dr Samuel Hoffman, and the same instrument turned up again under Bernard Herrmann for the flying-saucer cues of The Day the Earth Stood Still (1951). In both scores, the theremin's continuous glide between pitches and its complete absence of any attack transient do a job no orchestral instrument could: they remove the sense of a note being struck or bowed at all, so the line seems to arrive from nowhere and leave the same way. Rózsa and Herrmann were not treating the theremin as a flawed violin — they were exploiting exactly the qualities that made it useless as a mainstream orchestral instrument, and turning that limitation into a description of the uncanny.

That delivers: named technique (continuous pitch, no attack transient) · named era / instrument / composer · named works · sonic outcome · critical judgement (limitation repurposed as expressive device).

A STRONGER JUDGEMENT (COUNTER-ARGUMENT)

Not every laboratory instrument in this era can be heard as a performance. The RCA Mark II Sound Synthesizer, built at RCA's Princeton laboratories and delivered to Columbia University in 1957 — two years before Columbia and Princeton jointly founded the Columbia-Princeton Electronic Music Center around it — produced its pitches from punched paper tape rather than a player's hand, and the difference is audible immediately: Milton Babbitt's Composition for Synthesizer (1961) has none of the theremin's waver, glide or human unsteadiness. Its intervals land with a precision no wrist-mounted vibrato could match, because no wrist was involved. A sympathetic listener might argue that this makes it more purely synthesised than the theremin — mechanically exact where an instrument is normally expected to carry some grain. A sceptical one could reply that removing the player removed the music-making too, and that precision without gesture is closer to a specification than to a performance. Either argument is defensible; a strong essay answer states which side it is taking before it makes it.

ESSAY VOCABULARY: HOW THE PITCH WAS MOVED

A quick way to place an unfamiliar extract in this era is to ask how the performer altered a note once it had started, because three quite different answers were all current between the 1920s and the early 1960s. By touch, continuously, in real time (theremin, Ondes Martenot) — the note bends because a hand moved. By punched tape, precisely, in advance (RCA Mark II) — the note is exact because nobody moved anything while it sounded. And by splice, after the fact, at a workbench (Radiophonic Workshop) — the note was fixed the moment the tape was recorded, and everything audible afterwards is editing, not performance. Naming which of the three you are hearing, and why, does most of the descriptive work a listening answer needs.

Clara Rockmore, the One Exception

Almost nobody could really play these instruments, which is exactly why Clara Rockmore matters for a listening question. A violin prodigy whose training was cut short in her teens by illness, Rockmore approached the theremin with a violinist's ear for intonation, developing a precise aerial fingering technique — measuring pitch distances in the air the way a string player measures them on a fingerboard — that let her land intervals cleanly rather than sliding vaguely towards them. Set her recordings against an amateur theremin part in a low-budget film score and the difference is not the instrument, it is the player: the same continuous, gliding pitch control that sounds eerie and approximate in untrained hands becomes controlled and lyrical in hers. In an essay, naming Rockmore lets you make a sharper claim than "the theremin sounds wobbly" — the wobble is a limit of the player, not a property of the machine, and a virtuoso could all but erase it.

Instruments and Records to Know

THEREMIN	Leon Theremin (1920) · Clara Rockmore, virtuoso concert performer · Dr Samuel Hoffman, Hollywood's session theremin player (Spellbound, 1945; The Day the Earth Stood Still, 1951)
ONDES MARTENOT	Maurice Martenot (1928) · Olivier Messiaen, Turangalîla-Symphonie (1946–48) — the instrument's best-known orchestral outing
TELHARMONIUM	Thaddeus Cahill (1897/1906) · no recording survives — a reminder that "hearing" this instrument today means reading contemporary description, not listening
RCA MARK II SOUND SYNTHESIZER	RCA, Princeton (1957) · delivered to Columbia University; housed at the Columbia-Princeton Electronic Music Center from 1959 · Milton Babbitt, Composition for Synthesizer (1961)

BBC RADIOPHONIC WORKSHOP	Founded 1958 (Daphne Oram among its founders) · Delia Derbyshire's 1963 tape realisation of the Doctor Who theme (composed by Ron Grainer), assembled from tape loops of individually recorded test-oscillator notes

What the Spec Asks

BASIC WAVEFORMS	Spec: “sine, square, sawtooth, triangle”. This era is built almost entirely from the first of these — laboratory instruments produce close to a pure sine tone, with none of square’s hollowness or sawtooth’s brightness heard again until the next chapter, The Four Blocks.
PITCH CONTROL	Spec: “Frequency modulation and tuning”. Heard here as continuous, hand-controlled glide — portamento and glissando — rather than the fixed pitch-steps every later era takes for granted.
ATTACK TIME	Spec: “How quickly sound reaches peak level”. There is no attack time to measure yet — no envelope generator exists in this era, so the honest answer is always “however quickly the player’s hand allowed”.
MODULATION SOURCES	Spec: “Creating movement in synthesised sounds”. In these instruments the modulation source is the performer’s own hand or finger pressure, not a circuit — a human being standing in for the LFO that the next chapter, The Four Blocks, will formally introduce.

THE FRAMEWORK

The Four Blocks

Source, filter, amplifier — and a fourth job that never gets its own knob, because it moves all three.

Every synthesiser in this book, whatever decade it was built in and however it makes its first sound, does the same three jobs in the same order. A source makes a raw tone. A filter sculpts that tone's colour. An amplifier switches it on and shapes how it dies away. Everything else — knobs, patch cables, menus, plug-in windows — is a different way of reaching those same three stations, plus a fourth layer that moves them while you play.

This is the chapter you will come back to. From here on, every era in this book is really a story about who controls these four things, how many hands are needed to do it, and whether the settings survive after the power is switched off. Learn the names now and the rest of the book reads itself.

The Signal Path

Strip any synthesiser down to what the audio itself actually passes through — a 1970 Minimoog Model D, a Eurorack module patched together last week, a free softsynth running in a browser tab — and there are only ever three stages.

THE PATH SOUND ACTUALLY TAKES

1. **Source.** An oscillator or a noise generator makes the raw tone. Nothing downstream matters yet — there is no sound to shape until something exists to shape.
2. **Filter.** The raw tone is coloured — brightened, darkened or narrowed — before it goes any further.
3. **Amplifier.** The coloured tone is switched on, held and let fade. This is the last stop before the output socket.

Source → filter → amplifier. Everything else in this chapter is control — it does not sit in that chain, it reaches into it.

Block One: The Source

occasionally both mixed together. A waveform's shape is what a listener actually hears as its character, long before any filter has touched it. Four shapes turn up in almost every synthesiser built since the 1960s, and each has an unmistakable identity by ear.

- **Sine** — the purest tone available: smooth and round, with nothing extra riding on top of it. Close to a tuning fork, or a flute in its lowest register. Because it carries so little character of its own, it is rarely a whole patch by itself — more often it sits underneath a brighter wave, or supplies the shape for a slow-moving LFO (see below).
- **Triangle** — soft and slightly hollow, a gentler cousin of the square wave. Flute-like and mellow, good for a bass or pad that needs body without edge.
- **Square** — hollow and reedy, with a nasal, clarinet-like quality. Narrowing the width of the square itself sets the tone moving, sweeping between thinner and fuller versions of the same note as it changes.
- **Sawtooth** — the brightest and buzziest of the four, close to a bowed string or a brass section heard from a distance. It is the standard starting point for subtractive synthesis precisely because it gives the filter that comes next the most to work with.

Alongside the four pitched waveforms sits **noise** — a signal with no identifiable pitch at all, just energy spread across the spectrum. White noise sounds like an untuned radio between stations; filtered down and shaped by a fast envelope, it becomes a snare, a hi-hat, or the breath at the front of a flute patch. A source is rarely one waveform alone in practice: a little noise mixed under a sawtooth, or two oscillators detuned slightly against each other, is how a single thin tone becomes a full one.

Block Two: The Filter

If the source decides what raw material exists, the filter decides how much of it survives. This is the block that gives subtractive synthesis its name: rather than building a tone up from nothing, it starts with something harmonically busy and carves pieces away.

- **Low-pass filter (LPF)** — lets the low end through and rolls off the top. The most common filter type in synthesis: turning it down makes a bright sawtooth progressively darker, warmer and more muffled.
- **High-pass filter (HPF)** — the opposite: lets the top through and rolls off the low end. Thins a sound out and strips the mud from underneath it.
- **Band-pass filter (BPF)** — keeps a narrow window in the middle and cuts both above and below it. The result is nasal and telephone-like, useful for isolating one part of a sound's character.

Two further controls decide exactly how a filter behaves at the point where it starts working:

- **Cutoff frequency** — the point in the spectrum where the filter begins to act. Sweep it slowly upward on a low-pass filter and a muffled tone opens out into a bright one; this single movement, more than any other, is the sound of a synthesiser being played expressively rather than simply triggered.

into the squelchy, singing quality familiar from acid basslines. Pushed further still, most filters will self-oscillate: they generate a pure tone of their own, becoming, briefly, a second oscillator.

Block Three: The Amplifier and ADSR

An oscillator on its own never stops — it just drones. The amplifier is what turns that drone into a note with a beginning, a middle and an end, and its usual controller is an envelope with four stages: attack, decay, sustain and release, or ADSR.

- **Attack** — how quickly the sound reaches full volume once a note starts. Set close to zero, the result is percussive: a piano, a pluck, a drum hit. Stretched out, the note swells in rather than striking — a string pad fading up instead of switching on.
- **Decay** — how quickly the sound falls from that peak down to its held level. A short decay gives a note a sharp, accented front edge even before sustain takes over.
- **Sustain** — not a time at all, but a level: how loud the note stays while a key is held down. High sustain behaves like an organ, holding steady for as long as you press. Sustain set near zero behaves like a piano or a plucked string, decaying away to nothing even while the key is still held.
- **Release** — how quickly the sound fades once the key is lifted. A short release stops a note dead; a long one lets it ring on, like the natural decay of a room.

The same four letters, arranged differently, are why a synthesiser can imitate a piano's quick strike and slow fade in one patch, then an organ's instant-on, instant-off block chord in the next, without changing a single oscillator or filter setting.

The Fourth Block Isn't a Block

Source, filter and amplifier are places sound physically passes through. Control is different: it is not a fourth stop on the signal path, it is a layer of movement reaching into all three of the others at once, and it is what turns a static patch into something a musician can actually play.

- **The LFO** (low-frequency oscillator) is an oscillator too slow to hear as a pitch — typically below about 20 Hz — used instead as a wobble applied somewhere else. Aimed at pitch, it produces vibrato. Aimed at amplitude, it produces tremolo. Aimed at the filter's cutoff, it produces the rhythmic opening-and-closing wah, or wobble, heard on everything from a slow analogue pad to a modern bass drop.
- **A second envelope**, routed to the filter rather than the amplifier, is how a synthesiser gets a bright attack that settles into a duller sustained tone — or the reverse, a sound that starts dark and blooms open. The filter envelope and the amplitude envelope run independently, which is why a plucked bass can sound bright for an instant and dark for the rest of the note.
- **Velocity** — how hard a key is struck — is routed to loudness by default on almost every synthesiser, and very often to brightness too: hit a key harder and a velocity-sensitive patch does not just play louder, it opens its filter and sounds sharper, the way a real instrument responds when it is played harder.

— a physical lever pushed during a note, most commonly wired to vibrato depth, so expression can be added or withdrawn while a note is still sounding.

- **Glide (portamento)** lets one note slide into the next instead of jumping cleanly between pitches — a controlled smear between two notes rather than a clean cut.
- **Monophonic and polyphonic** describe how many notes a synthesiser can sound at once. A monophonic instrument plays one note at a time, however many keys are held — every era before Part Three of this book is built on that limitation. A polyphonic instrument can sound a full chord, each note running its own copy of the source–filter–amplifier chain simultaneously.

THE ONE FACT

A patch with no control layer at all is not silent, but it is dead: the same note, at the same volume, with the same tone, however hard or long you play it. Everything a synthesiser gets praised for — expression, movement, feel — lives in this fourth layer, not in the three blocks underneath it.

What the Specification Asks For

Source: Pearson Edexcel Level 3 Advanced GCE in Music Technology (9MT0), Component 4, Section 1.3: Synthesis.

OSCILLATORS & WAVEFORMS	Sine, square, sawtooth, triangle; pitch control by tuning; noise as a fifth, pitchless source.
FILTERS	Low-pass, high-pass and band-pass; cutoff frequency; resonance (Q).
ENVELOPES	ADSR — attack, decay, sustain, release — controlling amplitude, and very often filter cutoff too.
LFO	A modulation source; targets pitch, filter cutoff and amplitude; rate and depth as its two controls.

Keep these four blocks in mind as you read on. Every machine in this book — the ones with no memory at all, and the ones that hold ten thousand presets — is answering the same question with them: what does the source sound like, what does the filter do to it, how does the amplifier let it live and die, and whose hand is on the controls while it happens.

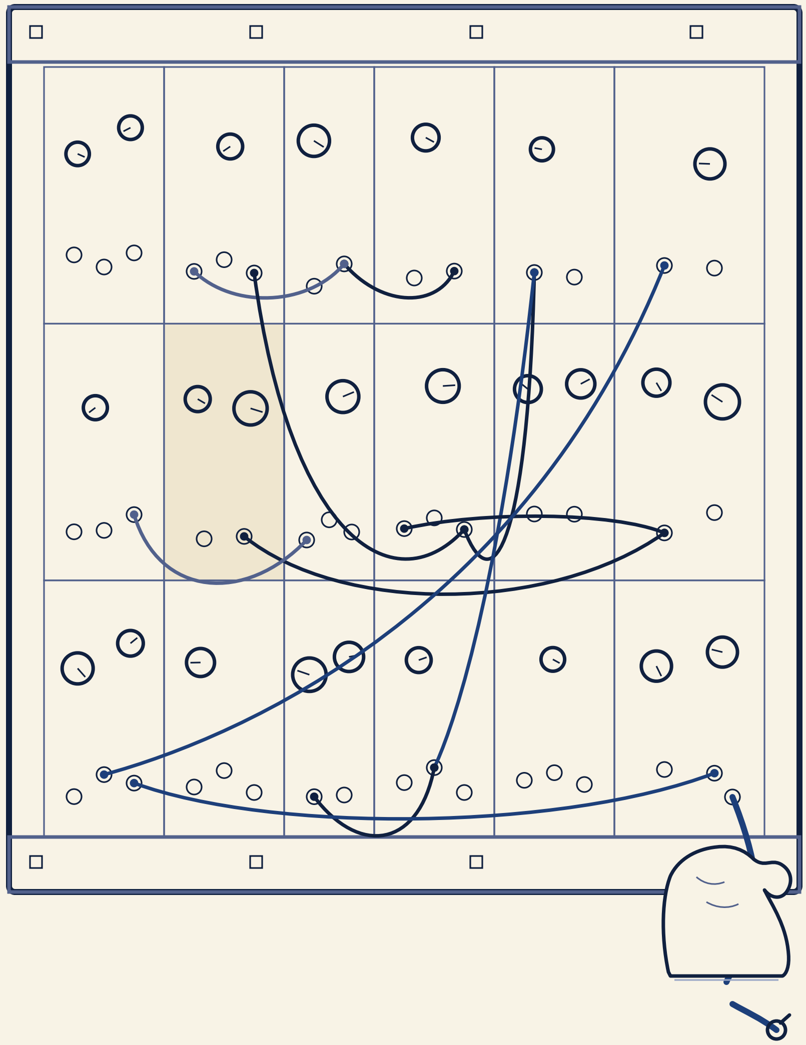

PART TWO · 1964–1977

The Patch-Cord Years

One voice, patched from scratch. Nothing can be saved.

CHAPTER THREE

One Voice, From Scratch

A theremin makes one pitch move under a hovering hand. An Ondes Martenot adds a ribbon and ranks of stops, all wired permanently into its own cabinet. Neither instrument could become a different instrument; the circuit inside each one was its whole and final idea. In the mid-1960s a different principle took hold, and it changed what "synthesiser" would mean from then on: treat pitch, tone and loudness as separate electrical problems, govern each one with a voltage, and let any module drive any other. Voltage control did not build one new instrument. It built a way of building instruments — one voice at a time, wired from nothing, and undone again by the end of the session.

Two Coasts, One Argument

The idea had two authors working at almost the same time on opposite sides of a continent, and they answered a single question in opposite ways: what does a musician actually do with their hands?

Robert Moog had been building and selling theremin kits by mail order since he was a teenager in the 1950s — the very instrument this book opened with — when composer Herbert Deutsch came to him in 1964 asking for something new: a module that could be played expressively, not just tuned. Moog answered with a piano-style keyboard wired to a control voltage of exactly one volt per octave, so that a given voltage always meant the same pitch interval on any module built to the same standard. He demonstrated the system at the Audio Engineering Society's convention that same year. It was a synthesiser designed to feel, immediately, like an instrument a working musician already knew how to hold.

Don Buchla, building his own modular system in California at almost the same moment, reached the opposite conclusion. Commissioned by composers at the San Francisco Tape Music Center to build them a flexible sound-generating instrument, he left the piano keyboard out on purpose. He believed it dragged eighty-eight fixed pitches and a hammer's percussive attack into a machine that owed the piano nothing, and he replaced it with touch-sensitive plates and a built-in sequencer — a way of stepping through a pattern of voltages rather than striking a key. Buchla's instruments made sounds a keyboard could never trigger, and asked the player a different question: not which note, but what shape of voltage.

MOOG · NEW YORK	Keyboard, one-volt-per-octave standard — built for a musician already trained on a keyboard.

BUCHLA · CALIFORNIA	No keyboard, touch plates and a built-in sequencer — built to escape everything a keyboard assumes.

A NAME FOR THE SPLIT — EAST COAST, WEST COAST

Writers later shorthanded this disagreement geographically: "East Coast" synthesis for Moog's keyboard-and-filter lineage, "West Coast" for Buchla's sequencer-and-touch-plate lineage. Neither term was printed on a manual in the 1960s — they arrived once the two approaches were far enough apart to need naming — but both are still in everyday use among manufacturers and reviewers, so both are worth knowing.

Neither position won outright, and it was never really a contest to be won: each builder had made a case for what a synthesiser was for, and the instrument itself carried that case in its wiring. The argument they started is still visible today, half a century on, in every modular case built since.

The Cable Is the Score

A modular synthesiser has no factory sound. Every cable moved from one socket to another rewires the instrument from nothing: an oscillator's raw wave routed into a filter, the filter's output routed into an amplifier, a hand-set envelope or a slow oscillator patched in wherever the player wants movement. The Source, Filter and Amplifier blocks from earlier in this book — with Control routed across all three — began life exactly this way: a literal chain of cables, laid out across a wall of sockets, different every time a player sat down.

And at the end of the session, all of it came apart. A patch existed only as long as the cables stayed plugged in. Nobody could save a sound to a disk, because there was no disk to save it to, and no memory inside the machine to hold a setting once the power went off. If a studio needed the same sound again the following week, the only record was whatever the player had written down — a hand-drawn patch sheet, or nothing at all — and the sound was rebuilt from memory, or simply lost.

THE ONE FACT THAT EXPLAINS THE ERA'S SOUND

Nothing beyond a scratch pad of pencil marks recorded a synthesiser's setting until the mid-1970s. Every session on one of these machines built its sound again from silence — which is why "patched from scratch" is not a figure of speech for this decade. It is exactly what happened, every time.

Monophony was not a stylistic choice either. Most modular systems of the period could sound only one pitch at once, because each voice needed its own dedicated oscillator, filter and amplifier, and a second voice meant patching an entire second rig alongside the first. A synthesiser part on a record from this never a chord played by one instrument.

Proof the Machine Could Play

For most listeners at the start of this decade, a synthesiser meant a sound effect: the theremin's science-fiction wobble, a whoosh or a bleep dropped into incidental music, a novelty rather than an instrument capable of a serious performance. One record changed that argument almost overnight.

In 1968 the composer and engineer Wendy Carlos, working with producer Rachel Elkind, released *Switched-On Bach*: a set of J. S. Bach's keyboard works, realised entirely on a Moog modular synthesiser, its single voice recorded and layered part by part until the album's dense keyboard textures existed as a finished, replayable performance. The engineering brought out warmth, articulation and dynamic shading — a performance rather than a demonstration — and the mainstream classical audience noticed. *Switched-On Bach* won three Grammy Awards the following year and became one of the best-selling classical albums ever made. A machine built from voltage and cable had made its case in the terms that mattered most: not as a sound effect, but as an instrument capable of music that critics judged on musical grounds, against Bach.

The Argument, Hard-Wired

Modular systems had solved the tone but not the price, the size, or the studio-bound patience they demanded. A player needed a piece of furniture-sized instrument, a spare afternoon and a memory for cable routes just to make a sound. In 1970 Moog answered that problem directly with the Minimoog Model D: a synthesiser with no patch cables at all.

Its signal path was fixed at the factory rather than wired fresh every time — three oscillators feeding a mixer, into a filter, into an amplifier stage, with two built-in envelopes shaping the filter and the loudness over time. It was the same Source, Filter, Amplifier and Control argument from earlier in this book, hard-wired into a single cabinet a musician could carry to a gig and switch on. Every control that mattered sat on the front panel, visible all at once. Nothing needed patching, because the one patch that mattered had already been chosen, permanently, by the people who built it.

The trade was real: a Minimoog could not become a different instrument the way a modular system could, one cable at a time. But it could travel in a car boot, and it could be played by a musician with no engineering background at all — and its particular filter sound, warm and slightly aggressive when driven hard, became one of the most recognisable, and most copied, tones in the whole history of the instrument.

Live and Unstable

Modular synthesisers were built for a studio bench, not a stage: heavy, liable to drift out of tune as their circuits warmed under stage lighting, and entirely dependent on whichever cable happened to still be plugged in. Keith Emerson of Emerson, Lake & Palmer toured with one anyway — a full Moog modular to gamble a live show on an instrument that had never been designed to survive one.

The BBC Radiophonic Workshop, whose engineers had built the electronic sound of *Doctor Who* out of tape loops and test oscillators through the previous decade — the “not a synthesiser” instruments this book opened with — spent the early 1970s finally acquiring the real thing: EMS synthesisers, patched and played by composers including Paddy Kingsland and Peter Howell. A decade after Delia Derbyshire had built a theme tune from borrowed oscillators and tape, the Workshop’s engineers were routing sound through the instrument voltage control had actually built.

CHAPTER FOUR

Hearing the Patch-Cord Years

This is your listening toolkit: place an unfamiliar synth sound by the technology you can hear.

Every constraint in Part Two's story leaves a fingerprint in the audio, and the fingerprint is remarkably consistent from a home-built EMS pin-matrix in London to a wall of Moog modules in a Los Angeles session room. One voice, patched from scratch, played by a pair of hands in real time: that sentence describes both the instrument and, if you listen for it, exactly what the recording sounds like.

What to Listen For

- Monophonic lines — one note at a time; a single, unmultitracked synth of this era cannot play a chord, so anything that sounds like block harmony is either an overdub or not a synthesiser at all
- Glide between notes — portamento slides the pitch smoothly from one note to the next instead of jumping, sometimes taking a noticeable fraction of a second to arrive
- Filter sweeps you can hear opening — the cutoff moving in real time, the tone brightening or darkening as it rises or falls, often by a hand on a knob rather than a fixed, repeating pattern
- A resonant peak riding the sweep — a vocal, sometimes whistling emphasis right at the cutoff, occasionally pushed so far that the filter self-oscillates into a pure, tuned tone of its own
- Detuned oscillators — two voices a few cents apart, beating against each other for a thicker, chorused unison rather than one clean pitch
- Hard-synced oscillators, where used — a searing, ripping edge as one oscillator's cycle is forcibly restarted by another, adding harmonics no single waveform contains on its own
- Hand-played expression — vibrato and bends driven live by a wheel under the player's hand, never quite the same depth or speed twice, because nothing about the performance was stored

Writing About What You Hear

A strong answer is a chain: technique → sonic outcome → judgement.

MODEL SENTENCE

"Herbie Hancock played the opening hook of Chameleon (Head Hunters, 1973) on an ARP Odyssey, not a bass guitar, working the line as a single held pitch at a time and using portamento to glide smoothly from note to note in a way no plucked string can duplicate. There is no pick attack and no string decay — only a continuous tone that holds at a fixed level for exactly as long as a finger stays on the key, and cuts off cleanly the moment it lifts. The synthesiser trades the natural transient of a struck string for something a bass guitar cannot do at all: an unbroken, mechanically smooth slide between pitches."

That delivers: named technique (portamento on a single-line ARP Odyssey part) · named record (Chameleon, 1973) · sonic outcome (continuous glide, no transient) · AO4 judgement (a trade, not a free upgrade).

A STRONGER JUDGEMENT

A synthesiser that can only sound one note at a time forces the player back into an instrumentalist's discipline: every pitch is chosen and shaped live, with nothing programmed to fall back on. When patch memory and polyphony arrive in Part Three, that discipline loosens — a preset can be recalled instead of re-patched, and a chord no longer has to be built one line at a time. A confident essay can argue that something real was lost along with the constraint: the physical urgency of a Moog or ARP line that is being played, in real time, rather than selected.

A LISTENING TRAP

Do not assume that a busy, chordal-sounding synth passage rules out a monophonic instrument. Wendy Carlos built the interweaving lines of *Switched-On Bach* one note at a time on a Moog modular, overdubbing each voice separately because the machine itself could never have played more than one at once. What sounds like a keyboard part with both hands playing chords may be a single hand, a single voice, and a great many passes of the tape.

Minimoog Bass or Bass Guitar?

Both instruments can hold a low, sustained note, so the giveaway is rarely the pitch — it is the shape. A bass guitar note starts with a pick or finger attack and then decays, however slowly, even under a held sustain; a synthesiser note sits at a fixed level for exactly as long as the key is down and stops on command when it is released. Listen for glide: a bass player's slide between frets is audibly imprecise and rarely spans more than a fret or two, while portamento on a Minimoog or ARP is mechanically smooth and can cover an octave or more without a single step in between. And listen for texture — string noise, fret buzz and pick attack are acoustic events a filter and envelope cannot produce, while a sweeping cutoff guitar's body and strings ever could, which is why a synth bass can sound simultaneously heavier in the low end and cleaner on top than the instrument it is standing in for.

Instruments and Records to Know

SYNTHESISERS	Moog modular (1964–, patch cords, no memory) · Minimoog Model D (1970, hard-wired monophonic) · ARP 2600 (semi-modular, built-in patch points) · ARP Odyssey (compact, no patch cords) · EMS VCS3 / Synthi A (UK, pin-matrix patching)
FILTER	Moog ladder filter, 24 dB/octave low-pass — a warm, musical resonance that self-oscillates into a pure tone at its most extreme setting
NOTABLE PLAYERS	Wendy Carlos · Keith Emerson · Herbie Hancock · Stevie Wonder (with TONTO programmers Malcolm Cecil and Robert Margouleff) · Kraftwerk
RECORDS TO KNOW	Wendy Carlos, *Switched-On Bach* (1968) · Emerson, Lake & Palmer, *Tarkus* (1971) · Stevie Wonder, *Music of My Mind* (1972, with TONTO) · Herbie Hancock, *Head Hunters* (1973) · Kraftwerk, *Autobahn* (1974)

What the Spec Asks

PITCH CONTROL: FREQUENCY MODULATION AND TUNING	Portamento glides continuously between notes; a pitch wheel bends a held note up or down without triggering a new one
FILTER CUTOFF FREQUENCY	An audible sweep in real time — the tone opens from a muffled hum towards a bright buzz, or closes the other way
RESONANCE (Q)	A vocal, whistling emphasis riding the sweep; at its extreme the filter itself rings out a pure sine tone
ADSR: ATTACK, DECAY, SUSTAIN, RELEASE	A fast filter-envelope attack gives a plucky bite before the tone settles to its sustained level and cuts off cleanly on release
LFO TARGETS: PITCH, FILTER CUTOFF, AMPLITUDE	Vibrato driven as often by a hand on the pitch wheel as by an LFO — the player, not a fixed rate, decides the depth

BASIC WAVEFORMS: SINE, SQUARE, SAWTOOTH, TRIANGLE	One bright sawtooth or square oscillator, or two detuned together for a thicker, beating unison

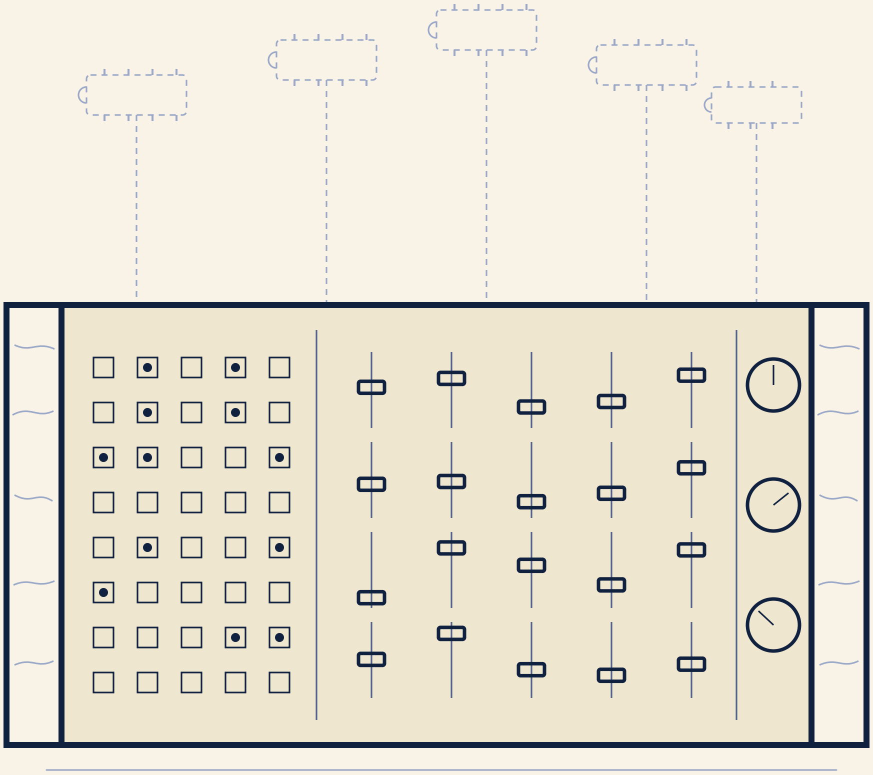

PART THREE · 1978–1995

Presets and Chips

Sound becomes storable — then becomes numbers.

CHAPTER FIVE

Patch Memory and the Peace Treaty

For the first fourteen years of this story, a synthesiser could hold one sound at a time, played by one finger at a time, remembered by nobody at all. Turn the machine off, and the patch went with it. Even once engineers found a way to make several notes sound together, that memory problem did not go away — it multiplied.

What Changed

This part of the story turns on two separate inventions arriving within six years of each other: a way to save a sound, and a way for one manufacturer's saved sound to mean anything to another manufacturer's machine at all. Between the two sat an instrument built to sound nothing like its analogue neighbours — and it became one of the best-selling synthesisers ever made, largely because almost none of its owners ever learned to programme it themselves.

- Yamaha's CS-80 (1977) proved a synthesiser could play more than one note at a time and still respond to the weight of a finger — but recalling a sound you had already built, quickly, at the press of one button, was still nobody's standard.
- Sequential Circuits' Prophet-5 (1978) solved that: a microprocessor read every knob's position and stored the result as numbers, recallable in an instant.
- MIDI (1983) let a Prophet 600 talk to a Roland JP-6 for the first time — a shared language that no single manufacturer owned.
- Yamaha's DX7, released the same year, generated its entire sound from numbers alone, with no voltage-controlled circuit anywhere inside it — and sold in quantities no synthesiser had reached before.

THE ONE IDEA THAT EXPLAINS THE ERA

Every invention in this chapter answers the same question from a different angle: how do you keep a sound once you have found it? Patch memory kept it inside one machine. MIDI let it travel between machines. FM synthesis stored it as numbers from the very start — and made the numbers so fiddly to edit that most owners simply kept the factory's instead of writing their own.

The Era’s Signature Moves

Genuine polyphony arrived first without solving the memory problem at all, because it arrived without much memory anywhere in the instrument. Oberheim's Four Voice (1975) reached four-note polyphony by the bluntest available method: four complete, independent monophonic synthesiser modules, each with its own oscillators, filter and envelope, wired to a single shared keyboard. Set every module identically and the result was one coherent four-note instrument; leave one filter knob slightly off from the other three and the fourth note in any chord quietly sat apart from the rest. There was no way to store the setting a player had found — only four panels of knobs to remember by hand, or leave exactly where they were and never touch the instrument again.

Yamaha's CS-80 (1977) folded that same problem into a single cabinet: eight true voices, each sensitive to how hard a key was struck, plus a ribbon controller under the keyboard's left hand for pitch bends drawn by finger rather than a wheel. It became the expressive standard-bearer for polyphonic analogue synthesis, prized for exactly the touch-sensitive control that the Four Voice's four separate modules could never coordinate. Recalling a sound already built, quickly and by pressing a single button, remained beyond it.

The Prophet-5 — sound becomes storable

Dave Smith's Sequential Circuits solved the recall problem a year later, and named the result plainly: the Prophet-5 (1978), five-voice polyphonic, with patch memory built in. Underneath the front panel it looked much like any analogue synthesiser — genuine voltage-controlled oscillators and filters gave each of its five voices the same continuously variable warmth a Minimoog player already knew. The difference sat between the knobs and the circuits: a microprocessor read every control's position and stored the readings as a set of numbers, up to forty patches deep. Press one button, and the whole instrument silently reconfigured itself — the knobs stayed exactly where they were, but the sound underneath them changed to whatever had last been saved in that memory slot.

This is the invention this part of the book is named for. A sound had never before been a fact a player could file away and retrieve intact; it had been a physical arrangement they either kept their hands on or lost. The Prophet-5 made a patch a piece of storable information for the first time, and the approach spread quickly — Roland's Jupiter-8 (1981) carried the same digitally-controlled-analogue idea into eight-voice polyphony within a few years, with its own bank of recallable sounds.

MIDI — the peace treaty between rival machines

The Prophet-5 solved memory inside one machine. It said nothing about getting a sound, or even a single note, from one manufacturer's instrument into another's. Robert Moog's one-volt-per-octave standard, adopted by Roland, ARP, Oberheim and Sequential Circuits, meant that machines from those four companies could at least trigger each other's oscillators. Yamaha and Korg used incompatible schemes of their own. A studio built from synthesisers by different manufacturers was, electrically, a room of instruments speaking past one another.

Ikutaro Kakehashi of Roland, the idea became MIDI, the Musical Instrument Digital Interface. Its first public demonstration, at the NAMM trade show in January 1983, connected a Sequential Prophet 600 to a Roland JP-6 down a five-pin cable, and had one instrument's keyboard play the other's voices. What made the moment remarkable was not the cable but the licence: Smith and Kakehashi gave MIDI away, royalty-free, to every competitor in the room. No single company owned the standard that let every other company's instruments understand each other. Within a couple of years, a MIDI socket was standard equipment on synthesisers from every serious manufacturer — a genuine peace treaty, signed by companies that spent the rest of the decade competing as hard as ever on everything else.

The DX7 — sound as pure numbers

Yamaha's DX7 (1983) took the idea of a number-based sound to its logical end: an instrument with no voltage-controlled circuit anywhere inside it. Its method, FM (frequency modulation) synthesis, had been discovered by accident, when Stanford researcher John Chowning pushed an oscillator's vibrato far beyond any musical depth in 1967 and heard entirely new tones bloom around the note rather than a wobble. Stanford licensed the discovery to Yamaha in the mid-1970s. Inside the DX7, six digital ‘operators’ — pure tone generators, with no wave to shape because there was no voltage to shape it with — could be wired into thirty-two different routings, called algorithms, in which one operator's pitch modulates another's at audio speed. The results were bright, glassy, bell-like and metallic in ways no analogue filter could reach, because nothing was being subtracted from a rich starting wave: the timbre was built directly out of the modulation itself.

The DX7 sold in numbers no synthesiser before it had reached. It also became a byword for an instrument almost nobody could programme. Its front panel carried a small LCD screen, a single data-entry slider and rows of membrane buttons — no dedicated knob per parameter, the opposite choice from the Prophet-5's one-control-per-function layout. Editing a single one of its six operators meant stepping through menus with that one slider; editing all six, with the routing between them, defeated most owners entirely. They kept the factory patches instead. One in particular, a bright, bell-like electric piano sound named E. PIANO 1, ended up on more records through the 1980s than almost any other single synthesiser sound — a preset nobody in the room had written themselves.

Analogue versus digital — the era’s real argument

Two machines released in the same year, 1983, embody the argument this part of the book is really about. A Prophet-5-style polysynth was still, underneath its memory chips, a rack of continuously variable voltage — warm, capable of drifting slightly as its analogue components warmed up, and controlled by a panel where one knob still meant one function. The DX7 was precise in a way no analogue circuit had ever managed, immune to that drift, and controlled by an interface that put its entire sound two menu-levels deep in numbers nobody outside Yamaha's engineering team had written.

The argument was never really about which one sounded better. It was about what a sound now was, and who got to touch it. Patch memory had already turned a sound into something a player could keep. Digital synthesis went one step further: it turned a sound into something nobody had to build at all, because the many hit records of the decade were made from sounds their own players had never opened.

CHAPTER SIX

Hearing Presets and Chips

Train your ear to place an unfamiliar track by the synth technology you can hear.

Two sound-worlds share the airwaves through this era, often on the very same record, and they do not sound remotely alike. One is warm, chorused, and drifts by a hair, because a real voltage is still doing the work behind the panel. The other is glassy, perfectly stable, and was built entirely from numbers. Sound becomes storable, then becomes numbers — and by the end of this chapter you should be able to tell which one you are hearing within a few seconds of a record starting.

What to Listen For

- Chorused, wide analogue pads and brass — several detuned oscillators run through an onboard chorus circuit, producing a warm ensemble sound that seems to shimmer even on a single held chord
- Bright, syncopated poly-brass hits — a buzzy sawtooth-based stab with a short filter sweep audible on its leading edge, thickened by more than one voice sounding at once
- Glassy FM electric piano — a hard, bell-struck attack followed by a high shimmer that decays faster than the low end underneath it; brighter and thinner than a real Rhodes or Wurlitzer, with none of a tine's mechanical thump
- FM bells and metallic pads — pure, clangorous, faintly inharmonic tones a filtered sawtooth could never reach on its own, however far its resonance is pushed
- FM slap bass — a percussive, plucked low note with a bright transient "pop" and a quick decay, borrowing the attack of a real slap bass without any string noise or fret buzz underneath it
- Velocity-sensitive brightness — hit the key harder and an FM patch does not just get louder, it gets audibly brighter and more metallic, with no filter anywhere in its signal path to account for the change
- Perfectly stable tuning — no drift, no thermal wobble between notes: a plain giveaway of a digital oscillator rather than a voltage-controlled one
- The same patch, twice — because a sound can now be saved and recalled, one factory preset can turn up, essentially unchanged, on records that otherwise share nothing at all

In a listening exam this split is usually the fastest identification available to you. If a pad drifts, breathes and slowly detunes against itself as it holds, a real voltage is doing the work. If it is glassy, static under sustain, and faintly metallic in a way no amount of filtering explains, a digital operator is.

Presets You Should Know by Ear

attack is hard and percussive, almost bell-struck, with a bright high partial that fades away noticeably quicker than the low end holding underneath it — the opposite decay shape from a real electric piano, whose bass tines actually ring longer than their treble. That mismatch is exactly why it reads as an electric piano to a casual ear and as something else entirely once you listen for it: no wood, no felt, no mechanical thump, just six sine-wave operators giving a very convincing impression. Whitney Houston's *Saving All My Love for You* (1985) is as good a place as any to hear it sitting under a vocal, on a record that shares nothing else with the DX7's other famous outings.

The same engine gave the era its other two signature timbres. FM bells push a modulator far enough away from its carrier that the resulting overtones sit outside any ordinary harmonic relationship altogether — a clangorous, glassy edge no filtered sawtooth can reach, whatever its resonance is set to. FM slap bass borrows the percussive pluck of a real bass guitar's attack, then decays too cleanly and too quickly to be mistaken for wood, metal string or finger: there is a pop, and then very nearly nothing.

Programming an operator algorithm from scratch was so unrewarding that a market grew up around avoiding it entirely: third-party patch libraries, sold on cassette or floppy disk, let a player load someone else's carefully tuned sound without ever opening a menu themselves. That is the real reason a single factory preset could end up defining a decade of records — most owners never wrote a patch of their own. They bought one, or simply kept whatever the factory had already loaded in.

HEAR IT HAPPEN

Toto's *Africa* (1982) opens with both of this chapter's sound-worlds standing side by side. The plucked, marimba-like pattern under the intro is FM — a Yamaha GS-1 factory preset, an early commercial FM instrument that predates the DX7 — while the sustained pad answering it is a Yamaha CS-80: fully analogue, chorused, slow to swell into place. One instrument sounds instantly stable and precise. The other visibly breathes, even standing still on a single note.

Toto, Africa — Toto IV, 1982.

That chorus is not incidental, and it is worth knowing where it comes from. Machines like the Roland Juno-60 (1982) and the CS-80 route each voice through a chorus circuit before it ever reaches the output socket, splitting the signal into two slightly delayed, slightly detuned copies and spreading them across the stereo field. The result is a pad that seems to move even though nothing on the keyboard is changing — width and shimmer built into the instrument itself, years before a mix engineer would reach for a chorus pedal to fake the same trick after the fact. Listen on headphones and the two copies audibly separate, one drifting a fraction sharp of the other; that drift is the whole effect.

Instruments and Records to Know

ANALOGUE POLYSYNTHS	Yamaha CS-80 (1977) · Sequential Circuits Prophet-5 (1978) · Roland Jupiter-8 (1981) · Roland Juno-60 (1982, built-in chorus)
FM & DIGITAL SYNTHS	Yamaha GS-1 (early FM, pre-dates the DX7) · Yamaha DX7 (1983, six-operator FM) · New England Digital Synclavier (FM, high end of the market)
SIGNATURE PRESET	Yamaha DX7 factory patch E. PIANO 1 — heard, largely unedited, across hundreds of otherwise unrelated 1980s records
RECORDS TO KNOW	Toto, *Africa* (1982) · Vangelis, *Blade Runner* (1982, film score) · Prince, *1999* (1982) · Whitney Houston, *Saving All My Love for You* (1985)

What the Spec Asks

ANALOGUE PADS & BRASS	Basic waveforms (sawtooth, square) shaped by a low-pass filter — cutoff frequency and resonance — for a darker or brighter tone.
FM ELECTRIC PIANO, BELLS, BASS	FM synthesis: frequency modulation using multiple oscillators; carrier and modulator; operator algorithms and their routing configurations.
VELOCITY-SENSITIVE BRIGHTNESS	How different synthesis parameters affect timbre and character.
CHORUSED WIDTH	Oscillator fine tuning — subtle detuning between oscillators, producing a rich, thick chorus effect.

PART FOUR · 1996–2009

Software and the Infinite Instance

The instrument becomes a file: free to duplicate, endless by default.

CHAPTER SEVEN

The Infinite Instance

In 1978, a Prophet-5 held five voices, one set of knobs, and no way to be in two places at once. There was exactly one of it in whichever room it happened to sit. By the end of this era, a musician could open a synthesiser, open nine more copies of the very same one, undo every decision back to a blank project, and never touch a soldering iron or a patch cable. The instrument had stopped being a thing you owned. It had become a file you ran.

NOTE

Section 1.3 tests the categories a synthesiser is built from — oscillator, filter, envelope, modulation — not the model numbers or years attached to them. Nothing below is required knowledge for the exam. But VST, Reaktor and Serum are very likely tools already installed on your own laptop, and knowing where they came from is what turns “I used a soft synth” into an answer that names what actually changed, and why it mattered.

What Changed

None of what follows needed a new circuit to be invented. Three separate shifts, arriving within a decade of each other, did the work instead:

- A keyboard learned to hold recorded sound and treat it as raw material for the same filter-envelope-modulation chain every synthesiser in this book has used since Part Two — sampling and synthesis met inside one affordable box.
- A plug-in standard let a synthesiser run as software inside an ordinary computer, so owning one stopped meaning owning a hardware object with a serial number.
- Once an instrument was a file, it could be copied, layered and undone without limit — which changed not only how a sound was made, but how a piece of music could be assembled in the first place.

THE ONE IDEA THAT EXPLAINS THE ERA

Every earlier era in this book was defined by a scarcity: one voice, one patch, one hardware unit passed between studios. Era 4 did not solve that scarcity — it removed it. A musician with a laptop had access to more polyphony, more running instances and more undo history than the best-funded studio of the previous decade owned outright, and had to develop new judgement about when to stop, now that nothing else was stopping them.

The Era's Signature Moves

The Korg M1 (1988) — sampling and synthesis meet in one keyboard

Korg's M1, launched in 1988, put a bank of recorded instrument sounds — piano, strings, brass, choir, plucked bass — onto ROM chips inside a keyboard, then ran every one of them through exactly the same filter, envelope and modulation stages any subtractive synthesiser in this book already uses. Musicians and manufacturers soon had a name for that hybrid: a ROMpler, a synthesiser whose oscillator is a fixed recorded sample rather than a waveform generated from scratch. Add an onboard sequencer and a set of built-in effects, and the M1 needed nothing else around it — the first true workstation synthesiser, and among the best-selling synthesisers ever made. One of its factory patches, a lush pad named Universe, turned up on enough late-1980s and early-1990s dance records that it briefly became a genre's default sound, long before most listeners had any idea a single preset was responsible.

THE BOUNDARY THIS BOOK WON'T CROSS

The M1's hybrid sits right on a line this book stops at deliberately. A true sampler lets a musician record or load any sound at all; the M1's samples were fixed at the factory, a menu rather than an open microphone. The fuller story of sampling — the machines that could capture anything, and what that did to the question of who owns a sound — belongs to a different volume on this shelf. What matters here is narrower: the M1 proved that a recorded sample and a synthesised waveform could pass through the same filter and envelope and come out equally at home in a synthesiser's argument.

VST (1996) — the instrument becomes a file

Steinberg gave that argument its clearest expression eight years later. Virtual Studio Technology, launched alongside Cubase in 1996, was a plug-in standard: a specification any software developer could write to, so a program running inside a computer could generate or process audio exactly as a piece of outboard hardware would. VST covered audio effects first; support for whole virtual instruments followed within the same format before the decade was out. The consequence reached further than the studio it started in. A synthesiser no longer had to be an object with a serial number, shipped in a box, sitting in a deleted without leaving a gap on a shelf. Every soft synth built since, including whichever one is bundled inside the DAW you already use, inherits that architecture.

Unlimited instances, unlimited undo

Every synthesiser earlier in this book existed as a fixed number of physical units, and every fixed number of units was itself a limit on what a piece of music could contain. A studio built around a Prophet-5 owned five voices, and if two songs booked into the same week both wanted synthesiser pads at once, someone waited their turn. A soft synth running inside a computer removed that arithmetic almost entirely. The same instrument could open as many times as the computer's processor allowed — a dozen tracks of the identical synthesiser stacked into one arrangement, each running its own patch, none of them borrowed from the others. Mistakes stopped being permanent, too: a patch programmed wrongly could be undone, step by step, back past the point where it went wrong, something no hardware synthesiser before this era had offered as a matter of course. What changed was not only the sound available to a composer, but the working method itself — music built by trying things and discarding what failed, rather than committed to in advance because committing was the only option on offer.

A decade of software instruments — Reaktor, Massive, Serum

Once the format existed, the instruments built for it grew ambitious quickly. Native Instruments released Generator in 1996: not one synthesiser but a modular building kit on screen, in which a musician wired oscillators, filters and effects together much as a Patch-Cord-years musician once wired a Moog or a Buchla — except every patch could now be saved, renamed and reopened exactly as it was left. Generator was renamed Reaktor within a few years, the name every musician who has opened it since actually knows it by. Massive, released by Native Instruments in 2006, became so tightly bound to one sound that the two are still spoken of in the same breath: a wavetable synthesiser whose growling, wobbling low end supplied much of that decade's dubstep and early EDM bass, closely enough that producers talked about "the Massive sound" the way an earlier generation talked about the DX7's electric piano. Serum, released by Xfer Records in 2014, took wavetable synthesis and put the mechanism on screen: earlier wavetable instruments hid their tables behind menus and numbers, while Serum let a musician see, and draw, the actual waveform moving under a note as it played — wavetable synthesis with the lid taken off.

KORG M1 · 1988	ROMpler workstation: onboard PCM samples run through a subtractive filter and envelope, plus a sequencer and effects in the same box.
VST · 1996	Steinberg's plug-in standard. Effects first; virtual instruments followed within the same format by the end of the decade.
REAKTOR · 1996	Native Instruments' modular software environment, launched in 1996 under the name Generator and renamed Reaktor a few years later — build a synthesiser from blocks on screen, then save the patch.

MASSIVE · 2006	Native Instruments' wavetable synth, closely tied to the wobbling low end of the dubstep and early-EDM years.
SERUM · 2014	Xfer Records' wavetable synth, with the waveform itself shown and editable on screen in real time.

Presets for sale — what an infinite instrument sells you

An instrument that is a file makes its patches into files too. A sound designer can programme a single Serum patch once and sell it to ten thousand people who will never open its oscillator page for themselves; preset packs became a small industry inside the larger one, built on someone else's sound-design hours rather than a factory ROM chip. It is the DX7's preset culture, met earlier in this book, at a different scale entirely — not one factory-programmed chip shared by everyone who bought that model of synthesiser, but any number of near-identical patches shared by anyone who paid for the same download, on any synthesiser able to open the file. The complaint that followed was predictable, and not entirely wrong: records built from the same three downloaded patches, a sound bought rather than built. Whether that counts as a shortcut, or simply a new place to spend the hours a synthesiser used to demand of a musician, is exactly the kind of question Section 1.3 rewards you for being able to argue either way.

None of it required more electricity, more voltage, or a single new physical component. It required only that a computer could run software fast enough to sound like an instrument in real time — and once it could, the limiting factor was no longer money, or hardware, or a single machine passed between studios. It was only ever how many ideas a musician had time to try before the session ended.

CHAPTER EIGHT

Hearing Software

This is your listening toolkit: place an unfamiliar synth sound by the technology you can hear — even once that technology stops telling you anything about which machine, or even which decade, made it.

Every earlier chapter in this book has traded on a simple rule: constraint leaves fingerprints. A theremin can only glide. A Minimoog can only play one note at a time. A DX7 patch, once written, is a fixed set of numbers almost nobody but its author fully understood. Part Four breaks that rule on purpose. A single laptop, running a single piece of software, can be an unlimited number of different instruments, held back by nothing but disk space and imagination — so the fingerprint this era leaves behind is not which machine was used, but what it sounds like when nothing stops a machine from doing anything at all.

What to Listen For

- Supersaw stacks — not one bright sawtooth oscillator but seven or more, tuned a few cents apart from a shared root and spread across the stereo field, beating gently against one another into a shimmer no single-oscillator patch can produce
- Wavetable motion — the timbre itself changing shape inside a single held note, without any hand visibly turning a knob: a pluck that starts hollow and grows metallic, a pad that seems to breathe, because the oscillator is scanning through a table of waveshapes rather than sitting still on one
- Sidechain pumping — a pad or bass audibly ducking in volume in time with the kick drum, a rhythmic "breathing" that has nothing to do with which notes are being played and everything to do with a compressor keyed to the beat
- Perfect tuning and perfect timing — pitches landing exactly on the cent and notes starting exactly on the grid, with none of the drift an analogue oscillator or a human hand introduces; when a synth line sounds suspiciously flawless, the flawlessness is itself the clue
- Stacked, wide unison — dozens of simultaneous voices where an older analogue instrument could offer, at most, a handful; width and thickness that cost nothing extra to add once the instrument is a file rather than a circuit
- Presets travelling unedited — the same factory patch, untouched, turning up on records that share nothing else in common, because thousands of producers own an identical copy of the same plug-in

Writing About What You Hear

MODEL SENTENCE

"Darude's Sandstorm (2000) builds its entire lead sound from a supersaw — not one sawtooth wave, as a Minimoog player would use, but several tuned a few cents apart from a shared root and spread wide across the stereo field, so each voice beats gently against its neighbours. There is little of the earlier eras' audible hand on a filter knob and no player's touch to speak of: the thickness itself is the performance, engineered once and then triggered identically every time the pattern repeats. This is the shift the whole of Part Four turns on — where Chapter Four's monophonic lead earned its texture from what one instrument's circuitry allowed, Sandstorm's lead earns its texture from what an unlimited number of digital oscillators cost nothing to add."

That delivers: named technique (supersaw, detuned oscillator stacking) · named record (Sandstorm, 2000) · sonic outcome (engineered thickness, little audible performance) · AO4 judgement (constraint-earned texture versus cost-free abundance).

AN EVALUATIVE COUNTER-ARGUMENT (LIFTS AO4)

Perfect tuning and perfect timing are easy to praise as professionalism, but not every ear hears them as an improvement. A synth part quantised to the exact frame and tuned to the exact cent can sound flawless and faintly inert at the same time — a performance with all the small errors that usually signal a human being is present removed from it. Deadmau5's Strobe (2009) makes the opposite case from inside the same technology: its central pad shifts almost imperceptibly across the track's roughly ten-minute length without a single audible edit or repeat, and the same unlimited software resources that let Sandstorm stack detuned oscillators into a hard, bright edge are spent here on patience instead of density. A confident essay can hold both claims at once — unlimited digital resources can make a track sound machine-perfect and characterless, or more patient and more detailed than any two hands could manage in real time. The technology decides neither outcome by itself; the person using it does.

When the Instrument Could Be Anything

Every previous chapter in this book has let you work backwards from a sound to a machine: a pure sine tone with no attack meant a theremin, because nothing else in Part One could make that shape; a fixed, gliding monophonic line meant a Moog or an ARP, because nothing else in Part Two could play it. That trick stops working here. A supersaw lead, a moving wavetable pad and a sidechained bass can each be built inside dozens of unrelated pieces of software, on a laptop bought last year or one bought in 1999, by a bedroom producer or a platinum-selling one — and the audio alone will not tell you which. Naming a specific plug-in or manufacturer in an exam answer about an unfamiliar Part Four extract is a guess, not a description, and an examiner cannot credit a guess.

timbre is one static wave or is itself moving inside a single note; whether the tuning and timing sit exactly on the grid or drift; whether the volume is pumping in time with a kick drum that has nothing to do with the notes being played. Describe those observable behaviours, precisely and in the spec's own vocabulary of waveform, cutoff, resonance and modulation, and the description survives regardless of which software actually made the sound. That shift — from naming the machine to naming the behaviour — is the exam skill this chapter has been building towards.

Instruments and Records to Know

SOFTWARE SYNTHESISERS	Native Instruments Reaktor (released 1996 as Generator, renamed Reaktor) — modular software instrument · Native Instruments Massive (2006, the wobble decade's default bass and lead engine) · Xfer Serum (2014, wavetable synthesis built around watching the wave itself move)
THE STACK'S ORIGIN	Roland JP-8000 (1996, hardware) — its Super Saw oscillator, seven detuned sawtooth waves in a single voice, gave the whole era's stacked lead and pad sound its name
RECORDS TO KNOW	Darude, *Sandstorm* (2000) · deadmau5, *Strobe* (2009) · Avicii, *Levels* (2011)

What the Spec Asks

BASIC WAVEFORMS	Spec: "sine, square, sawtooth, triangle". A supersaw is still, underneath the stacking, sawtooth waves — the spec's own vocabulary names the ingredient; only the quantity has changed.
PITCH CONTROL	Spec: "Frequency modulation and tuning". A digital oscillator holds a tuning exactly, indefinitely — the perfect stability earlier chapters treat as a giveaway is this same heading, taken to its limit.
MODULATION SOURCES / LFO TARGETS	Spec: "Creating movement in synthesized sounds" · "Pitch, filter cutoff, amplitude". Wavetable motion sits just outside this list — rather than one wave modulated by an LFO, the oscillator itself is the thing changing, frame by frame, inside a single note.

CROSS-REFERENCE

Sidechain pumping is not a synthesis parameter at all — it belongs with compression and gating, not with oscillators and filters — but it remains one of this era's clearest audible fingerprints, whichever part of the specification it is filed under.

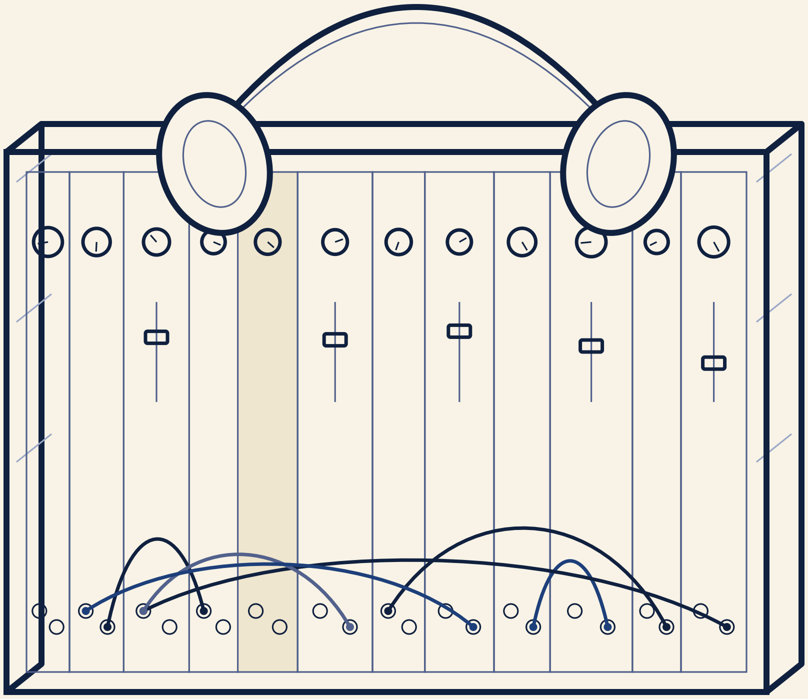

PART FIVE · 2010–NOW

The Modular Return

Musicians choose scarcity back on purpose.

CHAPTER NINE

Choosing Scarcity

By the middle of this decade, a musician could own a synthesiser with no memory at all — on purpose. No patch could be saved. No setting could be recalled. No mistake could be undone. Thirty years after the Prophet-5 made a sound something you could file away and thirty after the DX7 made it a string of numbers, a growing number of working musicians spent real money to give both of those inventions back. This chapter is about why.

The Format That Waited

The idea did not arrive new. In the mid-1990s, the German engineer Dieter Doepfer released a modular synthesiser format he called the A-100: a standardised rack size, a standardised power connector, and Moog's own one-volt-per-octave pitch standard underneath it all, so that a module from one small manufacturer would reliably talk to a module from another. Doepfer built it to put modular synthesis back within reach of working musicians at a fraction of the old cost, riding a renewed interest in analogue sound that was already building through the early 1990s. For its first fifteen years, Eurorack — the name the format eventually took, after the metric "3U" rack height it shares with professional broadcast and telecoms gear — stayed exactly that: a niche, serving builders and experimentalists who already knew what a patch cable did.

Then, from around 2010, the niche stopped being one. A wave of new manufacturers — small companies alongside hobbyists turned professional, releasing single modules rather than complete instruments — began building for the format in numbers Doepfer's original range had never approached. Cases that had been a specialist's furniture became something sold, assembled and demonstrated at ordinary music-shop counters and dedicated synthesiser expos. The format itself had barely changed since the 1990s. What had changed was who wanted it, and why.

THE ONE IDEA THAT EXPLAINS THIS CHAPTER

Every synthesiser since the Prophet-5 had competed on how much it could remember. Eurorack modules compete on the opposite ground: most hold no memory of any kind, on purpose, because their maker and their buyer have both decided that forgetting is part of the instrument.

Three described, forces a decision at every socket: this oscillator, filtered this way, controlled by this envelope, right now — and the decision cannot be half-remembered from a preset browser, because there is no browser to remember it. Musicians who choose Eurorack over a software synthesiser with infinite undo describe the same trade from several angles: a patch that cannot be saved has to be listened to and judged in the moment it exists, not filed away to compare later against forty other versions. A knob turned by hand, with nothing standing between the turn and the sound, responds instantly and continuously, rather than through a mouse click landing on one of a fixed set of values. And because nothing recalls a setting exactly, no two performances of the "same" patch are ever quite identical — small drifts and human error become part of the material rather than an accident to be corrected out.

The recall a Prophet-5 offered, and every polysynth after it perfected, is not being rejected as a bad idea. It is being set aside, deliberately, for particular music where the constraint itself is the point — the same logic, updated fifty years, behind Chapter Three's modular pioneers building a voice from nothing because nothing else existed yet. The difference this time is that plenty else exists. The scarcity is chosen.

Real Circuits, Lower Prices

A second force pushed in the same direction from below. Through the 2010s, budget manufacturers — the German company Behringer chief among them — began releasing low-cost hardware built around genuine voltage-controlled oscillators and filters, some of them close clones of specific classic circuits, others new designs built to the same analogue principles. A synthesiser that had cost a working musician a month's wages in the 1970s, or that survived only as an expensive vintage original by the 2000s, became available new, in a Eurorack module or a small standalone box, for a fraction of that price. Established manufacturers answered from the other direction: Moog itself returned the original Minimoog Model D to production in the mid-2010s, built to the original analogue circuit rather than modelled in software, alongside reissues of its earlier modular systems.

The effect reached students directly. A real voltage-controlled oscillator — not a digital model of one, an actual analogue circuit doing the same physical job the Minimoog's did in 1970 — sat on a school desk for the first time in numbers that mattered. Understanding the Source, Filter and Amplifier chain from earlier in this book stopped being a matter of reading a diagram or watching a demonstration video, and became a matter of turning a real cutoff knob and hearing a real filter open.

Two Rigs at Once

Almost nobody making electronic music today works in only one of this book's two most recent worlds. A DAW running on a laptop remains the fastest way to record, edit and arrange a piece of music, exactly as Part Four described; a Eurorack case or a hardware synthesiser sitting beside it remains the fastest way to find a sound nobody has heard from that laptop before, because no preset browser generated it. The two are routinely wired together rather than kept apart: a software sequencer or a clock-sync protocol such as Ableton Link keeps a modular rig's patterns locked in time with a DAW session running alongside it, and a hardware synthesiser's output is recorded straight into the same software timeline that hosts a dozen software instruments beside it.

them. The DAW supplies the parts of production that genuinely are better with unlimited undo — editing, arranging, mixing. The hardware supplies the parts that are arguably better without it — a first, unrepeatable encounter with a sound, made by hand, that then gets recorded before it can be recalled by anything except that recording.

Both Positions, Alive at Once

This part of the book is named for a choice, not a technology, and the choice sits at the centre of where the story of synthesis actually stands today. Software synthesis, described in Part Four, offers a genuinely infinite instrument: any number of copies running at once, every setting saved automatically, every mistake reversible with a single keystroke. Eurorack and its wider hardware revival offer close to the opposite of every one of those properties, at real financial cost, by choice. Both are commercially healthy. Both fill entire sections of any working studio's equipment list. Neither side of the argument this book opened with — what may a musician touch, save, copy or lose — has won. A century in, both answers are still being built, sold and played, often by the same musician in the same session.

Look closely at what a Eurorack case actually asks of the person patching it, and the shape underneath it is the one this book began with. A theremin gave a player open air and a hand held near an antenna, with nothing saved and nothing to fall back on but the player's own ear. A modular case gives a player a wall of sockets and a handful of cables, with nothing saved and nothing to fall back on but the player's own hands. A century of instruments that could be filed, copied, recalled and duplicated without limit has arrived back, deliberately, at an instrument that exists only for as long as someone is touching it.

CHAPTER CLOSE

Source: Pearson Edexcel Level 3 Advanced GCE in Music Technology (9MT0), Component 4, Section 1.3: Synthesis.

CHAPTER TEN

Hearing the Return

This is your listening toolkit: place an unfamiliar track by the technology you can hear.

Era Five is the first era in this book where a single record can genuinely have been made either way, and a listener has no reliable means of finding out which. A patched Eurorack case and a laptop running a careful software model of the same circuit can sit in the same session, sometimes feeding the same recorded track. Nothing below claims a foolproof test for telling hardware from software by ear — that test does not exist. What follows is vocabulary: precise enough that which side of the hardware line actually made the sound stops mattering to your mark.

What to Listen For

- Slight, uneven pitch instability between layered tones — two oscillators tuned to what should be the same note drifting a little against each other, never quite settling, producing a slow beating that reads as thickness or life rather than as an error
- The opposite: a layered tone that sits dead still, with no beating at all — a tell for a digital source holding its tuning with a precision no analogue circuit reaches, not a sign of anything worse
- A sequenced or generative pattern that repeats but never quite identically — one value nudged, one step arriving a fraction early, cycle after cycle, rather than a loop that repeats bar for bar exactly
- Filter movement that feels driven by a hand rather than drawn on a screen — small hesitations, an uneven rate, a sweep that slightly overshoots and corrects, as against the machine-even curve of a programmed automation line
- Saturation, or harmonic grit — extra overtones appearing on a tone that did not have them at the source, the sound of a signal driven harder than its path was designed for
- A resonant filter pushed to sing a pitch of its own and used as a second voice, not just a colour on the first one — audible as a whistling tone riding above or beneath the note that triggered it

Writing About What You Hear

A strong answer is a chain: technique → sonic outcome → judgement. In this era, describe what a sound is doing before you decide what made it.

MODEL SENTENCE

"Imagine a track whose opening minutes began life as an unedited hardware-modular improvisation, later arranged and layered inside software rather than composed on a grid from the outset. The rising pattern underneath it never repeats exactly — small values shift from one cycle to the next — and the filter opening across the introduction moves with the same small hesitations a hand makes on a real knob, not the even curve of a drawn automation line. The modular improvisation supplies an instability the software arrangement then preserves; a version built entirely inside a DAW from a fixed pattern would need that unevenness added back in on purpose, because a sequencer does not misbehave on its own."

That delivers: named technique (modular improvisation, later arranged in software) · sonic outcome (a non-identical generative pattern, a hand-driven filter sweep) · AO4 judgement (hardware supplies an instability software would otherwise have to fake). A real answer replaces this hypothetical opening with the actual track being described, named in full.

A LISTENING TRAP

Do not write "this must be real analogue hardware because it sounds warm." Software built to model analogue circuits now reproduces oscillator drift, filter self-oscillation and saturation on purpose, closely enough that trained ears — and often measuring equipment — struggle to separate a careful model from the circuit it is modelling. "Warmth" alone is not safe evidence of anything. What is safe is naming the specific quality you can actually hear — detuned drift, a particular resonant character, added harmonic grit — and describing its effect, without committing to a guess about which side of the hardware line produced it.

What "Analogue Warmth" Actually Means

Chapter Five's contrast between a Prophet-5's drifting voltage and a DX7's numeric precision already established the underlying fact: analogue circuits drift, digital circuits do not. This chapter turns that fact into three describable, physical things a listener can actually name, each with a real cause. First, drift: a voltage-controlled oscillator's pitch depends on components that respond very slightly to heat and time, so two oscillators tuned to the same note never sit in perfect unison, and the slow, uneven beating between them is what a listener hears as thickness rather than as a mistake. Second, filter character: analogue filter circuits behave a little differently at every point across their range and distort in their own particular way as resonance is pushed towards self-oscillation, a signature specific enough that individual filter designs are named for exactly this quality — a "Moog sound", a "TB-303 sound". Third, saturation: driving an analogue signal path harder than it was built for adds harmonics that were never in the original wave, a soft, musical kind of distortion rather than a hard digital clip.

getting better at reproducing exactly these behaviours in software, in the same way the DX7's FM operators built a convincing bell by an entirely different route to a bell's own physics. Even inside a Eurorack case the line blurs further: several well-known modules from manufacturers such as Mutable Instruments run on digital chips doing the calculation, built and tuned to sit in a rack of otherwise fully analogue modules and behave as if they belonged there. The safe move is the same in every direction: name the specific behaviour, not the hardware you are guessing produced it.

Instruments and Records to Know

MODULAR FORMATS	Eurorack (Doepfer A-100, mid-1990s) · budget analogue clones (Behringer) · Moog's reissued Minimoog Model D
MODULE MAKERS	Make Noise · Mutable Instruments — many of its modules are digital circuits built to behave like analogue ones, sitting in an otherwise analogue rack
HYBRID HARDWARE/SOFTWARE ARTISTS	Trent Reznor and Atticus Ross — extensive hardware-modular studios reported across their film and television scoring through the 2010s
RECORDS TO KNOW	Trent Reznor and Atticus Ross, *Watchmen* (2019) — reported to begin as live modular improvisation, later edited and arranged into the finished score

What the Spec Asks

PITCH CONTROL: FREQUENCY MODULATION AND TUNING	Real analogue drift and detune beating versus locked, perfectly still digital tuning — both audible within the same track
FILTER CUTOFF FREQUENCY	A sweep with a hand's small hesitations rather than a machine-even automated curve — evidence of real-time play, on either side of the hardware line
RESONANCE (Q)	A filter pushed to self-oscillation and used as a second voice; the specific character of the whistle is a fingerprint of the filter's design, not proof of analogue circuitry

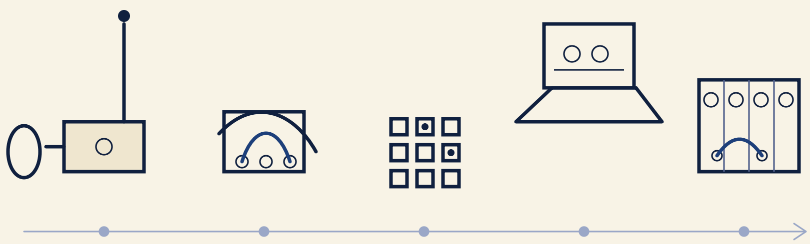

REFERENCE · THE BACK MATTER

The Back Matter

Part Six: five eras, one pair of ears. Everything the book has named, indexed for revisiting — and ten scenarios built to test what you can actually hear.

Glossary

Every technical term used in this book, in one alphabetical list. If a word turns up in a Hearing chapter and you cannot remember what it means, it is here — described the way this book always describes it, by what it does to the sound rather than by the circuit that does it.

A

Additive synthesis
Building a tone by stacking pure sine waves on top of one another rather than filtering a busy waveform down; clean, precise and organ-like.

ADSR (envelope)
The four-stage shape an envelope draws across a note — attack, decay, sustain, release — turning a raw, unchanging drone into something with a beginning, a middle and an end.

Amplifier (VCA)
The block that switches a source's tone on, holds it and lets it fade; the last stage before the output, usually governed by an ADSR envelope.

Analogue synthesis
Sound built from continuously variable electrical voltage rather than from numbers; warm, and capable of drifting slightly in pitch as its circuits warm up — the "warmth" players still pay for today.

Attack
How quickly a note reaches full volume once triggered. Near-zero for a pluck or a drum hit; stretched out for a pad that swells in rather than starts.

B

Band-pass filter (BPF)
Passes a narrow window of frequencies around a chosen point and rolls off everything above and below it; nasal, narrow, telephone-like.

C

Cutoff frequency
The point in the spectrum where a filter begins to act. Sweeping it upward on a low-pass filter opens a muffled tone into a bright one — the single most expressive move in subtractive synthesis.

D

Decay
How quickly a sound falls from its peak down to its held, sustained level, immediately after the attack.

Detune
Running two oscillators very slightly apart in pitch so they beat against each other, thickening one voice into something wider and richer.

edited through menus rather than one knob per function.

E

East Coast / West Coast synthesis

Shorthand for two rival philosophies born in the same mid-1960s moment: Moog's keyboard-and-filter lineage (East Coast) against Buchla's sequencer-and-touch-plate lineage, built to escape the keyboard altogether (West Coast).

Eurorack

A modular synthesiser format — a standard case height, module width and voltage rail — that lets modules from hundreds of different makers plug into the same rack.

F

Filter (VCF)

The block that sculpts a source's tone by removing or emphasising frequencies; the block that gives subtractive synthesis its name.

Filter envelope

A second ADSR routed to a filter's cutoff instead of the amplifier, so a sound's brightness can open and close independently of its volume.

Filter sweep

The audible movement of a cutoff frequency over time, whether pushed by hand, by an envelope or by an LFO.

FM synthesis

Frequency Modulation synthesis: one oscillator rapidly bending the pitch of another at audio speed, throwing off bright, glassy, often bell-like or metallic harmonics that no filter could carve from a single wave.

G

Glissando

A continuous slide through every pitch between two notes, with no discrete steps in between; the default behaviour of any instrument built with no keyboard to divide it.

H

High-pass filter (HPF)

Passes the top of the spectrum and rolls off the bottom; thins a sound out and clears away low-end mud.

Hybrid rig

Pairing hardware synthesisers or a modular system with a computer and DAW, keeping hands-on physical control while still recording, sequencing and processing on screen.

L

LFO (low-frequency oscillator)

An oscillator too slow to hear as a pitch, typically below about 20 Hz, used instead as a wobble aimed elsewhere: at pitch it is vibrato, at amplitude it is tremolo, at a filter's cutoff it is the rhythmic wah, or wobble, heard in a slow pad or a bass drop.

Low-pass filter (LPF)

darker and warmer as it closes.

M

MIDI

Musical Instrument Digital Interface: a shared digital language, given away free in 1983, that lets keyboards, sequencers and sound modules from different manufacturers talk to one another.

Mod wheel

A lever beside the keyboard, pushed during a note rather than set in advance, most often wired to vibrato depth so expression can be added while a note is still sounding.

Modular synthesiser

An instrument with no fixed signal path; oscillators, filters, amplifiers and controllers are separate modules, wired together fresh with patch cables for every new sound.

Modulation

Using one signal to move or control another — an LFO wobbling a filter, an envelope opening it, a hand turning a wheel; the general name for everything in the Control layer.

Monophonic

Able to sound only one note at a time, however many keys are held down.

N

Noise

A signal with no identifiable pitch, its energy spread across the spectrum rather than gathered at one frequency; filtered and shaped by a fast envelope, it becomes a snare, a hi-hat or the breath at the front of a flute patch.

O

Oscillator (VCO)

The source that generates a synthesiser's raw, repeating tone; its waveform shape decides most of a sound's character before anything downstream touches it. In an analogue instrument its pitch is set by a control voltage — a voltage-controlled oscillator, or VCO.

P

Patch / patching

Connecting modules together, usually by cable, to build a sound from nothing. On an instrument with no memory, a patch exists only for as long as the cables stay plugged in.

Patch memory

The ability to store a sound's settings as numbers and recall them at the press of a button, rather than rebuilding the sound by hand every time it is needed.

Polyphonic

Able to sound more than one note at once, each note running its own copy of the source–filter–amplifier chain.

Portamento (glide)

Sliding smoothly from one note's pitch to the next instead of jumping cleanly between them.

Preset

outlasted the machines they were written for.

R

Release

How quickly a sound fades once a key is lifted, from its sustain level back to silence.

Resonance

An emphasis right at a filter’s cutoff frequency. A little adds definition; a lot turns squelchy and vocal; pushed further still, the filter can self-oscillate into a pure tone of its own.

ROMpler

A keyboard or workstation that builds its sound from short recordings of real instruments stored in memory chips, rather than generating a waveform from scratch.

S

Sample-based synthesis

Building or shaping a sound from a short recording rather than from an oscillator; a different family from FM, subtractive or wavetable synthesis.

Sawtooth wave

The brightest and buzziest of the four basic waveforms, close to a bowed string or a distant brass section; the standard starting point for subtractive synthesis.

Self-oscillation

A filter driven with enough resonance that it generates a pure tone of its own, briefly becoming a second oscillator.

Sequencer

A device that steps automatically through a pattern of pitches or control voltages at a set rate, with no hand touching a key.

Sidechain compression

Routing one sound, often a kick drum, to duck the level of another, often a pad or bass, in time with it — the rhythmic pumping heard behind most dance-music breakdowns.

Sine wave

The purest tone available: smooth, round and carrying no harmonics on top of the fundamental.

Softsynth

A synthesiser that exists entirely as software, running inside or alongside a DAW rather than as a physical instrument.

Square wave

Hollow and reedy, with a nasal, clarinet-like quality; narrowing its width sets the tone moving.

Subtractive synthesis

Starting with a harmonically busy waveform and using a filter to carve away what is not wanted, rather than building a tone up from nothing; the family this book’s exam actually tests.

Supersaw

Several detuned sawtooth oscillators stacked and played as a single voice, producing the wide, churning stack of sound behind most trance and EDM leads.

Not a time but a level: how loud a note stays while a key remains held down.

Sync (oscillator sync)

Forcibly restarting one oscillator's cycle with another, producing a searing, ripping edge and harmonics that no single waveform holds on its own.

T

Tremolo

A regular pulsing in loudness, produced by an LFO aimed at amplitude.

Triangle wave

Soft and slightly hollow, a gentler, flute-like cousin of the square wave.

U

Undo

The software-era ability to reverse a change instantly and completely; a freedom no hardware synthesiser with real knobs has ever offered.

Unlimited polyphony

A software instrument's freedom from the old rule of one voice per physical circuit: as many notes at once as the computer can calculate, not as many oscillators as were built into the case.

V

Velocity

How hard a key is struck, routed by default to loudness and, on many patches, to a filter's brightness too.

Vibrato

A regular wavering in pitch, produced by an LFO or a hand aimed at an oscillator's tuning.

Voltage control

Governing a synthesiser's pitch, tone and loudness with electrical voltages rather than mechanical linkages. Standardised by Robert Moog at one volt per octave in 1964, it is the idea that made modular synthesis, and everything patched since, possible.

W

Wavetable synthesis

Scanning through a stored table of single-cycle waveforms and morphing between them, producing a timbre that moves and evolves inside a single held note.

White noise

A signal containing equal energy at every frequency across the audible spectrum; heard as a steady hiss, like an untuned radio between stations.

Plate Notes

The six plates in this book are original ink drawings, made specifically for this edition rather than adapted from any historical photograph. One artist drew all six, in the same hand throughout, so the century reads as a single unbroken line rather than six different illustrators' guesses.

Plate One. A theremin — a plain wooden cabinet with two upright antennae — drawn with a hand hovering near the pitch antenna, never touching it. It opens Part One, Electricity Finds a Voice, and states the era's whole argument in a single gesture: the first electronic instrument a musician could play was one they were never allowed to touch.

Plate Two. A modular synthesiser panel, its face crowded with patch cables, one cable lifted mid-air by a hand about to seat it in a socket. It opens Part Two, The Patch-Cord Years, where every sound began from an empty panel and ended the moment the power went off.

Plate Three. A long, wooden-ended synthesiser panel in the style of the Prophet-5, five voice chips drawn ghosted in the air above the case, a bank of preset buttons running along its front edge. It opens Part Three, Presets and Chips, where the ghosted chips picture the chapter's real subject: a sound that could finally be stored rather than merely played.

Plate Four. A laptop whose screen shows a single soft-synth window, with identical duplicate windows cascading behind it into the distance. It opens Part Four, Software and the Infinite Instance, and makes its point without a caption: once an instrument is a file, there is no limit to how many you can have open at once.

Plate Five. A Eurorack case mid-patch, a pair of headphones resting on top of it. It opens Part Five, The Modular Return, and the headphones do the work: someone has chosen to sit with this machine alone, patching it by hand, when software could have done the same job in a fraction of the time.

Plate Six. Five small emblems in a single row — a theremin, a tangle of patch cables, a bank of preset buttons, a laptop screen, a modular case — the whole century compressed into one wordless image. It opens Part Six, the book's back matter, where every term the five parts have introduced is gathered for revision.

Where to Go Next

Hand-picked outside reading and watching, gathered from all five parts. These are the sources worth your time if a particular era, or a particular machine, has caught your interest.

ELECTRICITY FINDS A VOICE · 1920–1963

The Bob Moog Foundation's collected interviews — engineers and performers describing, in their own words, the theremin's earliest concert tours and the workshop culture that grew up around Leon Theremin's invention.

Sisters With Transistors (documentary) — a history of women in electronic music that gives Daphne Oram and Delia Derbyshire's years at the BBC Radiophonic Workshop their fullest screen treatment yet.

THE PATCH-CORD YEARS · 1964–1977

Wendy Carlos's own technical notes on Switched-On Bach (wendycarlos.com) — a first-hand account of patching a Moog modular one note at a time, take by take, to build a full multi-voice recording from an instrument that could only ever play one note.

Vintage Synth Explorer: Minimoog Model D — the instrument's full panel layout and specification, useful for seeing exactly how few controls the Four Blocks needed once they were hard-wired into one portable case.

PRESETS AND CHIPS · 1978–1995

Sound on Sound's Synth Secrets series (Gordon Reid) — a long-running magazine series that works through patch memory, MIDI and FM synthesis from first principles, aimed at exactly the reader this book is written for.

Vintage Synth Explorer: Sequential Circuits Prophet-5 — the specification sheet for the instrument that first turned a patch into something you could save and recall at the press of one button.

SOFTWARE AND THE INFINITE INSTANCE · 1996–2009

Steinberg's own history of the VST format — the manufacturer's account of turning a hardware plug-in slot into a software standard nearly every DAW on the market still uses.

Native Instruments' own retrospective on Reaktor — how a modular synthesis environment was rebuilt entirely in software, and what changed once a patch cable became a line drawn with a mouse rather than an object you had to own.

THE MODULAR RETURN · 2010–NOW

Perfect Circuit's Eurorack history essays — a modular retailer's own long-form account of how Dieter Doepfer's A-100 format became the shared standard behind the 2010s hardware revival.

Sound on Sound's coverage of the Eurorack revival — reviews and features on why a generation raised on unlimited software chose to buy real, scarce hardware anyway.

This book's whole argument has been constraint on the instrument. Its twin argument — constraint on capture, on how a performance already in the air got recorded, blended and rebuilt — is told in The Story of the Studio, Volume One on this shelf.

The Listening Companion

Every era this book describes has a soundtrack. Play each track below and listen for the one thing named beside it — that is the ear-training half of the argument the chapters have already made in words. Read the list in order and the century plays out exactly as Parts One to Five tell it.

ELECTRICITY FINDS A VOICE · 1920–1963

MIKLÓS RÓZSA — SPELLBOUND (1945)	A wavering theremin line glides between pitches with no attack and no fixed reference point at all.
BERNARD HERRMANN — THE DAY THE EARTH STOOD STILL (1951)	Two theremins glide and swoop together, pitch entirely continuous, standing in for a flying saucer.
OLIVIER MESSIAEN — TURANGALÎLA-SYMPHONIE (1946–48)	The Ondes Martenot bends and swells inside a full orchestra, a controlled voice no strings can match.
DELIA DERBYSHIRE — DOCTOR WHO THEME (1963)	Every note began as a separately recorded oscillator tone, cut and spliced by hand, not synthesised.

THE PATCH-CORD YEARS · 1964–1977

WENDY CARLOS — SWITCHED-ON BACH (1968)	One monophonic Moog voice, overdubbed part by part, builds dense keyboard counterpoint no single performance could play live.
EMERSON, LAKE & PALMER — TARKUS (1971)	A monophonic Moog line, prone to drifting slightly out of tune between takes, plays one note at a time throughout — the same instrument Emerson would soon carry on tour.
HERBIE HANCOCK — CHAMELEON (1973)	An ARP Odyssey glides between held notes, one line at a time, with no pick attack.
KRAFTWERK — AUTOBAHN (1974)	A single synthesiser line locks to a steady, mechanical pulse, hand-played yet strikingly even throughout.

PRESETS AND CHIPS · 1978–1995

VANGELIS — BLADE RUNNER (FILM SCORE, 1982)	A CS-80's warm, breathy polysynth pads and glides score the whole film, weighted by a finger's touch.
SEQUENTIAL CIRCUITS PROPHET-5 — WARM, RECALLABLE ANALOGUE PADS (1978–)	A patch dialled in once is recalled instantly at the touch of one button, unchanged, session after session.
YAMAHA DX7 — FACTORY ELECTRIC PIANO, "E. PIANO 1" (MID-1980S)	A glassy, bell-like electric-piano tone, numerically exact and slightly cold, recurs unedited across dozens of ballads.
ROLAND JUPITER-8 — ANALOGUE POLYPHONY WITH DIGITAL PATCH MEMORY (1981–)	Eight true polyphonic voices hold a chord steady, each one brightening slightly under a harder-struck key.

SOFTWARE AND THE INFINITE INSTANCE · 1996–2009

DARUDE — SANDSTORM (2000)	Seven or more detuned sawtooth voices, spread across the stereo field, stack into one supersaw lead.
DEADMAU5 — STROBE (2009)	One sustained pad shifts shape almost imperceptibly across ten minutes, with no audible edit or repeat.
AVICII — LEVELS (2011)	A piano and pad audibly duck in volume with every kick drum hit, pumping in rhythmic time.
NATIVE INSTRUMENTS MASSIVE — THE WOBBLE DECADE'S BASS ENGINE (2007–)	A growling, filter-wobbling bass line opens and closes rhythmically, the software synth behind most of the decade's dubstep.

THE MODULAR RETURN · 2010–NOW

JON HOPKINS — SINGULARITY (2018)	A rising pattern underneath the track never repeats exactly, its filter opening with a hand's small hesitations.
TRENT REZNOR & ATTICUS ROSS — FILM AND TELEVISION SCORING (2010S)	Extensive hardware-modular studio rigs feed patient, evolving instrumental textures across a decade of soundtrack work.

EURORACK / DOEPFER A-100 — A PATCH THAT EXISTS ONLY WHILE PATCHED (1995–)	A patch built fresh from empty sockets exists only as long as the cables stay plugged in.
KAITLYN AURELIA SMITH — THE KID (2017)	A solo Buchla modular system sequences shifting patterns live, no keyboard anywhere, echoing the format's touch-plate origins.

The Era-Identification Drill

Ten unfamiliar sounds, none of them named. This drill rewards exactly that skill: place a recording by what its instrument would, or would not, let the player do — never by guessing at a date. Read each scenario, decide on an era and a reason, then check the smaller line beneath it.

1. You hear: a single, continuous tone with no attack at all, sliding smoothly between pitches, with an irregular waver that never quite repeats the same way twice. Which era, and why?

Era One — Electricity Finds a Voice, 1920–1963. No envelope generator exists yet, so there is no attack transient to hear; the uneven waver is a hand near an antenna, not a circuit, and continuous pitch is this era's default state, not an effect switched on for colour.

2. You hear: a short instrumental theme built from a handful of pure, evenly spaced pitches, each one starting and stopping with a faint click and a thin layer of hiss underneath. Which era, and why?

Era One — Electricity Finds a Voice, 1920–1963. The click is a tape splice and the hiss is analogue tape noise; mechanical evenness across a handful of discrete pitches is the signature of material built note by recorded note at a workbench, not played live by a hand.

3. You hear: one melodic line at a time, its pitch sliding smoothly into each new note rather than jumping cleanly, and a bright tone that audibly opens and darkens mid-note as though a hand were turning a dial. Which era, and why?

Era Two — The Patch-Cord Years, 1964–1977. Portamento and an audibly hand-swept filter cutoff are live, physical gestures — before patch memory existed, every movement inside a sound had to be performed at the moment of playing, not stored in advance.

4. You hear: a single sustained note thickened by two voices very slightly out of tune with each other, beating gently against one another, with no chord sounding underneath it anywhere. Which era, and why?

Era Two — The Patch-Cord Years, 1964–1977. A single voice, however thickened by detuning, is still only one voice; true polyphony, and the memory that would make a second identical voice easy to recall later, both belong to the era that follows this one.

5. You hear: a full chord held steady by several identical voices at once, each one growing brighter the harder it was struck, and the same sound recalled again a minute later at the press of one button. Which era, and why?

Era Three — Presets and Chips, 1978–1995. Several identical voices sounding together, each one brightening under a harder strike, needs genuine polyphony; recalling the whole setting again at the press of one button needs the patch memory this era invented.

6. You hear: a glassy, bell-like electric-piano tone with almost no warmth to it, its pitch and its timing landing with mechanical exactness across the whole recording. Which era, and why?

Era Three — Presets and Chips, 1978–1995. No analogue circuit holds pitch and timing this precisely; a glassy, bell-like tone built from modulation rather than filtering is FM synthesis's own signature, generated from numbers rather than voltage.

7. You hear: a huge, shimmering chord built from many near-identical detuned voices stacked on top of each other, filling the entire width of the stereo picture. Which era, and why?

Era Four — Software and the Infinite Instance, 1996–2009. Stacking this many detuned voices this wide costs nothing once a synthesiser is a software instance rather than a physical circuit with a fixed, limited number of oscillators inside it.

hollow, then buzzing — with no second note played and no hand visibly on a knob. Which era, and why?

Era Four — Software and the Infinite Instance, 1996–2009. A single note changing shape while it holds, with no second note or visible knob involved, is wavetable synthesis scanning through a stored waveshape in real time — a trick no single analogue oscillator can perform alone.

9. You hear: a warm, gently unstable pad, its pitch drifting almost imperceptibly and its filter breathing open and shut as if a hand were riding a knob live, with a soft grain under the tone. Which era, and why?

Era Five — The Modular Return, 2010–now. A hardware oscillator's slight pitch drift and a filter played live by hand are exactly what a musician is choosing to pay for, on purpose, after decades in which software made both effectively free to remove.

10. You hear: a rapidly repeating pattern of short notes, each one's pitch or tone subtly different from the last, evolving gradually over several minutes with no keyboard player visibly needed to keep it going. Which era, and why?

Era Five — The Modular Return, 2010–now. A self-generating pattern that keeps shifting with nobody visibly playing individual notes is a sequencer at work, not a keyboard part — the same idea Buchla built into his instruments back in Era Two, chosen again deliberately half a century later.

A Century: A Timeline

Five eras, one line for each instrument that changed what the argument allowed. Below, entries are grouped by the constraint they belong to, which is occasionally not the decade their date falls in — the Eurorack format is the clearest case, designed in the 1990s but left waiting for the return this book's fifth part describes. No era ends the moment the next begins; it just stops being where the newest instruments get built.

ELECTRICITY FINDS A VOICE · 1920–1963

1920	Leon Theremin builds the Theremin in Russia: two radio antennae, pitch and volume set by hand position alone, nothing touched.
1928	Maurice Martenot introduces the Ondes Martenot in Paris, adding a piano-style keyboard to the same heterodyne principle.
1929	RCA licenses Theremin's patent and manufactures the instrument for retail sale; even mass-produced, it finds almost no working musicians able to play it.
1957	RCA delivers the Mark II Sound Synthesizer to Columbia University; composers programme it with punched paper tape rather than playing it live.
1958	The BBC Radiophonic Workshop is founded at Maida Vale, London, with Daphne Oram among its co-founders.
1963	Delia Derbyshire realises Ron Grainer's *Doctor Who* theme from tape-spliced oscillator notes — tape technique, not a synthesiser.

THE PATCH-CORD YEARS · 1964–1977

1964	Robert Moog demonstrates a keyboard-controlled, voltage-controlled modular system in New York; Don Buchla, working in California at almost the same time, builds a touch-plate system with no keyboard at all.
1968	Wendy Carlos releases *Switched-On Bach*, recorded part by part on a Moog modular synthesiser.

1969	EMS, in London, releases the VCS3, a compact synthesiser patched with pins instead of cables.
1970	Moog releases the Minimoog Model D: a signal path fixed at the factory, with no patch cables at all.
1971	ARP releases the 2600, a semi-modular synthesiser that reaches schools and studios far beyond the handful of full modular owners.
1975	Oberheim's Four Voice reaches four-note polyphony by wiring four complete monophonic synthesisers to a single keyboard.
1977	Yamaha's CS-80 adds touch-sensitive polyphony and a ribbon controller — still with no way to recall a sound already built.

PRESETS AND CHIPS · 1978–1995

1978	Sequential Circuits releases the Prophet-5: a microprocessor stores every knob's position as numbers, recallable at the press of a button.
1981	Dave Smith proposes a shared digital instrument standard at the Audio Engineering Society convention — the idea that becomes MIDI.
1981	Roland releases the Jupiter-8, an eight-voice polysynth with its own bank of recallable sounds.
1983	MIDI's first public demonstration, at the NAMM trade show, connects a Sequential Prophet 600 to a Roland JP-6 — given away royalty-free to every manufacturer in the room.
1983	Yamaha's DX7 generates its entire sound from FM synthesis, with no voltage-controlled circuit inside it at all.
1987	Roland's D-50 pairs digitally sampled attack transients with an analogue-style filter, bridging FM's pure numbers and the older subtractive method.
1988	Korg's M1 plays sampled waveforms through a synthesiser's filter and envelope stage — the ROMpler workstation.

MID-1990S	Dieter Doepfer launches the Eurorack modular format, the A-100, in Germany; it stays a niche for over a decade before Part Five's return.

SOFTWARE AND THE INFINITE INSTANCE · 1996–2009

1996	Steinberg releases VST, a format that lets a synthesiser exist as software inside a computer rather than as a physical circuit.
1996	Native Instruments, founded in Berlin, releases Generator, a modular software synthesis environment later renamed Reaktor.
2000	Propellerhead releases Reason, a virtual rack of software instruments with no audio recording built in at all.
2007	Native Instruments releases Massive, the wavetable synthesiser closely associated with the decade's dubstep and electro "wobble" bass.

THE MODULAR RETURN · 2010–NOW

2010	Eurorack's slow niche growth accelerates; a wave of small manufacturers begins building single modules for the format in numbers Doepfer's original range never reached.
2014	Xfer Records releases Serum, a wavetable synthesiser with its waveform editing shown directly on screen.
MID-2010S	Moog returns the Minimoog Model D to production, built to the original analogue circuit rather than modelled in software.
TODAY	DAWs and Eurorack cases sit side by side in the same studios; a century after the theremin, the argument over what a musician may touch, save, copy or lose is still being played, not settled.

Hear These Machines With Your Ears

Everything above is a description of a sound. The tools below let you build and hear the thing itself — the same Source, Filter, Amplifier and Control blocks named throughout this book, made playable.

Subtractive Synth Explorer — build the source–filter–amplifier patches this book's framework chapter describes, one oscillator and one filter at a time.

resources.musictechstudio.co.uk/subtractive-synth-explorer

The Map Room — the whole of Component 4, hung in one room; drag to see how every synthesis term in this book connects to the exam that asks about it.

resources.musictechstudio.co.uk/map-room

Octave & Period Trainer — the frequency-and-period relationship behind an oscillator's pitch, and behind an LFO slow enough to hear as rhythm rather than tone, made interactive.

resources.musictechstudio.co.uk/octave-period-trainer

BPM & Delay Calculator — set a tempo and see the exact millisecond value behind a synced LFO rate, an arpeggiator division, or a sequenced pattern.

resources.musictechstudio.co.uk/bpm-delay-calculator

For a sense of how these five eras turn into exam questions, see **Exam Difficulty** — sherborne.musictechstudio.co.uk/exam-difficulty.