


Radish

Five modules for drones, for Voltage Modular



Radish is a series of five modules for drones: three oscillators that beat against one another at rates set in hertz, three resonant filters tuned the same way, a pair of ring modulators and frequency shifters, a tape loop that wears on every pass, and a six-lane sequencer drawn on a drum that turns once in anything from ten seconds to two hours.

This manual describes the ideas the five share, every control and jack of each, and some ways of patching them together. Its graphs are measured from the modules' own DSP code.

Contents

Introduction	2
1 · Littoral	6
2 · Fluvial	11
3 · Calderal	16
4 · Alluvial	21
5 · Glacial	28
6 · Drones	34

Introduction

What the series is for

Radish is a series of five modules for drones: sounds held for minutes or hours, in which the pitch hardly changes and yet the sound never settles. The music of the late Éliane Radigue is an inspiration for the series.

The modules at the heart of the series do simple things. An oscillator, a filter and a ring modulator each have a short mathematical description, and most players will have used many versions of all three. These differ in two ways. They are built for very fine relative tuning, in fractions of a hertz or in just ratios, so that the relationships between their voices are exact. And they bring in some of the impurities of real equipment, drift, mains hum and a faint colouring of the harmonics, so that the electrical workings breathe through the mathematical abstraction as they do in a circuit.

The two together are what keep a drone worth listening to: strict relationships, perturbed by something like the noise of their surroundings. Pure sines held at a constant pitch soon stop being interesting to the ear; a drone made this way does not.

A held tone becomes worth listening to through small differences: two oscillators a fraction of a hertz apart, a filter just short of ringing, a loop of tape losing a little of its sound on every pass. Each Radish module is built to make one kind of small difference precise, so that it can be set, repeated, and changed slowly.

Module	What it is
Littoral	Three oscillators tuned from one reference pitch, each set a number of hertz away from it, so that they beat at exact rates.
Fluvial	Three resonant filters tuned the same way. Past a mark on their resonance they ring by themselves, as sines, and beat as Littoral's voices do.
Calderal	Two ring modulators. Each can keep one sideband, which makes it a frequency shifter: every partial of a sound moves by the same number of hertz.
Alluvial	A stereo tape loop of up to a minute, which loses a little of its sound on every pass, as tape does.
Glacial	Six lanes of control voltage, drawn on paper wrapped round a drum that turns once in anything from ten seconds to two hours.

Littoral and Fluvial make sound; Calderal and Alluvial transform it. Glacial orchestrates slow change on purpose: it moves the other four's controls more slowly than a hand would, along curves drawn in advance. Alluvial layers drones on a loop of tape and wears what builds up there, a little loss and distortion on every pass. The chapters follow that order. Chapter 6 gives some patches that use them together.

Beating

Two sine waves of nearly the same frequency, added together, are not heard as two pitches. They are heard as one pitch whose loudness rises and falls. Where the two are in phase they add; half a beat later they cancel. The sum swells and fades as many times a second as the difference between their frequencies: tones at 100 Hz and 100.2 Hz beat 0.2 times a second, once every five seconds.

The rate of a beat is a difference in hertz. A tuning control in cents cannot set it, because the same interval in cents is a different number of hertz at every pitch: 10 cents is 0.38 Hz at 65 Hz, and 3 Hz at 520 Hz. So Littoral and Fluvial tune each voice from a shared reference in the usual way, and then add a last offset in hertz, **BEAT**. Two voices that differ only in BEAT beat at exactly the difference between their two settings, whatever the reference, and the same settings give the same beats every time.




Figure 1. Half a second of two of Littoral's voices on sine at A1 (55 Hz), voice 2 with **BEAT** at +4 Hz, and their sum at **MIX**. Each voice is steady; the sum swells and dies away four times a second. A beat this fast is heard as a flutter; the slow beats a drone is made of are the same shape drawn out over seconds.

Tones that are not pure sines beat in more than one place. If two tones each have harmonics, their second harmonics are twice as far apart in hertz as their fundamentals, their thirds three times as far, and so on. Each harmonic beats at its own rate in its own register: two tones 0.2 Hz apart beat at 0.2 Hz on the fundamental, 0.6 Hz on the third harmonic and 1 Hz on the fifth. Littoral's sine is made from a triangle, and keeps a little of the triangle's third and fifth harmonics for this reason (Figure 2). Two pure sines beat only once, and a drone made of them sounds still in a way that one made of Littoral's does not.




Figure 2. Two of Littoral's voices on sine, 0.2 Hz apart, with **DRIFT** fully left. Each line is the amplitude of one harmonic of the two voices' sum, as a share of its own largest. The fundamental swells once every five seconds, the third harmonic three times as often, the fifth five times.

Three voices beat in three pairs. With **BEATs** of 0, +0.1 and -0.15 Hz the pairs beat at 0.1, 0.15 and 0.25 Hz, and the three pulses fall in and out of step with one another, coming back to where they started every 20 seconds. Settings less simply related take far longer to come round.

Beating also comes from places other than two oscillators. A filter ringing by itself a fraction of a hertz from one of a drone's harmonics beats against it (Fluvial). A ring modulator's sidebands carry the beats of what goes into it to new pitches (Calderal). A faint hum near a harmonic beats against it (Littoral's **MAINS**). And a frequency shift moves every partial of a beating drone by the same number of hertz, which leaves every beat as it was.

Two places look as though they should beat and do not. A filter set a fraction of a hertz off one of a drone's harmonics, with the drone going through it, does not beat against it: driven, a filter sounds at the frequency that drives it, whatever its own tuning, and its **BEAT** only sets how much of the harmonic passes. Only a filter ringing by itself, past OSC and beside the drone rather than under it, beats (Fluvial). And two voices of one drone ringed against each other do not beat: their product is a steady tone at twice their pitch and a difference too slow to hear. The beat comes back from their harmonics, from a third voice, or from the dry signal beside the ring (Calderal).

Movement within movement within stasis

A Radish drone moves at several speeds at once, each inside the one slower than itself.

- **Seconds:** the beats. Between about a hundredth of a hertz and a few hertz, they are the fastest movement in the drone, and the most audible.
- **Minutes:** the beats’ own rates change. Littoral and Fluvial drift slowly, as analogue oscillators do; Glacial’s lanes move BEATs, levels, cutoffs and shifts over minutes. At 100 Hz a cent is about 0.06 Hz, so a drift of a few cents is enough to slow a 0.2 Hz beat, stall it and speed it up again.
- **Hours:** a loop of tape wears pass by pass, darkening and quieting what was laid down on it, and Glacial’s slowest turn takes two hours.

The pitch of the drone, meanwhile, stays where it was set: that is the stasis. Within it the beats move, and within the beats’ movement their rates move. Each layer is too slow to be heard as motion from the one above it, and so the changes are heard as the sound changing rather than as something being done to it.




Figure 3. The time scales the series works at, from the ranges of its controls. Wow and flutter are the tape’s speed wavering, the fastest movement here. How long a layer takes to wear into the tape depends on the loop’s length and Alluvial’s settings.

What the circuits added

The drones this series is made for were first made with analogue oscillators, filters, ring modulators and tape, and much of their timbre came from what those circuits and that tape did besides their nominal job. A sine made by rounding off a triangle kept some of the triangle. Pitch wandered as the circuits warmed up. The power supply’s hum leaked into the audio. Tape lost a little of its high end and gained a little noise on every pass. Each of these contributed to the timbres the series exists to recreate, and each is modelled where it did:

What	What it contributes	Where
The shaper’s residue	a third and fifth harmonic on the sine, which beat three and five times as fast as the fundamental	Littoral
Drift and warm-up	slow wandering in pitch, shared by every voice in a module more than it differs between them, so the beats wander less than the pitch	Littoral, Fluvial
Mains hum	faint sidebands 100 or 120 Hz either side of each harmonic, which beat slowly against the harmonics of a low drone	Littoral
The reset tick	a faint click once a cycle, heard at low pitches as a click train beating against the other voices’	Littoral
A multiplier’s leakage	a little of each input under the ring’s sidebands	Calderal
Tape	loss of highs, hiss, saturation, print-through, wow and flutter and dropouts, each written onto the tape so that it builds up pass by pass	Alluvial

Every one of them has a control, and every one can be turned down to nothing where a clean signal is wanted. Each Littoral and Fluvial is also a unit of its own. The details that vary from one analogue circuit to the next (the levels of a shaper’s harmonics, the course its drift takes, the faint noise a filter’s ringing starts from) are drawn when

the module is made and saved with the patch. Two Littorals in one patch sound like two units rather than one twice, and a patch sounds the same each time it is loaded.

What the five share

The reference. Littoral and Fluvial tune their three voices from one reference pitch, set by **REF** as 1 V/oct where 0 V is C2 (65.41 Hz), with **V/OCT** added to it. **REF** out carries the reference, with the module's shared drift on it, and **REF IN** takes one. A cable from one module's REF to another's REF IN puts both on the same reference and the same drift, as though they were in one cabinet.

Tuning a voice. Each voice is tuned by **OCT** in octaves, **COARSE** in semitones, just intervals or hertz, as its switch says, **FINE** in cents, and **BEAT** in hertz, which is added last. The just intervals, one for each semitone from the reference, are 1/1, 16/15, 9/8, 6/5, 5/4, 4/3, 7/5, 3/2, 8/5, 5/3, 7/4, 15/8 and 2/1.

CV and AMT. Every CV input has an amount knob, **AMT**, beside it. AMT is off at its centre, where it loads, scales the CV up to its full effect as it turns right, and inverts it as it turns left. CVs are on a ± 5 V scale: at full AMT, ± 5 V moves a control across the whole of its travel. Glacial's lanes give -5 to $+5$ V, or 0 to $+5$ V, so a lane can take any control in the series across its range.

Levels. Each of Littoral's and Fluvial's voices and each of Calderal's channels ends in a VCA, with a **LEVEL** knob and a CV, so a patch needs no separate VCAs or mixer for them. The level follows a square law, so a slow fade moves evenly. Audio is at ± 5 V.

The faces. Outputs are on vermilion blocks. Hovering over a knob shows its value, in hertz, cents, decibels or seconds; a BEAT or SHIFT tooltip can be clicked and a value typed in. Parts of a module that nothing is listening to are not computed.


1 • Littoral

littoral /ˈlɪt(ə)rəl/

1 *adjective* of, on or near the shore of a sea or a lake.

2 *noun* the zone of a shore between the high and low water marks.

Littoral is three oscillators for drones. All three are tuned from one reference pitch, and each can be set a precise number of hertz away from it with its **BEAT** knob, so that the voices beat against one another at exact rates. Each voice ends in its own VCA, and **MIX** sums the three.



- ① **REF** the reference, C2 where it loads
- ② **V/OCT** added to the reference, 1 V/oct
- ③ **REF IN** replaces REF
- ④ **DRIFT** how far the pitch wanders
- ⑤ **MAINS** 50 or 60 Hz hum
- ⑥ **Trim** how much hum
- ⑦ **WARM** restarts the warm-up, on a rise
- ⑧ **Button** restarts the warm-up, by hand
- ⑨ **TICK** the reset tick's level
- ⑩ **REF** the reference, out
- ⑪ **MIX** the three voices



- ⑫ **OCT** whole octaves, -2 to +2
- ⑬ **COARSE** semitones, just steps or hertz
- ⑭ **Hz, SEMI, JUST** what COARSE counts in
- ⑮ **FINE** ±50 cents
- ⑯ **BEAT** ±5 Hz, added last
- ⑰ **CV** on the beat
- ⑱ **AMT** the beat CV's amount
- ⑲ **FM, x1** the beat CV's range
- ⑳ **WAVE** saw, square, pulse, triangle, sine
- ㉑ **LEVEL** the voice's VCA
- ㉒ **CV** on the level
- ㉓ **AMT** the level CV's amount
- ㉔ **OUT** the voice

The three voices are the face's three rows; each has the same controls. With every control where it loads, the three play the same sine at C2, and **MIX** carries them.

The reference

REF sets the pitch all three voices are tuned from, three octaves either side of C2, as 1 V/oct where 0 V is C2. Its tooltip gives the frequency and the nearest note, as Reference: 65.41 Hz (C2).

V/OCT adds to REF, a volt an octave, so the whole bank can follow a keyboard or a sequencer. Turning REF glides over a few milliseconds; V/OCT and REF IN do not, so a note lands at once and a chained bank follows the first exactly.

REF IN replaces REF. Patch another Littoral's REF into it and the two banks share one reference: six voices beating against one pitch. A chain of three gives nine. V/OCT still adds.

Each voice

Tuning

OCT moves the voice by whole octaves, two down to two up.

COARSE moves it further, in what its switch says:

- **SEMI**, the middle, where it loads: equal-tempered semitones, an octave either way.
- **JUST**, down: just intervals over the reference, one for each semitone. Up from the centre they are 1/1, 16/15, 9/8, 6/5, 5/4, 4/3, 7/5, 3/2, 8/5, 5/3, 7/4, 15/8 and 2/1; left of centre the same an octave lower. Switching from SEMI to JUST moves a voice to the nearest just interval. A just fifth is exact: a voice on 3/2 over a reference of 100 Hz runs at 150 Hz.
- **HZ**, up: a frequency added to the voice, up to 100 Hz either way.

Its tooltip gives the setting: Coarse: +7 semitones, Coarse: 3/2, Coarse: +12.5 Hz.

FINE tunes the voice up to 50 cents either way.

BEAT sets the voice a number of hertz from where the controls above put it: up to 5 Hz either way. It is added after everything else, so two voices that differ only in BEAT beat at exactly the difference of their BEATs, whatever the reference. Its law is a cube, very fine near the centre: a tenth of the way from the centre to an end is five thousandths of a hertz, and half way is 0.625 Hz. Its tooltip gives the rate to a thousandth of a hertz, finer below that; click it and type a rate in hertz to set it exactly.

A voice's tuning from its knobs, OCT, COARSE, FINE and BEAT, glides over a few milliseconds when any of it moves, so turning COARSE does not click, and the glide settles exactly.

BEAT and **FINE** do different jobs. FINE is in cents, so the beat it makes changes with pitch: 10 cents is 0.38 Hz at 65 Hz and 3 Hz at 520 Hz. BEAT is in hertz, so a beat set at one pitch is the same beat at any other.




Figure 4. BEAT's law, 5 Hz times the cube of the knob's distance from its centre. The flat middle is where slow beats are set.

BEAT's CV

CV and **AMT** under BEAT: a CV on the beat, through its amount. AMT loads at its centre, where the CV does nothing; at full right each volt adds a hertz, so ± 5 V moves the beat as far as the knob does; left of centre inverts it. The CV does not glide: it arrives as it is. A slow Glacial lane into it composes a voice's beat over an hour.

The small switch beside AMT sets the CV's range. Down, at **x1**, it is a hertz a volt. Up, at **FM**, it is twenty: ± 5 V at full AMT moves the voice ± 100 Hz, and a CV at audio rate is linear FM, deep enough to be heard as FM rather than a shimmer. A low voice driven that far goes through zero and runs backwards, and comes back the same way, without a stall or a click.

WAVE

WAVE chooses what the voice puts out, drawn round the knob: saw, square, pulse, triangle and sine. The five share one phase. The pulse is a quarter of the cycle wide, with no control of its width.




Figure 5. One cycle of each waveform at full LEVEL, measured at 120 Hz.

The sine

Each voice makes its sine as analogue oscillators did: a sawtooth core, folded into a triangle, and the triangle rounded off by a shaper. The rounding is never exact. It leaves a third harmonic about 36 dB under the fundamental, a fifth about 46 dB under it, a trace of the second, and nothing above the fifth. That residue is why two of Littoral's sines beat in three registers at once (Figure 2). The residue rises a little with pitch, as the shaper was driven harder at the top of the range: the third harmonic from -37 dB at 30 Hz to -35 dB at 2 kHz.

No two shapers matched, and no two of Littoral's do. Each voice's harmonics differ from the others' by a decibel or two, and each Littoral's differ from every other's. They are drawn when the module is made and saved with the patch.




Figure 6. The sine's harmonics, for voice 1 of two different Littorals. Above the fifth, what remains is the reset tick, about 90 dB down.

TICK, the small knob beside WARM, sets the tick the core makes as it resets, which leaks through the shaper once a cycle. At its centre the tick is about 67 dB under the tone at 50 Hz; fully left it is gone; fully right it is 12 dB louder. Its tooltip gives its level against the centre's, as Tick: $+6.0$ dB. The tick is the same size every cycle, so it is a larger share of a high tone than a low one, about 3 dB more for each octave; but only at low pitches, where the cycles are long, is it heard as a faint click train, beating against the other voices' ticks. It is on the triangle and the sine.

Levels

Each voice's output passes through its own VCA, on the right of its row.

LEVEL sets the voice's level. It loads at 80%. The level follows a square law, close to how loudness is heard, so a slow fade moves evenly rather than rushing at the end.

CV and **AMT** beside it: a CV on the level, through its amount, off at the centre. At full right, $+5$ V adds LEVEL's whole travel and -5 V takes it away. LEVEL and AMT glide when they are turned; the CV is not smoothed, so it takes anything from a slow fade to AM at audio rate. With LEVEL at zero, a Glacial lane on POS fades the voice from silence to full.

OUT is the voice after its VCA: ± 5 V at full level.

Drift and warm-up

The pitch of an analogue oscillator wanders as its circuit warms and as the room changes, and oscillators in one cabinet wander together more than they wander apart. That wandering is heard in the beats: at 100 Hz a cent is 0.06 Hz, so a drift of a few cents is enough to make a 0.2 Hz beat slow, stall and quicken over the course of a piece. Littoral's drift has two parts: one shared by all three voices, which rides on the reference, and a smaller one for each voice. The beats, which depend on the differences between voices, wander less than the pitch.

DRIFT sets how far the pitch wanders. At its centre, the shared drift wanders a few cents over several minutes and each voice's own about a cent. Fully right doubles both; fully left holds every voice still, and the beats are exactly what the **BEAT** knobs say. Its tooltip reads **Drift: 50%**.

A Littoral starts cold when it is made and whenever its patch loads. The bank starts 8 cents flat, each voice up to 2 cents either side of that, and rises towards pitch over the next twenty minutes or so, most of it in the first ten. **WARM**, the jack, starts the warm-up from cold on each rise above 0 V, and the small button beside it does the same by hand.

DRIFT sets the warm-up as it sets the wandering: fully right the bank starts 16 cents flat, and fully left it starts in tune and stays there. A patch saved with **DRIFT** fully left is exactly as it was set from the moment it loads. The warm-up's clock runs whatever **DRIFT** says, so **DRIFT** turned up later brings in the warm-up as far as it has got.

A bank whose **REF IN** is patched takes its shared drift from the cable, so a chain drifts as one cabinet. Its own **WARM** then restarts only its voices' own warm-ups.




Figure 7. Half an hour from cold, **DRIFT** where it loads. The bank warms up towards pitch, and then wanders; the voices wander with it, and apart from one another by a few cents at most.

Mains

The power supplies of analogue equipment leaked a faint hum into the audio, at twice the mains frequency. On a sustained tone the hum appears as sidebands either side of each harmonic, and near the harmonics of a low drone they make slow beats of their own. **MAINS** puts it back.

The switch sets the supply: **50**, up, for sidebands 100 Hz either side of each harmonic, or **60**, down, for 120 Hz. The small knob beside it sets how much. At its centre the sidebands are 46 dB under each component of the tone: faint. Fully left there is none; fully right they are 34 dB down, and the hum is plain. Its tooltip reads **Hum: -46 dB**.




Figure 8. A voice's sine at 55 Hz, MAINS at 50 and its knob where it loads. The sidebands of the fundamental fall at 155 Hz and at 45 Hz (-45 Hz, folded), 46 dB under it.

Outputs

REF, on the vermillion block: the reference, as 1 V/oct, with V/OCT and the shared drift on it. Patch it into another Littoral's or a Fluvial's REF IN to share them.

MIX: the three voices, after their VCAs, each at a third, so three voices in phase at full level peak at 5 V.

Jacks

Jack	Direction	What it carries
V/OCT	in	added to the reference, 1 V/oct
REF IN	in	the reference, replacing REF, 1 V/oct from C2
WARM	in	restarts the warm-up, on each rise above 0 V
CV (BEAT)	in	the voice's beat, 1 Hz a volt at full AMT, or 20 on FM
CV (LEVEL)	in	the voice's level, ± 5 V across LEVELs travel at full AMT
OUT	out	the voice, ± 5 V at full level
MIX	out	the three voices, a third each
REF	out	the reference with V/OCT and the shared drift, 1 V/oct from C2

Every input reads 0 V when nothing is patched.

Saving


The knobs and switches are saved by Voltage Modular. Each Littoral also saves one number, a random seed, that makes it a unit of its own. From that seed come each voice's shaper (the levels of its harmonics and its tick) and the path its drift takes from cold. A patch that loads starts the bank cold again, and it warms up and wanders the same way each time.

2 • Fluvial


fluvial /'flu:vɪəl/

adjective of, found in, or made by a river.

Fluvial is three resonant filters for drones, tuned as Littoral's oscillators are: from one reference pitch, each set a precise number of hertz away from it with its **BEAT** knob. Past the **OSC** mark on its **Q** a filter rings by itself, as a sine at exactly the pitch it is set to, so two ringing filters beat as two of Littoral's voices do. Below the mark, it rings when something passes through it or strikes it.



- | | |
|---------------------|--|
| 1 REF | the reference, C2 where it loads |
| 2 V/OCT | added to the reference, 1 V/oct |
| 3 REF IN | replaces REF |
| 4 DRIFT | how far the tuning wanders |
| 5 CHAIN, BANK, FREE | what IN 2 and IN 3 hear with nothing in them |
| 6 REF | the reference, out |
| 7 MIX | the three voices |



- | | | | |
|-------------------|-----------------------------|-----------|---------------------------------------|
| 8 Lamp | lit while the voice runs | 9 IN | the audio the filter works on |
| 10 PING | strikes the filter | 11 Button | strikes it by hand |
| 12 OCT | whole octaves, -2 to +6 | 13 COARSE | semitones, just steps or hertz |
| 14 HZ, SEMI, JUST | what COARSE counts in | 15 FINE | ± 50 cents |
| 16 FREQ | a CV on the cutoff, 1 V/oct | 17 AMT | FREQ's amount |
| 18 AUDIO, CTRL | how often FREQ is read | 19 BEAT | ± 5 Hz, added last |
| 20 CV | on the beat | 21 AMT | the beat CV's amount |
| 22 Q | resonance; rings past OSC | 23 CV | on Q |
| 24 AMT | Q's CV's amount | 25 MODE | low-pass, band-pass, high-pass, notch |
| 26 LEVEL | the voice's VCA | 27 CV | on the level |
| 28 AMT | the level CV's amount | 29 OUT | the voice |

The filters are clean. Nothing in the resonance loop saturates, so turning up the resonance does not pull the pitch, and a ringing filter holds its level without a limiter bending it: its pitch is exact, and two ringing filters beat at the difference of their BEATs to a thousandth of a hertz.

Only a filter ringing by itself beats. Driven, a filter below OSC sounds at the frequency that drives it, and its BEAT only sets how much of it passes, so a filter set just off a drone's harmonic, with the drone through it, does not beat against it. For a beat, ring the filter past OSC, on FREE, beside the drone rather than through it.

The reference

REF sets the pitch all three filters are tuned from, three octaves either side of C2, as 1 V/oct where 0 V is C2. Its tooltip reads Reference: 65.41 Hz (C2).

V/OCT adds to **REF**, a volt an octave. **REF IN** replaces **REF**. Patch a Littoral's **REF** into it and the filters are tuned from the oscillators' reference, drift and all: a filter set to the same pitch as a voice stays there as both drift. **V/OCT** still adds to it. **REF** glides over a few milliseconds when turned; **V/OCT** and **REF IN** are read about 1,500 times a second, and do not glide.

Each voice

The three voices are the face's three rows; each has the same controls.

IN and PING

IN is the audio the filter works on. What **IN 2** and **IN 3** hear with nothing in them is **NORMAL**'s choice (below); the arrow into each is that normal, and a cable in an **IN** breaks it.

PING strikes the filter. Each rise above 0 V strikes it, as hard as the trigger is high, up to 5 V. The small button beside the jack strikes it by hand, at full. A struck filter rings as a sine from zero, whatever its Q. Below OSC it dies away as fast as its Q lets it; past OSC it swells and settles back to its own level (Figure 13).

Tuning

The tuning is Littoral's. **OCT** moves the filter by whole octaves, from two down to six up: further up than Littoral's, because a filter is tuned to a drone's harmonics as well as its fundamental. **COARSE** counts in **SEMI**, **JUST** or **HZ**, as its switch says, with the same just intervals. **FINE** tunes up to 50 cents either way. **BEAT** adds up to 5 Hz either way, last, very finely near the centre; its tooltip gives the rate, and takes a typed one. Its **CV** and **AMT** move the beat a hertz a volt at full **AMT**; there is no FM range here, as a filter cannot be tuned through zero. The tuning from the knobs glides over a few milliseconds when it moves, and settles exactly.

FREQ

FREQ and its **AMT**, under **FINE**: a CV on the cutoff, a volt an octave at full **AMT**, off at the centre, inverted to the left.

The small switch beside them says how often **FREQ** is read:

- **CTRL**, down, where it loads: about 1,500 times a second. That is plenty for a slow CV, a Glacial lane or an LFO; a fast CV is heard to step.
- **AUDIO**, up: every sample, so a CV at audio rate is exponential FM of the filter. A ringing filter becomes an FM oscillator, its sine growing sidebands as **AMT** comes up; a filter below OSC smears what passes through it round a moving cutoff. With nothing in **FREQ**, **AUDIO** does nothing.

On **AUDIO** the CV is smoothed above about 5 kHz, so a saw or square drives it cleanly. Exponential FM raises a filter's average pitch as it deepens: a sine an octave deep raises it by about an eighth, two octaves by more than a half. So a filter under FM no longer beats at what its **BEAT** says.




Figure 9. A ringing filter on AUDIO, under a 5 Hz sine of growing depth: how far its average pitch rises.

MODE

MODE chooses the response, drawn round the knob: low-pass, band-pass, where it loads, high-pass and notch. At a modest Q, low-pass, band-pass and high-pass all peak at the cutoff by the same amount. A ringing filter's low-pass and band-pass are the same sine a quarter of a cycle apart; the notch takes the ring out, so past OSC it passes only its input.




Figure 10. The four responses, set to 1 kHz with a Q of 2. Each falls at 12 dB an octave.

Q

Q is the resonance. Up to the heavy **OSC** mark it is a Q, from 0.5 at the left to 1,000 just short of the mark, evenly in ratio. At 50 Hz a Q of 1,000 rings for about six seconds before it falls to a third. Q loads a little before the mark, at about 220. Its tooltip gives the Q, as Q: 218.7.

A band-pass's or low-pass's gain at the cutoff is its Q: +14 dB at a Q of 5, +54 dB at 500. A filter driven so hard that its output would pass about 5 V damps itself, as a limiter in its resonance would, and holds the level rather than clipping; at a high Q this happens with only a few millivolts in. To filter a drone at a high Q without it, keep what goes in quiet: about 5 V divided by the Q.




Figure 11. The band-pass at 1 kHz with four Qs, measured with a small signal. Its peak gain is its Q.

Past the mark the filter rings by itself, building up from silence. Turning further makes the ring start sooner, come up louder, from 2.5 V at the mark to 5 V fully right, and become a little less pure, its third harmonic rising from about

50 dB down to 30. At C2 a ring just past the mark takes about six seconds to build, nine-tenths of the way round about a second, and fully right a third of a second; a higher filter builds faster in proportion. The tooltip gives the level, as Q: oscillating, 3.8 V. However far it is turned, the ring's pitch does not move.




Figure 12. Left, Q's law below the mark. Right, past it: the ring's level, and its third and fifth harmonics against it, measured at C2. The ring's level is against 5 V.




Figure 13. Left, a filter at C2 building up from silence, past the mark. Right, the PING button at full on a filter at C2 below the mark, at three Qs, and on one already ringing, which jumps and settles back.

CV and **AMT** under Q: a CV on the resonance, off at the centre. At full AMT ± 5 V moves Q across its whole travel, so a Glacial lane can take a filter over the mark and back.

Every voice has a faint noise floor, 120 dB down, drawn for each unit and saved with it. A filter past OSC starts ringing from that floor, by itself.

NORMAL

NORMAL, the slide switch in the left column, says what IN 2 and IN 3 hear with nothing in them:

- **CHAIN**, up: IN 2 hears voice 1's filter, and IN 3 voice 2's, so the three are in series. The normal is taken before each voice's VCA, so LEVEL sets only what a voice gives its OUT and MIX.
- **BANK**, the middle, where it loads: IN 2 and IN 3 hear IN 1, so the three filter one input in parallel: three resonances on one sound.
- **FREE**, down: IN 2 and IN 3 hear nothing, so each voice is what its PING and its own resonance make of it.

A cable in an IN breaks its normal in any position. The tooltip says what the switch is doing, as Normal: CHAIN, each IN hears the filter above.

Levels

Each voice ends in its own VCA, as Littoral's do. **LEVEL** loads at 80% and follows a square law. Its **CV** and **AMT** move it across its whole travel at ± 5 V, full AMT; the CV is smoothed only enough that a step does not click, so a tremolo still comes through. **OUT** is the voice after its VCA. The filter is linear to ± 5 V in and out, and rounds off gently beyond, towards ± 7 V; a struck filter that is already ringing can reach that knee for a moment.

The red lamp beside each voice's number is lit while the voice runs. A voice runs only while something hears it: its OUT or MIX is patched and its LEVEL is up or its LEVEL CV patched, or, on CHAIN, the voice below it is running and hears it. A voice that is not running costs nothing; when it comes back it starts from silence, so a ringing voice swells back in.

Drift

DRIFT works as Littoral's does: a shared drift on the reference and a smaller one for each filter, a few cents and about a cent at its centre, doubled fully right, and none fully left, where there is no warm-up either. Its tooltip reads *Drift: 50%*. A Fluvial starts cold, 8 cents flat and rising to pitch over twenty minutes or so. One whose REF IN is patched takes its shared drift from the cable, so a Littoral and a Fluvial chained together drift as one cabinet.

Outputs

REF, on the vermilion block: the reference, as 1 V/oct, with V/OCT and the shared drift on it. **MIX**: the three voices after their VCAs, a third each.

Jacks

Jack	Direction	What it carries
V/OCT	in	added to the reference, 1 V/oct
REF IN	in	the reference, replacing REF, 1 V/oct from C2
IN	in	the audio the filter works on, linear to ± 5 V
PING	in	strikes the filter on each rise above 0 V, as hard as it is high
FREQ	in	the cutoff, 1 V/oct at full AMT; every sample on AUDIO
CV (BEAT)	in	the filter's beat, 1 Hz a volt at full AMT
CV (Q)	in	the resonance, ± 5 V across Q's travel at full AMT
CV (LEVEL)	in	the voice's level, ± 5 V across LEVEL's travel at full AMT
OUT	out	the voice, after its VCA
MIX	out	the three voices, a third each
REF	out	the reference with V/OCT and the shared drift, 1 V/oct from C2

Every input reads 0 V when nothing is patched.

Saving

The knobs and switches are saved by Voltage Modular. Each Fluvial also saves the course of its drift and the noise its rings start from, so a patch sounds the same each time it loads.

Notes


- **Ringings filters cross-patched on AUDIO lock.** Two ringing filters, each one's OUT into the other's FREQ, or a Fluvial voice and a Littoral voice patched the same way, settle into a steady tone over a range of settings, and turning one's pitch hardly moves it. Past some point the lock breaks and the sound turns chaotic. Small moves near that edge are the interesting ones.

3 • Calderal

calderal /kɑːlˈdɛrəl/

adjective of a caldera: the wide basin left when the summit of a volcano falls into the emptied chamber beneath it.

Calderal is two ring modulators side by side, made for drones. Each multiplies one signal by another, cleanly: two sines in give their sum and their difference and nothing else. Each can keep just one of the two sidebands, which makes it a frequency shifter, moving every partial of a sound by the same number of hertz. Each has a carrier of its own for when nothing is patched into B, a feedback path, a blend from dry to ring, and a VCA. There are no gates and no envelopes: everything moves continuously, so everything can move slowly.



- | | |
|--------------|-----------------------------------|
| 1 A IN | the ring's first side |
| 2 B IN | the second side; SHIFT when empty |
| 3 DRIVE | the multiplier's soft limit |
| 4 SIDE | DIFF, ring, SUM |
| 5 SHIFT | the carrier's rate in hertz |
| 6 CV | on SHIFT |
| 7 AMT | SHIFT's CV's amount |
| 8 WIDE, FINE | SHIFT's range |
| 9 BLEND | DRY, AM, RING |
| 10 CV | on BLEND |
| 11 AMT | BLEND's CV's amount |
| 12 BLEED A | A's leak, in or out of phase |
| 13 BLEED B | B's leak |
| 14 LEVEL | the channel's VCA |
| 15 CV | on LEVEL |
| 16 AMT | LEVEL's CV's amount |
| 17 CV | on FEEDBACK |
| 18 AMT | FEEDBACK's CV's amount |
| 19 FEEDBACK | the ring back into A |
| 20 OUT 1 | the channel |

The two channels

The two channels are the two halves of the face, with the same controls; the face shows channel 1. Only the outputs are numbered on the face, **OUT 1** and **OUT 2**; in the tooltips channel 2's controls and jacks carry a 2, as Shift 2 or A 2.

With nothing in channel 2's **A IN**, channel 2 hears OUT 1, after channel 1's VCA. The line on the face from OUT 1 up between the channels is that connection. A drone can be shifted and then ringed, or shifted twice. A cable in channel 2's **A IN** breaks the normal, and the two channels are then independent.

A channel nobody hears is not computed. Channel 2 runs while OUT 2 is patched; channel 1 while OUT 1 is patched, or while channel 2 hears it through the normal.

The inputs

A IN and **B IN** are the two sides of the multiplier. They are interchangeable as far as the ring is concerned; A is the side BLEND fades from and FEEDBACK goes back into, and B is the side SHIFT stands in for. Both are on the series' ± 5 V scale: two full-scale inputs make a full-scale ring.

Two voices of one drone, one in A and one in B, do not beat. Two of Littoral's voices a fraction of a hertz apart give a steady tone an octave up and a difference too low to hear, which is filtered out. The beats come back three ways: from the voices' third and fifth harmonics, whose products beside the octave beat at two and four times the rate; from two beating voices in A against a third in B, whose sidebands keep the beat and move it in pitch; and from BLEND's dry side.

DRIVE sets how hard both sides push into the multiplier's soft limit. Fully left it is clean; where it loads, a fifth of the way, it is nearly so, a full-scale sine's third harmonic 62 dB down. Turned up, loud peaks round off, a sine gains its own third harmonic, and the ring grows faint sidebands of its own: at half way the third harmonic is 34 dB down, fully right 15 dB. A full-scale signal's peaks stay where they were, but quieter signals come up: at full DRIVE a half-scale sine is 4.5 dB louder. Its tooltip reads Drive: 20%.




Figure 14. DRIVE: the third harmonic a 100 Hz sine gains, against its fundamental.

SIDE chooses which sidebands come out. At its centre it is a ring modulator: the sum of A's and B's frequencies and their difference, each at half strength. Turned right towards **SUM**, the difference fades out, until fully right only the sum is left, at full strength. Turned left towards **DIFF**, the same for the difference. The tooltip reads Side: ring, sum and

difference, Side: 75% sum, Side: sum only.

With SIDE at an end, a sine at 200 Hz in A and a sine at 30 Hz in B give 230 Hz alone on SUM, or 170 Hz alone on DIFE. Every partial of A moves by B's frequency: that is a frequency shifter. The sideband left behind is more than 51 dB down across A's range when SHIFT is the carrier. A sine in B does nearly as well, at least 46 dB down from about 7 Hz up, but not below: at 4 Hz it is only 27 dB down, and at 1 Hz there is hardly any separation. For slow shifts, use SHIFT.

DRIVE acts before the sidebands are separated, so its harmonics are separated with them, and turning DRIVE up does not bring the other sideband back.




Figure 15. The sideband SIDE leaves out, fully right on SUM, for a 1 kHz sine in A.

SHIFT

With nothing in **B IN**, B hears **SHIFT**: a carrier of its own, at a rate in hertz, made so that it shifts cleanly at any speed, through zero and out the other side.

SHIFT sets the rate. At its centre the rate is exactly zero and the carrier stands still at +5 V, so the ring is A unchanged and the channel is a VCA. Either side of the centre it is very fine, as BEAT is, and widens towards the ends. Right of centre shifts up on SUM; left of centre is a negative rate, which shifts down. Its tooltip gives the rate to a thousandth of a hertz, as Shift: +0.200 Hz, or Shift: off (VCA) at the centre; click it and type a rate to set it exactly.

The small switch beside SHIFT's AMT sets the range: **FINE**, down, up to 5 Hz either way; **WIDE**, up, up to 500 Hz either way.

CV and **AMT** beside SHIFT: a CV on the rate, off at the centre. At full AMT each volt adds a hertz on FINE, or 100 Hz on WIDE, so ± 5 V moves the rate as far as the knob does; left of centre inverts it. The rate runs through zero and on, so a slow Glacial lane into it shifts a drone up, stops it, and shifts it down again. The tooltip shows the knob's rate, not the CV's.




Figure 16. SHIFT's law on WIDE. FINE is the same curve a hundredth as tall. The centre is a flat zero, the channel a VCA there.

A cable in B IN takes SHIFT out of play. SHIFT, its AMT and its switch grey out, and the line from the SHIFT row into B IN shows why. SHIFT's tooltip then reads Shift: off, B is patched.

When the rate is exactly zero, the carrier glides back to rest at +5 V over about a second, so a stopped shift always leaves A at its full level and in phase. A CV sweeping through zero passes through it without stopping the carrier. Where the carrier stands is saved with the patch.

What SHIFT is for. Shifts of a few hertz and up, on WIDE, make a harmonic drone inharmonic: every partial moves by the same number of hertz, so they no longer keep their ratios, and the tone turns metallic, bell-like or glassy as the shift grows. A shift leaves the distances between partials alone, so a beating drone shifted beats exactly as before. A slow shift on its own is only a tremolo: on the ring, at 0.2 Hz, a swell that dies to silence and comes back inverted (Figure 17). With BLEND at AM and SIDE on SUM, the drone and a copy 0.2 Hz above it beat together, the whole drone rising and falling.




Figure 17. A Littoral sine at A1 in **A IN**, nothing in B, **SHIFT** at +2 Hz on FINE and SIDE at its centre. With **BLEND** on **RING**, the ring dies to silence four times a second, and between the silences the tone is turned upside down. With **BLEND** at **AM**, the dry tone is added at equal power: the sum swells and sinks almost to nothing twice a second, and is never inverted.




Figure 18. A Littoral sawtooth, and the same shifted by +30 Hz on WIDE with SIDE on SUM. Each partial moves 30 Hz at its own level; they are no longer multiples of one frequency.

BLEND and the bleeds

BLEND, the channel's large knob, crossfades from the dry input to the ring. Fully left, at **DRY**, it is A alone. Fully right, at **RING**, where it loads, it is the ring alone. Halfway, at **AM**, it is A and the ring at equal power: with **SIDE** at the centre and B a sine, that is amplitude modulation, A swelling and shrinking with B. Its tooltip reads Blend: 50% ring. The dry side is delayed as the ring is, so mixing the two does not comb.

CV and **AMT** under **BLEND**: a CV on the blend, off at the centre; at full **AMT** ± 5 V moves it across its whole travel.

BLEED A and **BLEED B**, the two small trims. An analogue multiplier lets a little of each input through to its output, and trims are fitted to null it, never perfectly. Here each trim is nulled at its centre, lets a little of its input through in phase to the right and inverted to the left, and is very fine near the null. Fully over, each lets through 12 dB under its input. They load a little either side of their nulls, about 43 dB down, as a trimmed multiplier is left. The tooltips read Bleed A: -43 dB, Bleed B: -43 dB, inverted, Bleed A: nulled. **BLEED A** puts a little of the plain tone under the sidebands; with nothing in B, **BLEED B** lets **SHIFT**'s own carrier through.

FEEDBACK

FEEDBACK sends the ring back into A, ahead of **DRIVE**. Fully left there is none. Each pass shifts what was already shifted, so every partial grows a ladder of copies spaced **SHIFT** apart, each fainter than the last. On **SUM** the ladder climbs; on **DIFF** it descends; on the ring it spreads both ways. At 0.2 Hz the ladder is a slowly beating cluster round each partial; at 50 Hz on **WIDE** it is an inharmonic spectrum climbing from every partial of the drone.

The loop has a soft limit at ± 5 V, so it stays bounded at full **FEEDBACK**, whatever comes in. A loud input reaches that limit, which shortens the ladder and adds faint products of its own; a quiet one makes a long ladder, a quarter-

volt sine's rungs falling only half a decibel each at full. Two high-pass stages at 20 Hz in the loop keep a descending ladder that folds through zero from rumbling, and the top rolls off from about 6 kHz.

The loop runs through the stages that separate the sidebands, and they delay each frequency by a different amount. So with SHIFT at zero, FEEDBACK colours the sound as a resonance would: at full, peaks of 15 to 20 dB near 90 Hz and 1.4 kHz, and shallow dips between.

Its tooltip reads Feedback: 60%, or Feedback: off. **CV** and **AMT**, FEEDBACK's pair on the bottom row: a CV on the feedback, off at the centre, across its whole travel at ± 5 V and full AMT.




Figure 19. A full-scale 220 Hz sine, SIDE on SUM, SHIFT +50 Hz on WIDE: the ladder at 270, 320, 370 Hz and up, with the loop's limit shortening it at full.

The loop is about a sample long, so the whole ladder sounds at once. For a spiral in time, in which each copy arrives a step after the last, put a delay in the loop: Alluvial, patched round the channel (Chapter 6).

The VCA

Each channel ends in its own VCA. FEEDBACK is taken before it, so LEVEL does not change the loop.

LEVEL sets the channel's level. It loads at 80%, which on its square law is -3.9 dB. Its tooltip reads Level: 80%. **CV** and **AMT** beside it: a CV on the level, off at the centre, across LEVEL's whole travel at ± 5 V and full AMT, smoothed only enough that a step does not click.

OUT 1 and **OUT 2**, on the vermillion blocks: each channel after its VCA.

Jacks

Each channel has its own set.

Jack	Direction	What it carries
A IN	in	the ring's first side, ± 5 V full scale; channel 2's hears OUT 1 when empty
B IN	in	the ring's second side, ± 5 V full scale; SHIFT's carrier when empty
CV (SHIFT)	in	the rate, 1 Hz a volt at full AMT on FINE, 100 Hz on WIDE
CV (BLEND)	in	the blend, ± 5 V across BLEND's travel at full AMT
CV (LEVEL)	in	the level, ± 5 V across LEVEL's travel at full AMT
CV (FEEDBACK)	in	the feedback, ± 5 V across FEEDBACK's travel at full AMT
OUT 1, OUT 2	out	each channel, after its VCA

Every other input reads 0 V when nothing is patched.

Notes


- **Calderal tunes nothing.** It has no V/OCT and no ratios; the pitches are set at the sources.

4 • Alluvial

alluvial /ə'lu:viəl/

adjective of or made of alluvium: the clay, silt and sand that flowing water lays down, layer over layer, on a river's floodplain or delta.

Alluvial is a stereo tape loop for drones: a loop of up to a minute, with an erase head, a record head and a play head, and a little loss on every pass. Feeding what a tape plays back into what it records is one of the oldest ways to make a drone. Each layer is laid over what is already there, and each pass wears what came before a little further: darker, softer, quieter, a little noisier, until it settles into the tape or fades out under what came after.



- | | | | | | |
|--------------------|-----------------------|--------------------|---------------------------|-------------------|---------------------|
| 1 IN L | the left track | 2 IN R | the right track | 3 REC | records while high |
| 4 Button | records, a toggle | 5 CLEAR | wipes the tape, on a rise | 6 Button | wipes it by hand |
| 7 STOP | stops while high | 8 Button | stops, a toggle | 9 OUT L | the left track |
| 10 OUT R | the right track | 11 SPLICE | a trigger once a loop | 12 Display | the whole loop |
| 13 LENGTH | 1 to 60 s | 14 TAP | length from a clock | 15 SPEED | ½x to 2x |
| 16 GLIDE | the capstan's inertia | 17 SNAG | damage drags the tape | 18 CV | on SPEED |
| 19 AMT | SPEED's CV's amount | 20 WOW | slow wavering | 21 FLUTTER | fast wavering |
| 22 REV, FWD | the direction | 23 TRIG | turns the tape round | 24 INPUT | recording level |
| 25 ERASE | the erase head | 26 CV | on ERASE | 27 AMT | ERASE's CV's amount |
| 28 GAP | play head's distance | 29 FEEDBACK | playback laid back on | 30 MIX | DRY to TAPE |

31	HIGHS	bias loss a pass	32	SAT	saturation	33	HISS	the tape's noise
34	GRAIN	modulation noise	35	PRINT	print-through	36	WEAR	NEW to OLD, all losses
37	DROPOUT	starting damage	38	DISINTEGRATE	how fast it grows			

Tape lost a little of its sound every time it passed the heads, and the losses built up, pass on pass. Alluvial models that loss as tape does it: the losses are written back onto the tape, not put on the playback, so they accumulate. Every loss is counted in passes, not seconds, so a short loop wears faster than a long one: a two-second loop goes round thirty times a minute, a sixty-second loop once.

The machine

The column on the left is the machine itself.

IN L and **IN R**: the input, ± 5 V full on the tape. With nothing in IN R, IN L feeds both tracks.

REC records. The red button is a toggle, lit while recording; the **REC** jack records while it is above 0 V; the tape records while either is on. Punching in and out crossfades over a few milliseconds, so neither clicks. With REC off the erase and record heads are off, and the loop plays as it is.

CLEAR fades the whole tape to silence over a quarter of a second, as a bulk eraser would, and the tape comes back blank, its damage back to a fresh start. The small button does it by hand; the **CLEAR** jack on each rise above 0 V.

STOP brakes the motor: the tape slows to a halt, its pitch falling, and its sound fades as it nears a stop. Turn STOP off and the tape runs up to speed again. The button is a toggle; the **STOP** jack stops the tape while it is above 0 V; the tape stops while either is on.

OUT L and **OUT R**, on the vermilion block: the two tracks, as MIX sets them against the input. The right track's heads sit an eighth of a millisecond after the left's, as heads slightly out of line, which widens the sound.

SPLICE: a trigger, +5 V for one millisecond, each time the splice passes the heads, forwards or backwards: once a loop. Patch it into Glacial's RESET and the drum starts again with the loop.

The tape is not saved with the patch. A patch loads on blank tape. The knobs and switches are saved, as on any module.

The display

The display along the top is the whole loop, running right to left past fixed heads. The red line, **REC**, is the record head, with a tick for the erase head just before it. The black line, **PLAY**, is the play head, GAP's distance behind it. The dashed line is the splice, travelling with the tape.

The waveform is each part of the tape's level as the play head would hear it: the peak in a pale ink, the average inside it in full ink, on a decibel scale down to 48 dB below full, so a fade over hundreds of passes stays visible. Damaged tape is washed in red.

TAPE

LENGTH sets how long the loop is, from 1 second fully left to 60 seconds fully right, evenly in ratio. It loads at about 21.6 s. A new length is spliced in when the tape next comes to the splice, never by moving a head, so turning LENGTH does not click. What lies beyond a shortened loop stays on the tape, unplayed until the loop is long enough to reach it again. Its tooltip reads Length: 21.6 s.

TAP sets the length from a clock: the time between two rises above 0 V, from 1 to 60 seconds, spliced in as any new length is. Until a second rise, and whenever TAP is empty, LENGTH sets the loop. LENGTH's tooltip then shows both, as Length: 8.00 s from TAP (the knob: 21.6 s).

SPEED is varispeed: from $\frac{1}{2}x$ fully left through 1x at the centre to 2x fully right. It moves the recording as well as the playback, so material recorded at half speed plays an octave up at 1x, and layers laid at different speeds sit at

different pitches. Slower tape holds less of the top: at 1x and WEAR's centre the tape holds up to about 16 kHz, at $\frac{1}{2}x$ an octave less; at OLD, half that. Its tooltip reads Speed: $\times 1.000$.

GLIDE sets how slowly the tape follows a change of speed, as the inertia of a capstan: from 5 ms fully left to 5 seconds fully right. It loads at about 40 ms. The pitch glides as the speed changes, and GLIDE also slows the tape's stop, start and turn.

CV and **AMT** beside **SPEED**: a CV on the speed, off at the centre, for vibrato and chorus rather than transposing: at full AMT ± 5 V swings the pitch half a semitone either way, and a quarter of the way out it is ± 3 cents. AMT's tooltip gives the depth, as Speed CV: +12.5 cents at 5 V.

WOW and **FLUTTER**: the tape's speed wavering, as an eccentric reel and a capstan made it waver. Wow is slow, half a cycle to two cycles a second, its depth wandering between half and all of its setting from cycle to cycle. Flutter is faster, 6 to 16 a second, with a little scrape in it. WOW's most is $\pm 3\%$ of the speed, FLUTTER's $\pm 1\%$, and the first half of each knob is the range a good machine lives in. They load low: wow $\pm 0.12\%$, flutter $\pm 0.023\%$. Both tracks share them. The reel and capstan turn with the tape, so at $\frac{1}{2}x$ both waver an octave slower, and as STOP brakes the tape they slow with it.




Figure 20. The tape's speed against the motor's. Where the two knobs load, the wavering is $\pm 0.14\%$ at most, a fraction of a cent.

SNAG sets how much damaged tape drags on the heads. Where the tape is damaged it is held back a moment, and its pitch dips. Fully left, damage only loses sound. Where it loads, bare tape drags the speed 8%, about a semitone and a half; fully right, 16%. Its tooltip reads Snag: bare tape drags 8.0%, -1.44 semitones.

REV and **FWD**: the tape's direction. Turning round is not instant: the tape brakes, stops and runs up the other way, the pitch falling to nothing and rising again, in about half a second at 1x, longer with GLIDE slow. The **TRIG** jack under the switch turns the tape round on each rise above 0 V, from whichever way the switch has it.




Figure 21. Left, STOP held from 0.5 s to 2.5 s. Right, TRIG at 0.5 s turning the tape round.

HEADS

INPUT sets how hard the input is recorded: off fully left, the input at its own level at the centre, and four times, +12 dB, fully right. It loads at +5.8 dB. Fresh input meets a soft limit on its way onto the tape, from 7.5 V towards a ceiling of 10 V.

ERASE sets how much of each pass the erase head wipes, while recording. Lows sit deeper in a tape's coating and are harder to erase, so ERASE takes the highs faster than the lows, about twice as many decibels a pass: what survives is darker as well as quieter. Fully left nothing is erased, and a loop recorded over keeps every layer; the bottom of the travel is fine, a fraction of a decibel a pass; fully right the tape is wiped each pass and the loop is a plain delay. Its tooltip gives the loss a pass at 1 kHz and how long that takes to bring a sound down 60 dB at the loop's length and speed, as Erase: -0.19 dB a pass, 60 dB in 1:53:54, where it loads.

CV and **AMT** beside ERASE: a CV on the erase, off at the centre, across its whole travel at ± 5 V and full AMT.




Figure 22. What the erase head takes each pass, by frequency, with WEAR at NEW so that it acts alone.

GAP sets how far the play head sits behind the record head, as a share of the loop: on the record head fully left, round the whole loop to meet it again fully right. It loads at half. A share rather than seconds, so the echoes **FEEDBACK** makes fall on the same divisions of the loop whatever **LENGTH** is: a quarter is always four to the loop. Moving **GAP** moves the head as a head would be moved, the pitch bending as it goes. Its tooltip reads Gap: 50% of the loop, 10.78 s.

FEEDBACK lays what the play head hears back on at the record head, while recording, and what is already there gives way by as much, so the loop never runs away at any setting. It loads off. With **ERASE** fully right it is a tape echo, echoes **GAP** apart, each **FEEDBACK**'s share of the last. At full, the play head's sound replaces the loop, which becomes **GAP** long.

MIX, from **DRY** to **TAPE**: how much of **OUT L** and **OUT R** is the tape, and how much the input. It loads on **TAPE**.

WEAR

Each knob in this row is a loss or a noise on the tape. They act on each recording pass; with **REC** off, playing wears the tape too, by a tenth as much, without saturation or recorded noise.

HIGHS: what the record head's bias takes from the top of the band each pass, as a shelf above about 3 kHz. It compounds, so the corner of what is left falls as the loop ages. There is a very small loss across the whole band as well, which lets a loop that is never erased fade at all, and the record path's high-pass at 2 Hz takes a little from the lowest notes each pass. Its tooltip reads Highs: -0.60 dB a pass where it loads.




Figure 23. What survives on the tape after recording passes with nothing new coming in, ERASE at zero and the WEAR row where it loads. The top goes first; 1 kHz loses about a tenth of a decibel a pass, 100 Hz half that.

SAT: the tape's saturation, on what is already on the tape, each recording pass. Up to its knee the tape keeps what it holds exactly; above it, it bends softly towards a ceiling. Turning SAT up lowers the knee a little, from 5 V, and brings the ceiling down hard. Where it loads, the knee is 4.7 V and the ceiling 7.6 V. Its tooltip reads Sat: from 4.7 V, at most 7.6 V.

HISS: both of the tape's noises. Playback hiss is the tape's own grain, heard and never written; HISS sets it even on new tape, at 60 dB under a volt where it loads, Hiss: -60 dBV. Record noise is written on each recording pass and worn like everything else, so it settles where what each pass adds equals what each pass takes, within about a hundred passes, as a dark wash under the loop: where the knobs load, about 50 dB under a volt, louder than the playback hiss.

GRAIN: modulation noise. A tape's surface is never quite even, so a tone on it carries a faint fuzz that moves with it. It is heard on playback and written at half the depth on each recording pass, so it builds generation by generation. Its tooltip reads Grain: -41 dB of noise on the signal.

PRINT: print-through. On a reel, the layers of tape lie against each other, and each copies a faint ghost of itself onto its neighbours. Here each recording pass copies a soft, faint echo of the material a second of tape either side, or a third of the loop on loops shorter than three seconds; the copy heard before the material is a little louder than the one after. Its first half is the faint ghost real tape has. Its tooltip reads Print: -62 dB a pass.

WEAR, from **NEW** to **OLD**, scales the four above and the damage below together. At **NEW** the tape is perfect whatever they are set to: nothing is lost and nothing written, and only the playback hiss remains, at half. At its centre each is what its own knob says. At **OLD** each is four times that. Its tooltip reads Wear: 50%, the losses x1.00.

WEAR also wears the machine. The play head's contact with the tape is never quite steady, so the level sways slowly, the same on both tracks: about a quarter of a decibel at the centre, around a decibel at **OLD**, none at **NEW**. What **FEEDBACK** and recording over lay back down carries the sway, so it builds in a loop recorded over. And the capstan's flutter grows lopsided, its cycle no longer a plain sine, though never wider than **FLUTTER** sets.




Figure 24. A 1 kHz tone recorded once, then recorded over with nothing new coming in, at five settings of WEAR and ERASE at zero. On a 20-second loop, 300 passes is 100 minutes.

Damage

DROPOUT: how damaged the tape starts. A scatter of spots, 5 to 50 ms long, where the oxide is thin and the play head reads the tape quieter and duller. The spots are on the tape, not in time: they fall in the same places every pass and become part of the loop's rhythm. Low on the knob they are few and faint; turning it up brings more and deepens them. Where it loads there are about five a minute of tape, Dropout: 4.8 spots a minute of tape.

DISINTEGRATE: how fast the damage grows. Each time a damaged spot passes the heads it deepens a little and spreads into its neighbours, so holes widen pass by pass, and now and then a new speck appears. At full, with WEAR at its centre, a loop is more than half lost in about 25 passes; at a quarter, after 150 passes, about a quarter of it. It loads off, and the damage then holds. It grows whether recording or not.

Sound lost in a hole is lost once: what is recorded over it later is lost the same way, but a hole does not take more each pass it has been there. CLEAR puts the tape back to its starting damage, a fresh scatter with none of what grew.




Figure 25. How much of the loop is damaged to half its level or worse, pass by pass, with DROPOUT and WEAR where they load.

Jacks

Jack	Direction	What it carries
IN L	in	the left track, ± 5 V full on the tape
IN R	in	the right track; IN L feeds both when it is empty
REC	in	records while above 0 V, as the REC button does
CLEAR	in	fades and wipes the tape on each rise above 0 V
STOP	in	stops the tape while above 0 V
TAP	in	the loop's length, from the time between two rises above 0 V
CV (SPEED)	in	the speed, ± 50 cents over ± 5 V at full AMT
CV (ERASE)	in	the erase, ± 5 V across ERASE's travel at full AMT

TRIG	in	turns the tape round on each rise above 0 V
OUT L	out	the left track, against the input as MIX sets it
OUT R	out	the right track, the same
SPLICE	out	+5 V for 1 ms each time the splice passes the heads

Every input reads 0 V when nothing is patched.

Notes

- **A loop never runs away.** FEEDBACK gives way rather than adds, the erase and the bias never gain at any frequency, and SAT's ceiling holds the rest. A loop recorded over for an hour with nothing erased stays bounded.
- **Recording over a loop with ERASE at zero still wears it:** the record head's bias is most of the darkening in sound-on-sound. HIGHS sets how much.
- **Stopping and reversing are heard.** The tape slows and speeds up as a motor's tape does, and its sound falls away as it nears a stop.


5 • Glacial

glacial /ˈgleɪf(ə)l/

1 *adjective* of ice, or of glaciers.

2 *adjective* so slow as to seem not to move at all, as a glacier advances.

Glacial is paper wrapped round a slowly turning drum. The paper has six lanes; points placed in each lane make a curve, and as the drum turns, each lane's output follows its curve. One turn takes anything from ten seconds to two hours, so a lane can open a filter over twenty minutes, move a beat across an hour, or bring a voice in and out once a turn for as long as the patch runs.



- 1** **RATE** one turn: 10 s to 2 h
- 2** **ZOOM** how much of the turn is shown
- 3** **TURN** runs the motor
- 4** **REV, FWD** the motor's direction
- 5** **GATE** runs the motor while high
- 6** **RESET** to the seam, on a rise
- 7** **Button** to the seam, by hand
- 8** **RATE** the speed, 1 V/oct
- 9** **AMT** RATE's amount
- 10** **PHASE** moves the drum, a turn per 5 V
- 11** **AMT** PHASE's amount
- 12** **SEAM** a trigger once a turn
- 13** **PHASE** where the drum is, 0 to 5 V
- 14** **The paper** six lanes, and the playhead
- 15** **RANGE** BI or POS, each lane
- 16** **OUT** each lane, 1 to 6

The paper

The paper runs down the module. The lanes are its columns, numbered **1** to **6** beneath it. A point's place across its lane is its value: the lane's left edge is the bottom of its range, its right edge the top. A point's place down the paper is its place on the drum.

The red line a third of the way down is the playhead. What is under it is what the outputs are playing now. As the drum turns forwards the paper scrolls up, so what is coming is always below the playhead and what has passed is above it.

The paper has no beginning and no end. It joins itself at a seam, the dashed line, and each time the seam crosses the playhead the drum has made one turn. Every lane is a loop: its last point joins its first across the seam, and a curve that crosses the seam is as smooth as one that does not (Figure 29).

The darker strips beside and between the lanes are the drum's edge. Their ticks move as the drum turns, so it can be seen turning with nothing drawn.

A point's place is a fraction of a turn, not a time. Change **RATE** and the whole score stretches or squeezes, as it would on a drum turned faster or slower; nothing has to be redrawn. A lane with no points plays 0 V, and a lane with one point plays that point's value all the way round.

Drawing

Do this	Where	And
click	an empty place in a lane	a point appears there, held by the pointer
drag	a point	across sets its value; up and down its place on the paper
shift-click	a point	it is deleted
double-click	a point	it is deleted
right-click	a point	the curve leaving it takes the next shape
drag	the drum's edge	the drum turns by hand (see below)

The pointer shows what a press will do: a four-way arrow over a point, with a small square round the point; a pointing hand over an empty place in a lane; a dragging hand over the drum's edge. While a point is dragged it shows its value and its time, as -2.50V $0:42$: -2.5 V on the lane's range, 42 seconds into the turn at RATE's setting. A point stays between its neighbours. Dragged past the next one, it stops just short, so the order of the points round the drum never changes. It can be dragged across the seam if its neighbours allow. A lane holds up to 256 points.

While the paper is edited, the outputs glide. A point dragged across a lane moves its output in small steps, and a step on a filter or a VCA is heard as a click. So while points are added, moved, deleted or reshaped, the outputs follow the paper through a short smoothing, and settle a fifth of a second after the last edit. Away from an edit they follow the paper exactly, and a step segment steps.




Figure 26. A point dragged from -2 V to $+3\text{ V}$ in five steps 40 ms apart. The output rounds each step off and trails the paper by about 40 ms; it is within a millivolt of it 0.1 s after the last step.

The shape of a curve

Each point sets the shape of the curve leaving it, towards the next point. Right-clicking a point moves it on to the next shape:

Shape	How it moves
S-curve	eases out of one point and into the next, so the lane has no corners. A new point starts with it.
Linear	a straight line, at a steady rate.
Exponential	slow at first and fast at the end, rising or falling.
Step	holds the point's value until the next point, then jumps to it.

The S-curve is where each point starts because a slow straight sweep turns a corner at each point, and on a pitch or a filter the corner is heard.




Figure 27. One lane through the same three points, in each shape. The step reads 0 V before the first point: it is still holding the last point's value, from across the seam.



Figure 28. The same lanes as Glacial draws them, by its own code, the whole turn in view: S-curve, linear, exponential and step, in lanes 1 to 4; lanes 5 and 6 are blank, at 0 V. The drum stands an eighth of a turn past the seam, the dashed line above the playhead.




Figure 29. A lane over a turn and a half. The curve from its last point to its first crosses the seam without a corner.

RATE, ZOOM and TURN

RATE sets how long one turn takes: from 10 seconds fully left to 2 hours fully right. Each part of its travel covers the same ratio of times. It loads at one minute. Its tooltip reads Rate: 1:00 a turn.

ZOOM sets how much of the drum the paper shows: from an eighth of a turn, fully left, to the whole turn, fully right. It loads at a quarter of a turn, Zoom: 0.25 of a turn. **ZOOM** changes only what is shown, and how far a drag on the drum's edge turns it; the outputs are the same at any zoom.

TURN runs the drum's motor. While it is lit the drum turns at **RATE**; when it is not, the drum stops where it is. **REV** and **FWD**, the switch beside it, set the motor's direction. It loads on **FWD**, which scrolls the paper up. On **REV** the drum turns backwards, an exponential plays fast and then slow, and a step holds the value of the point above the playhead.




Figure 30. **RATE**'s law: the time for one turn. It loads at about 0.27 of its travel, one minute.

Turning by hand

Press on the drum's edge, beside or between the lanes, and drag. The drum follows the hand: drag up and the paper goes up, as if pushed round. The outputs follow too, so turning by hand plays the lanes at whatever speed the hand moves. A drag turns the drum further at a wide **ZOOM** than a narrow one: one height of the paper is always the span it shows.

While the hand is on the drum the motor is ignored. Let go, and the drum keeps the speed the hand had, and eases back to the motor's speed, or to a stop if **TURN** is off, with a time constant of a second. A flick sends it spinning.




Figure 31. The drum turned by hand at half a turn a second, then let go. It keeps the hand's speed, averaged over its last 50 ms, and eases back to the motor's with a time constant of one second: about a third of the way there after a second, and nearly all by three.

Inputs

GATE runs the motor while it is above 0 V, as **TURN** does: the motor runs while **TURN** is lit or **GATE** is high. Patch a gate from a sequencer or a clock so that the drum turns only when something else is playing.

RATE, the input, changes the drum's speed through its **AMT**: at full right, each volt up doubles the speed and each volt down halves it, up to five octaves either way. A one-minute turn goes as fast as 1.9 seconds or as slow as 32 minutes; at the ends of **RATE**'s knob, from 0.3 seconds to 64 hours. **AMT** loads at its centre, where the input does nothing, and inverts it left of centre.

RESET puts the drum back at the seam on each rise above 0 V, and the button beside it does the same by hand. Either one fires **SEAM**. The drum turns on from the seam even if the jack stays high; while it is high, the button does nothing more.

PHASE, the input, moves the drum from another module, through its own AMT. It adds to wherever the motor and the hand have put the drum. At full right, 0 to +5 V takes the drum a whole turn forwards and 0 to −5 V a whole turn back. With **TURN** off, **PHASE** alone plays the drum: a slow LFO into it scans the lanes back and forth; a ramp plays them through once per ramp. **RESET** does not undo **PHASE**: it puts the drum at the seam, and **PHASE** still adds to it.

Each CV is limited after its AMT, so at full AMT the **RATE** input moves the speed five octaves and **PHASE** one turn, whatever the voltage; at half AMT, ± 10 V reaches the same limits.

Outputs

OUT, lanes 1 to 6, each under its lane, with its **RANGE** switch above it:

- **BI**, up, where every lane loads: −5 V at the lane's left edge to +5 V at its right.
- **POS**, down: 0 to +5 V.

SEAM gives a trigger, +5 V for one millisecond, each time the seam passes the playhead: once a turn while the motor runs, in either direction, by hand, or because **PHASE** moved the drum across it. **RESET** and its button fire it too, and so does any jump of more than half a turn. Use it to clock something once a turn, or to reset another Glacial so the two stay in step.

PHASE, the output, is where the drum is, from 0 V at the seam to +5 V a turn later, including what the **PHASE** input adds. It rises once a turn going forwards and falls going backwards. It is on the scale **PHASE** in takes at full AMT, so one Glacial's **PHASE** out into another's **PHASE** in, with the second's **TURN** off, locks the second drum to the first: twelve lanes on one drum.




Figure 32. **PHASE** out over two turns from the seam, forwards and backwards. Each drop or jump is the seam passing the playhead, where **SEAM** fires. Starting backwards from the seam fires **SEAM** at once.

Jacks

Jack	Direction	What it carries
GATE	in	runs the motor while above 0 V
RATE	in	the drum's speed, 1 V/oct at full AMT
RESET	in	the drum to the seam, on each rise above 0 V
PHASE	in	the drum's position, a turn either way across ± 5 V at full AMT
SEAM	out	+5 V for 1 ms each time the seam passes the playhead
PHASE	out	the drum's position, 0 to +5 V round a turn
1 to 6	out	each lane, −5 to +5 V on BI , or 0 to +5 V on POS

Every input reads 0 V when nothing is patched.

Saving

The paper is saved with the patch: every point, its value and its shape, in every lane. It is saved with each of the patch's variations, so one Glacial can hold several scores, chosen with the variation selector.

Notes

- **Nothing moves:** the motor runs while TURN is lit or GATE is high. If the drum turns but an output is flat, that lane is empty. If PHASE seems to do nothing, its AMT is at the centre.
- **Very slow is the point.** At two hours a turn the paper moves about one pixel every few seconds. The PHASE output shows the drum creeping on, and the ticks on the drum's edge show it, given time.
- **Points very close together** are allowed. If two are closer than the drum moves in one sample, the lane jumps from one to the next.
- The time shown while dragging a point is at RATE's knob; it does not include the RATE input.

6 • Drones

Each patch below starts from every control where it loads. The drawing shows the cables, and the numbered controls are the ones to set, as the key beneath it says. Cables that leave the drawing go to a mixer or to the rack's outputs. Each patch ends with what to listen for, and where to take it next.

Most of these patches take minutes to show what they do. Leave them running.

Three voices beating



① BEAT, voice 2 +0.100 Hz


② BEAT, voice 3 -0.150 Hz

③ DRIFT fully left, for now

Littoral alone, its three voices on sine at C2. Type the BEATs into their tooltips. The pairs beat at 0.1, 0.15 and 0.25 Hz, and the third and fifth harmonics of each pair three and five times as fast, in their own registers: a slow pulse low down, a quicker shimmer above it. With DRIFT fully left the pattern repeats exactly every 20 seconds.

Now turn **DRIFT** back to its centre and leave it. Over the next minutes the beats slow, stall and quicken as the voices wander apart and back, and for the first twenty minutes the whole bank rises slowly towards pitch as it warms up. Then try voice 3 on **JUST**, **COARSE** at 3/2: the fifth's harmonics beat against the root's where they meet, at its third and sixth harmonics.

Rings struck in turn



- | | | | |
|-------------------|-----------------------|-------------------------|------------------------------------|
| 1 RATE | 20 seconds a turn | 2 ZOOM | fully right, to see the whole turn |
| 3 TURN | on | 4 RANGE, lanes 1 to 3 | POS |
| 5 NORMAL | FREE | 6 Switches, voices 2, 3 | JUST |
| 7 COARSE, voice 2 | 5/4 | 8 COARSE, voice 3 | 3/2 |
| 9 Q, each voice | a little short of OSC | | |

A lane can play gates as well as curves. On POS, a lane's left edge is 0 V. Click a point into the lane, right-click it until it is a step, and drag it across: it holds its value until the next point. Put a second step point just below it, at the left edge, and the lane is a gate, high for the stretch between the two and at 0 V everywhere else. The drawing shows three lanes of four gates each.

Each lane strikes one of Fluvial's voices. **PING** strikes on each rise above 0 V, as hard as the rise is high, so a gate's place across its lane is how hard it strikes: the right edge, 5 V, is the hardest. How long it stays high does not matter. Each voice, struck, rings as a sine at its pitch and dies away as slowly as its Q lets it: here the root, the just third above it and the just fifth, each rung four times a turn, at uneven gaps and strengths, so that the three interleave and each ring sounds over the tails of the others (Figure 33).





Figure 33. One turn of this patch, run through Glacial's and Fluvial's own code. Above, the three lanes' gates; below, the size of each voice's ring. Each strike is as hard as its gate is high, and each ring is still sounding when the next voice is struck.

Move the gates and the pattern changes on the next turn; drag one across its lane to accent it. Set RATE slower and the same pattern spreads out, the rings dying away between strikes. Set the three BEATs a few hundredths of a hertz

apart and the rings beat as they fade. Glacial's **SEAM** gives one trigger a turn, which strikes all three together if it is patched to each PING.

A beat composed over an hour




- | | | | |
|-------------------------|---|-------------------------------|--------------------------|
| ① RATE | one hour a turn; ZOOM fully right to see it all | ② RANGE, lanes 2 and 3 | POS |
| ③ TURN | on | ④ AMT, voice 2's BEAT | a little right of centre |
| ⑤ LEVEL, voices 2 and 3 | zero | ⑥ AMT, voices 2 and 3's LEVEL | fully right |
| ⑦ COARSE, voice 3 | JUST, 3/2 | | |

Glacial plays the score drawn on its paper (above, the whole turn in view; the dashed line is the seam). Lane 1 moves voice 2's beat against the root, through zero and out the other side: as the lane crosses its centre the beat slows to nothing and comes back. Lanes 2 and 3 bring voices 2 and 3 in and out, each swell taking ten minutes or more.

Nothing in this patch is heard to move. Listen to it for a few minutes, leave it, and come back.

A drone moved off its harmonics




- | | | | |
|------------------|---------------------|------------------|--------------------------------|
| ① WAVE, voice 1 | saw | ② LEVEL, voice 2 | unchanged: not used |
| ③ LEVEL, voice 3 | unchanged: not used | ④ SIDE | fully right, SUM |
| ⑤ Range | WIDE | ⑥ SHIFT | +20 Hz, typed into its tooltip |
| ⑦ FEEDBACK | from off, slowly up | | |

A sawtooth at C2 has partials at 65.4, 130.8, 196.2 Hz and so on, all on one harmonic series. Shifted up by 20 Hz they sit at 85.4, 150.8, 216.2 Hz: the same distances apart, no longer in the same ratios, and the tone turns metallic. As **FEEDBACK** comes up, each partial grows a ladder of copies 20 Hz apart above it. Turn **SIDE** to DIFF and the ladders descend.

Then switch the range to **FINE** and set **SHIFT** to +0.2 Hz. Every partial now beats against its copies at 0.2 Hz, the whole spectrum rising and falling together.


A stereo pair, turning



- | | | | |
|--------------------|-------------|--------------------|-----------------|
| ① BEAT, voice 2 | +0.100 Hz | ② BEAT, voice 3 | -0.150 Hz |
| ③ SIDE, channel 1 | SUM | ④ SHIFT, channel 1 | +0.070 Hz, FINE |
| ⑤ SIDE, channel 2 | SUM | ⑥ SHIFT, channel 2 | -0.070 Hz, FINE |
| ⑦ BLEND, channel 1 | AM, halfway | ⑧ BLEND, channel 2 | AM, halfway |

The same drone goes to both channels: a cable in **A IN 2** breaks its normal from **OUT 1**. Each channel mixes the drone with a copy of itself shifted 0.07 Hz, up on the left and down on the right, so the two sides swell and fade out of step with each other, and the sound moves slowly between the speakers, beneath the voices' own beating.

Laid on tape




- | | | | |
|-------------------|---------------------------|-----------------|-------------|
| ① BEAT, voice 2 | +0.100 Hz | ② BEAT, voice 3 | −0.150 Hz |
| ③ LENGTH | about 20 s | ④ REC | on |
| ⑤ AMT, ERASE's CV | fully right | ⑥ DISINTEGRATE | a little up |
| ⑦ RATE | any; SPLICE sets the turn | ⑧ RANGE, lane 1 | POS |
| ⑨ TURN | on | | |

Alluvial records Littoral's drone onto a loop about 20 seconds long, and **SPLICE** resets Glacial's drum each time the loop comes round, so lane 1 falls in the same place on the tape on every pass. Lane 1 drives the erase head: high where the loop is wiped and rewritten, low where each pass is laid over the last. The stretches that are never erased build up layer on layer, darkening as the record head's bias and the wear take their highs, while the erased stretch stays fresh.

After a few minutes, press **REC** again to stop recording, and change the drone. The loop plays on, wearing slowly, its damaged spots spreading if **DISINTEGRATE** is up, until it sinks into the tape.

The whole series



All five together, in a patch built to test the series as a whole. Everything is tuned from Littoral's reference, A1 (55 Hz), passed to Fluvial by REF. One output may feed several inputs: Calderal's OUT 1 goes to two of Fluvial's.

From	To
Littoral MIX	Calderal A IN
Littoral REF	Fluvial REF IN
Fluvial OUT , voice 1	Calderal B IN
Calderal OUT 1	Fluvial IN , voices 2 and 3
Fluvial MIX	Alluvial IN L ; Alluvial's OUT L and OUT R to the mixer
Glacial lane 1	Littoral CV (BEAT) , voice 3
Glacial lane 2	Calderal CV (BLEND)
Glacial lane 3	Fluvial FREQ , voice 2
Glacial lane 4	Fluvial CV (BEAT) , voice 1
Glacial lane 5	Littoral CV (LEVEL) , voice 2
Glacial lane 6	Alluvial CV (ERASE)

Module	Settings
Littoral	All three switches on JUST. Voice 1 on sine at the root, BEAT +0.070 Hz, LEVEL 0.8. Voice 2 on triangle, COARSE 3/2, BEAT -0.110 Hz, LEVEL about half, its LEVEL AMT a little under halfway right. Voice 3 on sine, OCT +1, BEAT +0.190 Hz, its BEAT AMT a touch right of centre, so lane 1 swings it about 0.3 Hz either way.
Fluvial	NORMAL on FREE, all switches on JUST. Voice 1 a little past OSC, band-pass, OCT +1, COARSE 7/4, BEAT +0.800 Hz, its BEAT AMT a hair right of centre, so lane 4 sharpens it by up to a quarter of a hertz: it rings by itself, and is Calderal's B. Voice 2 low-pass, OCT +3, Q a little under halfway, FREQ AMT a little right. Voice 3 band-pass, OCT +2, COARSE 5/4, Q a third of the way.
Calderal	DRIVE a quarter, BLEND a little past AM, its AMT a little under halfway right, LEVEL near full. Nothing else moved.

Alluvial	LENGTH 13.7 s, REC on, INPUT a little past centre, ERASE's AMT about two-thirds right, FEEDBACK a little past halfway, MIX a little past halfway, HIGHS a little under halfway.
Glacial	RATE four minutes a turn, TURN on, every lane BI, each lane a slow S-curve of three or four points across the turn.

Calderal rings Littoral's three voices against Fluvial's ringing voice, B, at 192.5 Hz. Because every voice is a just interval over one root, the sidebands land on the voices' own partials, and each pair beats at the difference of what it carries. A harmonic carries its voice's BEAT times its number; a sideband carries B's BEAT plus or minus its partner's. With the drum stopped and DRIFT fully left, four pairs stand out:

Where	What meets	Beat
82.5 Hz	B less the octave, and the fifth	$0.80 - 0.19 + 0.11 = 0.72 \text{ Hz}$
110 Hz	B less the fifth, and the octave	$0.80 + 0.11 - 0.19 = 0.72 \text{ Hz}$
247.5 Hz	B plus the root, and the fifth's third harmonic	$0.80 + 0.07 + 3 \times 0.11 = 1.20 \text{ Hz}$
275 Hz	B plus the fifth, and the root's fifth harmonic	$0.80 - 0.11 - 5 \times 0.07 = 0.34 \text{ Hz}$

The deepest by far is at 247.5 Hz, where the triangle's strong third harmonic meets a sideband nearly as strong; the others swell and fade by a tenth or less. With the drum turning, lane 1 moves the octave's BEAT by up to 0.3 Hz either way and lane 4 sharpens B, so the two 0.72 Hz beats range from about half a hertz to a little over one, and the others quicken by as much as B sharpens. Drift moves them all a little further.

Fluvial's other two voices filter the ring: a resonant low-pass that Glacial sweeps across the sidebands, lighting each as it passes, and a band-pass on the cluster round 275 Hz, so that its slow pulse stands out. Alluvial lays the whole on a loop that never lines up with Glacial's turn, and wears it in.

Glacial's six lanes move the octave's beat, the blend, the low-pass, the ringing voice's beat, the fifth's level and the erase over four minutes, so no two turns of the drum sound the same.