



VERDALIS
Natural sound synthesis

ChirpParade

Synthesised birds

Version 0.9.0

CLAP instrument for Linux and Windows

Part of the Verdalis Plugin Suite

2026-10-04

Contents

Introduction

- What is in this manual

Installing

- Linux
- Windows
- Hosts
- Your own presets
- Building from source

A first sound

- Four things to try first

How it works

- The contour is the bird
- Where the contours come from
- The four contour controls
- The valve, and why a symmetric oscillator cannot be a goose
- The partials: reaching for the recording
- Three families, found rather than chosen
- The tract: a tube and a beak
- Rasp is irregularity, not more harmonics
- A phrase, and why its gap matters
- A trill is not a modulation
- The flock is individuals, not a rate
- Drumming is not a voice
- Distance

The plugin window

- Gestures
- The mixer

Parameter reference

- Syllable
- Timbre
- Phrase
- Flock
- Drum
- Place
- Filter
- Envelope
- Output

The preset library

Using it in a host

- Notes
- Automation and modulation
- Reproducible renders
- CPU cost
- Rendering and tails

What is still being fitted

Reference and limits

- Where the model comes from
- Notes and limits

Versioning
License

Introduction

ChirpParade is a CLAP instrument that generates birds. Every syllable is computed while the plugin plays — nothing is played back. No two calls are ever alike.

It does not work the way a bird synthesiser usually does, and it does not work the way the first version of it did either.

A bird syllable is its frequency contour. Not a pitch envelope on an oscillator — the actual shape the pitch traces, which for a real chirp changes direction between two and forty times in fifty milliseconds and slews at up to four hundred octaves a second. That scribble is the bird. Smooth it out and you have a whistle.

So ChirpParade does not compute those shapes. It **measures** them. 67 frequency contours were extracted from real recordings, fitted as formulas, and they drive the oscillator directly. *Contour* chooses which one. There is no audio in the plugin — a contour is forty numbers describing a curve, which is the same thing van Hunter Adams gets when he reads a cardinal's trace off a spectrogram and writes down

```
f(x) = -260 sin(-pi x / 5200) + 1740
```

except with forty terms instead of one, because one term cannot draw a scribble.

Above the contour sits a **valve**. Air passes through a bird's syrinx only while the two membranes are apart, and *Voice* is the fraction of each cycle they are shut. At the bottom the valve never closes and what comes out is a pure sine — which is what 59 % of the library's syllables are. Close it and the airflow becomes a one-sided pulse with the harmonic stack a goose has, evens as well as odds. That is not a filter; there is nothing between the two settings because there is nothing between them in the bird.

Above *that* sits the trachea — a closed tube resonating at `c/4L`, with the beak both raising the resonance and following the pitch with it the way a songbird's gape does — and above that, the part that makes birds sound like birds: syllables into phrases, phrases separated by silence, and a flock of individuals calling and answering.

Every layer came out of measurement, and the measurement is checkable: the analysis code that measured the recordings also measures the plugin's own output and compares the two. `tools/analysis/README.md` in the source tree records all of it, including the first version's failure and the two bugs found in the measurement code itself.

ChirpParade 0.9.0 replaced the syllable model. Version 0.1.0 modelled the syrinx from first principles and was fitted to averages; it measured correctly on twenty quantities and sounded nothing like a bird. Presets and saved state from 0.1.0 do not load meaningfully. Nothing was released at that version. *What is still being fitted* at the end of this manual lists what is known not to match yet.

What is in this manual

Installing and *A first sound* are enough to get birds out of the plugin. *How it works* explains the model, and is worth reading before spending time in the parameter list — the layers are what the controls are named after. *Parameter reference* and *The preset library* are generated from the plugin itself, so they always describe the version on the cover. *Using it in a host* covers notes, automation, reproducible renders and CPU cost.

Installing

ChirpParade is one self-contained `.clap` file plus its `presets` folder. There is no installer, no shared library to place and nothing to register.

Linux

CLAP hosts scan `~/ .clap`. Copy the plugin folder from the release archive into it so that you end up with:

```
~/ .clap/ChirpParade/ChirpParade.clap
~/ .clap/ChirpParade/presets/
```

Then rescan plugins in your host. X11 and Cairo are the only external dependencies of the plugin window, and any Linux machine that can run a DAW has both.

Windows

CLAP hosts scan the common CLAP folder. Copy the plugin folder into:

```
C:\Program Files\Common Files\CLAP\ChirpParade\
```

so that `ChirpParade.clap` and `presets\` sit inside it, and rescan. There are no DLLs to install beside it — the window and its Cairo are linked in.

Keep the plugin and its presets together. The plugin finds its factory presets by looking for a `presets` directory next to its own binary. Moving the `.clap` file on its own leaves it with no factory presets at all.

Hosts

Tested with Bitwig Studio and Reaper. In Bitwig, `~/ .clap` is scanned by default; after installing, restart it or rescan under *Settings* → *Locations* → *Plug-in Locations*. ChirpParade then appears as an instrument under the vendor *Ravetracer*, and the factory presets are indexed through CLAP’s preset discovery, so they show up in the host’s own browser as well as in the plugin window.

Your own presets

`SAVE` in the plugin window writes to a user preset folder, which is created on first use:

PLATFORM	LOCATION
Linux	<code>\$XDG_CONFIG_HOME/ChirpParade/presets</code> , or <code>~/ .config/ChirpParade/presets</code>
Windows	<code>%APPDATA%\ChirpParade\presets</code>

User presets appear in the browser immediately, listed after the factory ones. They are plain text, so they can be copied between machines, kept in version control, or edited in any editor.

COLLECTIONS

A preset library is a shelf, and a long one is worth tidying. Your own presets are filed in **collections**: folders one level under the user preset directory. The browser lists the library down its left side — *All*, the factory presets, each of your collections, and *Unfiled* for presets saved into none — each with the number of presets in it. Clicking one lists just that one.

IN THE BROWSER	WHAT IT DOES
<code>NEW</code>	Makes a collection. It is listed at once, empty until something goes into it

RENAME	Renames the selected collection; its presets go with it
DELETE	Deletes the selected collection and every preset in it, after asking
Drag a preset onto a collection	Moves it there — or onto <i>Unfiled</i> to take it out of one
Right-click one of your presets	Renames it, changes its description, or deletes it

SAVE asks where a preset goes: pick a collection from the list — or *New collection...* to make one on the spot — then give the preset a name and, if you like, a **description**. The description is what the browser shows while the pointer is over the preset and what the help line shows while it is loaded: a sentence for whoever loads it next. Saving over one of your own presets offers its collection and its description again.

The factory presets are locked. They are part of the plugin, not files of yours, so they cannot be moved, renamed or deleted; to change one, load it, adjust it and save it as your own.

Collections are one level deep on purpose: a preset library is a shelf rather than a filesystem, and a tree deep enough to get lost in is one somebody will get lost in.

PRESET PACKS

A whole collection can be written out as a single file — a **preset pack**, extension `.chirpparadepack` — so a library can be handed to somebody else, or moved between machines, as one file rather than a directory.

With a collection selected in the browser:

BUTTON	WHAT IT DOES
EXPORT	Writes the selected collection — its presets and nothing else — as a pack into the <code>packs</code> directory beside <code>presets</code>
EXPORT AS...	The same, to a location you choose
IMPORT...	Lists the packs already in <code>packs</code> , plus <i>Other file...</i> for one from anywhere
REVEAL	Opens the folder the pack was just written to in the system's file browser

Importing adds the pack's presets to the library as a new collection named after the pack. Nothing is ever overwritten: importing the same pack twice gives two collections rather than a mixture of both versions in one.

A pack is the preset format again with a separator line between the presets, so it can be read, diffed and edited by hand like everything else here. It carries each preset's text verbatim rather than a re-serialised copy, which is what makes a round trip through a pack lossless.

EXPORT AS..., *Other file...* and **REVEAL** need the desktop's own file chooser — `zenity` or `kdiallog` on Linux, `xdg-open` to reveal, and the operating system's own on Windows. Where none is installed the buttons that need one are not drawn, and the plugin's own `packs` directory, which is where **EXPORT** writes, is enough on its own.

Building from source

The suite builds with a C++17 compiler and CMake 3.16 or newer. From the plugin folder:

```
./install.sh
```

That configures, builds, runs the self-test and installs to `~/.clap/ChirpParade/`. `ChirpParade_PREFIX=/some/where` `./install.sh` overrides the destination, and `-DCLAP_INCLUDE_DIR=/path/to/clap/include` points the build at a CLAP checkout it cannot find by itself.

A first sound

1. Load ChirpParade on an instrument track.
2. Play a short note. One bird sings one phrase, where you put the note.
3. Now hold a long one. The same phrase fires, and then a flock carries on by itself for as long as the note lasts.
4. Open the preset browser by clicking the preset name and try **Single Chirp**, then **Dawn Chorus**, then **Goose Flight**.

Those two behaviours are the two halves of the plugin, and they are separate controls. *Shot Level* is the phrase a note fires — the deliberate one, on the beat you chose. *Flock Level* is the birds that sing unprompted while the note is held. Turn either right down and you have the other on its own: **Single Chirp** is the first with no second, and **Dawn Chorus** is the second with no first.

Four things to try first

Step *Contour* through, on Single Chirp. Each position is a different real syllable, measured off a recording. This is the control that decides what kind of bird it is, more than anything else in the plugin.

Then sweep *Detail* on the same note. At 100 % you hear the measured contour; at the bottom it smooths into a glide. Nothing else changes. That difference is the whole reason this version of the plugin exists.

Turn *Voice* from bottom to top, on a low pitch. At the bottom it is a pure sine. Somewhere in the middle a second harmonic appears, then a third, and by the top it is a rasping stack. Nothing was filtered — the valve is closing, and air only passes while it is open.

Change *Species*. Ten of them, and each carries its own set of measured contours as well as its own pitch register, length, richness, roughness and rate — because in the recordings those move together. *Pitch* at its default is the median of the whole library, so selecting Goose at the default lands at 743 Hz, which is a goose. Cuckoo is the odd one out: its call is two different notes, so a phrase plays its measured “cu” and then its “coo”, and *Motif* sets the fall between them – 4.2 semitones in the recordings. Load **Cuckoo Call** to hear it.

Play a chord, then a melody. Each note is its own bird with its own flock. For melodies, load **Melody Bird**: one syllable a note, almost no variation, a flat contour and *Sweep* pulled down, so what you hear is the key you pressed.

How it works

The contour is the bird

Put a recording of a chirp on a spectrogram and the syllable is a line. That line is what you recognise: not its average pitch, not how wide it sweeps, but its **shape**. Two birds with the same pitch, the same duration and the same sweep width sound like different birds if their lines differ.

Measured at a third of a millisecond, a real syllable is far busier than it looks on a normal spectrogram:

	MEASURED THROUGH A 21 MS WINDOW	AT 0.33 MS
peak pitch slew	2.7 oct/s	20 to 440 oct/s
direction changes	0.25	2 to 40
octaves travelled	0.35	0.2 to 6.8 — against an end-to-end range of 0.2 to 1.0

A syllable travels three to seven times further than its range, because it doubles back. That is the sound of a bird, and it is why the first version of ChirpParade — which drew each syllable as one smooth sinusoidal gesture — did not work, whatever its physics said.

Where the contours come from

67 of them, extracted from the reference library:

1. Every well-isolated syllable is found and its dominant partial tracked at 0.33 ms resolution.
2. The pitch track and the level track are each fitted as a cosine series in normalised syllable time. **Forty terms** lands within 30 cents of the real curve; eight is out by 82; one — which is what a single gesture is — is out by more than that before it starts.
3. The contours are clustered per species, and the **medoid** of each cluster is kept: one real measured syllable, never an average. Averaging two contours that zig-zag out of phase gives a smooth glide, which is exactly the thing being avoided.

5749 syllables passed the quality gate; 72 became archetypes, eight for each of the first nine species. The cuckoo was added later from four recordings of its own: 38 whole calls, and eight archetypes again — four of each note, kept as two groups so that a phrase can play them in calling order. What ships is **80 archetypes, 15360 numbers, 60 KB, and no audio at all**.

Each archetype keeps its own measured duration as well as its shape, and the engine plays it at that duration rather than stretching every curve to one length. A single nightingale recording spans 21 ms to 1296 ms, and that refusal to repeat itself is most of what makes it sound like a bird.

The four contour controls

Contour walks across the archetypes of the current species, lowest-sitting first. It is the most important control in the plugin.

Detail is how much of the contour's fine motion survives. At 100 % the measured curve passes through; turn it down and the scribble smooths towards a glide. It is worth sweeping once on **Single Chirp** just to hear what the difference is — the bottom of this knob is where the first version of the plugin lived.

Sweep scales how far the contour travels: 100 % is the measured curve, and it goes to 300 % because the range is wanted.

Skew bends the syllable's own time axis, crowding the contour towards the start or the end. 0.5 plays it at its measured pace — the asymmetry the library shows, a 24 ms rise against a 41 ms fall, is already in the curve.

Pitch transposes the whole thing. The contour is anchored at the syllable’s **loudest moment**, so what you hear is what the knob says; and because the default is the library median, every species at the default sings in its own register.

The valve, and why a symmetric oscillator cannot be a goose

The tone comes from a phase accumulator at the contour’s frequency, through a one-sided valve. *Voice* is the fraction of each cycle the valve is shut:

VOICE → CLOSURE	0.00	0.13	0.26	0.65	0.78	0.91
harmonics	1	2	3	4	5	5

At the bottom the valve never closes and a **pure sine** comes out, which is what 59 % of the library’s syllables are. Close it and the airflow becomes a one-sided pulse.

That one-sidedness is not a detail. It is where every **even** harmonic comes from. A symmetric oscillator — and the syrinx equation in the physics papers is symmetric — has energy at f, 3f, 5f and nothing between, and no amount of drive makes it a goose: a 234 Hz fundamental comes out as 234, 656 and 1125 Hz. Air passes only while the membranes are apart, and rectifying at that point is the same step that makes a human glottal pulse rich rather than sinusoidal.

Radiate is the last part of it. A small source radiates the *rate of change* of the flow rather than the flow, which tilts the harmonics up by 6 dB an octave, referenced to a fixed frequency so it stays a tilt and not a gain that changes with the note.

The partials: reaching for the recording

The valve gives a plausible harmonic stack. What it cannot give is a stack that **changes shape** as the syllable goes — and real ones do. Measured over 1777 syllables, the balance between the first partials moves **4.4 dB** across a single syllable. A fixed valve through a fixed tract is flat by construction.

So each archetype also carries six measured amplitude curves, one per harmonic, and *Partials* crossfades the valve into them. They are phase-locked to the same oscillator, because a harmonic source is.

Two things to know about it:

It is scaled by how much the measurement is worth. Each archetype records what share of its own energy fell inside the harmonic comb. A syllable that was inharmonic, or had a second bird in it, has a low share and barely responds — which is honest, rather than asserting a balance that was never measured.

It thins the low, rough voices, and that is the library’s fault rather than the engine’s. The contour quality gate selects for tonality, so the archetypes are each species’ *cleanest* syllables. Turn *Partials* up on a Goose and its timbre starts evolving the way a real one’s does — but its harmonic count falls, because that particular syllable really did have fewer harmonics than a goose’s average. On the whistlers and warblers, turn it up freely.

Three families, found rather than chosen

Sorting every syllable in the library by harmonic count alone splits it into three groups that barely overlap:

FAMILY	SHARE	FUNDAMENTAL	ROUGHNESS
whistle, ≤ 1 harmonic	59 %	median 3329 Hz	−30 dB
stack, 2–5 harmonics	25 %	median 1325 Hz	−26 dB
rich, ≥ 6 harmonics	17 %	median 399 Hz	−22 dB

Harmonic richness, roughness and pitch move together and downwards: the rich voices are the low ones. So the ten species are not ten presets of the same thing — they sit in genuinely different places, and choosing one moves pitch, length, richness, roughness, rate and the whole archetype set at once, because in the recordings those move together.

Breath is turbulent air past the membranes. It is calibrated: 13 dB of spectral flatness per decade, with 5 % putting a sparrow on the library's median roughness of -28 dB.

The tract: a tube and a beak

Above the syrinx is the trachea, closed at one end and open at the beak, so it resonates at $c/4L$ like any such tube. *Tract Length* is L , in centimetres.

The default is 4.9 cm, and it is measured rather than chosen: across the library's harmonic syllables, the loudest harmonic is **not the first in 80 % of cases**, and where it sits has a median of 1749 Hz. A tube resonating at 1749 Hz is 4.9 cm long, which is a plausible mid-sized songbird.

Beak is how far it is open, and it does two things. An open beak shortens the effective tube and damps it, so the resonance rises and broadens. It also makes the resonance **follow the pitch**: songbirds track the frequency they are producing with their gape, and a resonance that stays put while the fundamental sweeps past it hands the loudest harmonic from one to the next in the middle of a syllable, which sounds like a fault.

Formant is how much of that colouring is applied; at zero the syrinx is heard raw. *Breath* is turbulent air past the labia — and it is also what *starts* the oscillation, because the noise is injected into the oscillator rather than added to its output. That is why a syllable's onset is never twice the same.

Rasp is irregularity, not more harmonics

A corvid's rasp is a contact that is never the same twice. *Rasp* varies the valve's closure from one cycle to the next, which is broadband in a way that no harmonic stack is — the roughest recordings in the library measure 9 dB flatter in the spectrum than the cleanest, and stacking harmonics does not get there.

A phrase, and why its gap matters

Above the syllable, the structure is measured. 34 of the 58 references have a **bimodal** gap distribution: short gaps inside a phrase, long gaps between them, sitting a factor of six apart — 52 ms against 0.31 s.

That factor is why *Syllable Rate* and *Phrase Gap* are separate controls. One rate would give a stream; two give a phrase you can hear as a unit.

Legato is how much syllables run into each other. It defaults high, because **70 % of the library's syllable pairs have no silence between them at all**. A phrase of separated notes is the exception, not the rule.

Motif steps the pitch from one syllable to the next, which is what turns a repeated syllable into a figure. *Variation* is how much each syllable differs from the last, and it is the difference between a bird and a sequencer.

A trill is not a modulation

What the ear calls a trill in these recordings is syllables arriving too fast to separate — up to 22.8 a second — with no silence between them. So there is no trill oscillator: a trill is *Syllable Rate* high and *Legato* high.

A nightingale, though, is not a trill. Across 45 s of reference it sings 254 syllables from 21 to 1296 ms and 1644 to 5731 Hz, whose shapes differ from one another by a median pairwise distance of 5.22 — so **Nightingale Song** runs *Variation* wide enough to cross the whole archetype set instead of settling on one shape and repeating it.

The fine motion *inside* a syllable is real and abundant — two to forty direction changes — but it is not periodic, which is a different thing, and the measured contours carry it directly.

Amplitude pulsing within a syllable is also real: a quarter of the library's syllables have it, at a median 13 Hz and a depth of 0.82. That is now in the measured level contours, which is why *Pulse Depth* defaults to nothing — it is there to push further, not to supply what is missing.

The flock is individuals, not a rate

Birds sets how many there are, and each one keeps its own pitch offset, position, distance, syllable length, sweep, contour and voice for as long as the note lasts. A flock has individuals in it rather than one bird moving about.

They call as a **Poisson process** — the waiting time between phrases is exponential rather than a clock with jitter on it — at *Flock Rate*, in syllables a minute, which is how the library was measured (87 to 711, median 273). *Restless* drifts that density slowly, because no dawn chorus is uniform.

Answer is the one behaviour that is not a statistic. Real birds reply to each other, so with this up, one bird's phrase provokes another from somewhere else a fraction of a second later. It is most of what makes a flock sound like a conversation instead of a random process.

Drumming is not a voice

A woodpecker drumming is *sonation*: a bill against wood, with the syrinx not involved at all. So it is a separate layer, and it is modelled as a resonance excited by a contact rather than as an oscillator.

The measurements say two things that look contradictory and are both true. The strike spectra have a bandwidth of about 935 Hz at a centroid of 1251 — a resonance that broad has rung out in a fifth of a millisecond. Yet the same strikes take a measured 8 ms to fall 20 dB. Both hold: the body is broad and the *contact* is not instantaneous. So *Knock* is the wood's first mode and *Ring* is how long the contact lasts.

And one measurement contradicts the usual description outright. A woodpecker's roll is normally said to slow down towards its end. In this library **22 of 35 rolls speed up**: the interval drifts by a median -23 % across a roll, and the last third runs at 0.79 of the interval of the first. *Accelerate* therefore defaults to positive.

Drum Rate at zero means the woodpecker only answers a note, which is how to put a single roll exactly where it is wanted. **Woodpecker Drum** is set that way.

Distance

Birds are point sources, unlike wind or surf, so distance does two things: it takes the level down by the inverse of the distance, and it takes the top off through air absorption. *Air* is how much of that top survives — cold dry air keeps more.

Depth is how much the birds differ in distance from each other, and it is most of what makes a wood sound deep rather than flat. The library's median L/R correlation is 0.72, noticeably more correlated than wind or surf, which is what *Width* defaults to reflect: a bird sits somewhere rather than being everywhere.

The plugin window

ChirpParade draws its own window with X11 and Cairo on Linux, and Cairo on the win32 backend on Windows. There is no toolkit dependency, which is what keeps the plugin one self-contained file. It is embedded in the host's own window through CLAP's GUI extension and repainted from the host's timer, and it resizes by zooming one layout, from half size to four times.

The window is the suite's, in ChirpParade's own finch-gold theme. Across the header, behind the wordmark, is a **sonogram** — because that is what the plugin makes, and a sonogram is what birds are read off. One stroke per syllable, entering at the right and scrolling off the left, its height and its shape deterministic in the syllable number: the same syllable always draws the same stroke. Beside the preset name is the activity meter, showing how many syllables are sounding and how many have been sung.

Clicking the version label fires one free phrase, with no note behind it. It is the quickest way to hear what a setting did.

Gestures

GESTURE	EFFECT
Drag a knob up or down	Change the value
Shift-drag	Fine control, a fifth of the travel
Double-click a knob, or right-click anything	Back to the default
Scroll wheel over a control	Step the value
Click a knob's value	Type a value in, with units
Click a selector's name	Open its list and pick a value
Click a selector's ◀ or ▶	Step one choice
Click the preset name	Open the preset browser
◀ or ▶ beside the preset name	Previous or next preset
Click the version label	Fire one phrase, with no note
SAVE	Save the current settings as a user preset
MIXER	Open the layer mixer

The line at the bottom of the window shows what the control under the pointer does — the same one-line explanation that appears in the *What it does* column of the parameter reference.

The mixer

Every layer's level lives on the panel that layer belongs to, which is right for editing one layer and wrong for balancing them against each other. **MIXER** opens the same controls arranged the other way round: one strip per layer, its level as a fader, its pan and its width as slim sliders under it. They are the plugin's own parameters, not a second set — moving a fader moves the knob on the panel, and the host sees the same automation it always would.

M mutes a layer and **S** solos one, which is how a preset gets built from nothing: solo one layer, get it right, bring the next one back. Several layers can be soloed at once.

Mute and solo are **not** parameters and are not saved with a preset. They work by holding a layer's level at the bottom of its range and remembering what it was, so while either is active the preset bar shows a **SOLO ON** or **MUTE ON** chip — click it to release every hold at once. **SAVE** releases them first, so a muted layer can never be written into a preset as a silent one, and so does loading a preset. Touching a held layer's fader also releases the holds rather than fighting the hand on it. Closing the mixer does not: soloing a layer and then going to its knobs on the panels is what the mixer is for.

Parameter reference

There are 62 parameters, grouped the way the plugin window groups them. Every one is automatable and modulatable from the host, and every one can be typed in directly by clicking its value in the window. The name in `code` is the key used in preset files.

SYLLABLE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Pitch <code>pitch</code>	200 Hz ... 8.00 kHz	2.60 kHz	Where the syllable sits. The fundamental across the library runs 292 Hz to 5.9 kHz with a median of 2580, which is the default: corvids at the bottom, a robin's whistle at the top. The contour is measured about its own centre, so this transposes it – and because the default is the library median, every species at the default sings in its own register.
Sweep <code>sweep</code>	0.00 % ... 300.00 %	100.00 %	Scales how far the contour travels. 100 % is the measured curve exactly; below that it flattens towards a held note, above it exaggerates. A multiplier rather than a width, because the contour already carries the real excursion – the archetypes travel 0.16 to 1.16 octaves of span and up to 3.2 octaves of path. It goes past 100 % on purpose: the range is wanted, and a percentage that stops at 100 cannot reach it.
Contour <code>contour</code>	0.0 % ... 100.0 %	50.0 %	Which syllable. Each species carries up to eight contours measured off real recordings – the medoids of its clustered syllables – and this walks across them, lowest-sitting first (for Cuckoo, across each of its two notes). It is the most important control in the plugin: a syllable's identity is its frequency contour, and these are measured curves rather than a shape derived from summary statistics.
Detail <code>detail</code>	0.0 % ... 100.0 %	100.0 %	How much of the contour's fine motion survives. A real syllable changes direction 2 to 40 times and slews at up to 440 octaves a second, and that scribble is most of what makes it a bird rather than a whistle. At 100 % the measured curve passes through; lower it and the contour smooths towards the glide that a summary statistic would have given – which is exactly what the first version of this plugin could only do, and why it did not work.
Length <code>length</code>	15.0 ms ... 800 ms	96.0 ms	How long one syllable lasts. Measured 27 to 481 ms, median 96.
Skew <code>skew</code>	0.0 % ... 100.0 %	50.0 %	Bends the syllable's own time axis: below a half it crowds the contour towards the start, above it towards the end. 0.5 plays the measured curve at its measured pace, and the asymmetry the library shows – a 24 ms rise against a 41 ms fall – is already in it.
Jitter <code>jitter</code>	0.0 % ... 100.0 %	12.0 %	How much the pitch wanders off its own contour. A syrinx is not a synthesiser: without this every syllable of a phrase is identical, which no recording in the library is.
Pulse Rate <code>pulse_rate</code>	4.00 Hz ... 150 Hz	13.2 Hz	Rate of the amplitude pulsing inside a syllable – the flutter of a trill held on one breath. Measured 8.5 to 66 Hz, median 13.2.
Pulse Depth <code>pulse_depth</code>	0.0 % ... 100.0 %	0.0 %	How deep that pulsing is, added on top of whatever the measured level contour already does. Defaults to nothing: the contours carry the real amplitude modulation, so this is for pushing further rather than for supplying what is missing.

TIMBRE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Species <code>species</code>	Whistler, Sparrow, Warbler, Budgie, Woodpecker, Crane, Goose, Screech, Piper, Cuckoo	Sparrow	What kind of bird. Each entry biases pitch, sweep, length, harmonic richness, roughness and syllable rate together, because in the references they move together. Every one is the median of the recordings of that bird; only Screech is invented. Cuckoo is the one whose call is two different notes: a phrase plays its measured "cu" and "coo" in turn, and Motif sets the fall

between them – measured at 4.2 semitones.

Voice voice	0.0 % ... 100.0 %	30.0 %	How much of each cycle the syrinx is shut. Air passes only while the labia are apart, so this is not a filter: at the bottom the valve never closes and a pure sine comes out, which is what 59 % of the library's syllables are, and closing it makes the airflow a one-sided pulse with the harmonic stack a goose has – evens as well as odds, which a symmetric oscillator cannot produce at all. 16 % of the library has six harmonics or more, so both ends are real.
Breath breath	0.0 % ... 100.0 %	5.0 %	Turbulent air past the labia. It is also what starts the oscillation: the noise is injected into the oscillator rather than added to its output, which is why a syllable's onset is never twice the same. Calibrated against the library: rendering a whistle at a sweep of this and measuring its spectral flatness gives 13 dB per decade, and 5 % is where a sparrow lands on the library's median roughness of -28 dB. The chosen species scales it, from 0.42 for the cleanest to 2.8 for the roughest.
Tract Length tract	1.00 cm ... 24.0 cm	4.90 cm	Length of the trachea above the syrinx, which resonates at $c/4L$ like any closed tube. In the library's harmonic syllables the loudest harmonic is not the first in 81 % of cases, and it sits at a median 1741 Hz – a 4.9 cm trachea, which is the default.
Beak beak	0.0 % ... 100.0 %	50.0 %	How far the beak is open. An open beak shortens the effective tube and damps it, so the resonance rises and broadens.
Formant formant	0.0 % ... 100.0 %	55.0 %	How strongly the tract colours the source. At zero the syrinx is heard raw.
Rasp rasp	0.0 % ... 100.0 %	0.0 %	How irregular the closure is from one cycle to the next. A corvid's rasp is not extra harmonics, it is a contact that is never the same twice, and that is broadband in a way no harmonic stack is. The roughest recordings in the library measure 9 dB flatter in the spectrum than the cleanest.

Radiate radiate	0.0 % ... 100.0 %	35.0 %	How much of the output is the labial velocity rather than its displacement. The radiated pressure of a small source follows the rate of change of the airflow rather than the flow, and this is the balance between the two: the derivative term tilts the harmonics up by 6 dB an octave, referenced to the library's median pitch so it stays a tilt and not a gain that changes with the note. It is what makes a small source sound small.
Partials partials	0.0 % ... 100.0 %	60.0 %	How much of the timbre is the archetype's own <i>measured</i> balance between its first six partials rather than the synthetic valve. This is the one control that reaches for the recording itself: across 1777 measured syllables the balance between the partials moves 4.4 dB over the course of a single syllable, and a fixed valve through a fixed tract cannot do that at all. At zero the voice is Voice, Breath, Rasp and the tract, all of which stay alterable; turning it up hands the spectral evolution back to the bird. Scaled by how much of that syllable's energy the measurement actually accounts for, so an archetype with a second bird in it does not pretend to know.

PHRASE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Syllables syllables	1 ... 24	3	Syllables in one phrase. Measured 1 to 22, median 3.
Syllable Rate syllable_rate	0.50 /s ... 40.0 /s	7.50 /s	How fast they come inside the phrase. Measured 2.0 to 21.6 a second, median 7.5. Past about 15 the ear stops counting them and hears a trill.
Rate Drift rate_drift	-100.0 % ... 100.0 %	0.0 %	Whether the phrase speeds up or slows down as it goes. Positive accelerates.
Legato legato	0.0 % ... 100.0 %	70.0 %	How much syllables run into each other rather than being separated by silence. 70 % of the library's syllable pairs have no silence between them at all, which is why this defaults high: a phrase of separated notes is the exception.
Motif motif	-12.00 st ... 12.00 st	0.00 st	How far the pitch steps from one syllable to the next. Zero repeats one syllable; a few semitones makes the phrase a figure that goes somewhere.
Variation variation	0.0 % ... 100.0 %	35.0 %	How much each syllable of a phrase differs from the last in pitch, length and contour. This is the difference between a bird and a sequencer.
Phrase Gap phrase_gap	0.02 s ... 20.0 s	0.31 s	Silence after a phrase before the next repeat. Measured 0.11 to 1.63 s, median 0.31 – and it is 5.9 times the gap inside a phrase, which is what makes the phrase audible as a unit.
Repeats	1 ... 8	1	How many times a phrase is repeated before the bird stops.

repeats repeats	1 ... 8	1	How many times a phrase is repeated before the bird stops.
---------------------------	---------	---	--

FLOCK

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Shot Level shot_level	-inf dB ... 6.00 dB	0.00 dB	Level of the phrase a note fires. This is the one-shot half of the plugin: play a note and one bird sings once, where you asked for it. Turn it right down for a flock with nothing deliberate in it.
Flock Level flock_level	-inf dB ... 6.00 dB	-12.00 dB	Level of the birds that sing unprompted while a note is held. This is the drone half: at -60 dB there are none and every sound is one you played.
Flock Rate flock_rate	2.00 /min ... 1.20 k/min	273 /min	How often somebody in the flock says something, as a Poisson process rather than a clock. Measured 87 to 711 syllables a minute across the library, median 273.
Birds birds	1 ... 16	5	How many birds there are. Each keeps its own pitch, position, distance and voice for as long as the note lasts, so a flock has individuals in it rather than one bird moving about.
Pitch Spread pitch_spread	0.00 oct ... 3.00 oct	0.80 oct	How far apart in pitch the birds are. Zero is a flock of clones.
Voice Spread voice_spread	0.0 % ... 100.0 %	30.0 %	How far apart their voices are: length, contour, richness and roughness.
Answer answer	0.0 % ... 100.0 %	35.0 %	How often one bird's phrase provokes a reply from another, shortly after and from somewhere else. Real birds do this and it is most of what makes a flock sound like a conversation instead of a random process.
Restless restless	0.0 % ... 100.0 %	40.0 %	How much the flock's rate drifts. At zero the density is constant, which no dawn chorus is.
Max Voices max_voices	4 ... 64	24	Ceiling on syllables and drum strikes sounding at once, so the cost is bounded.

DRUM

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Drum Level drum_level	-inf dB ... 6.00 dB	-inf dB	Level of the woodpecker layer. Off by default: 9 of the 58 references hold drumming and the other 49 do not.
Drum Rate drum_rate	0.00 /min ... 120.00 /min	6.00 /min	Rolls a minute from the flock. At zero the woodpecker only answers a note, which is how to place a single roll exactly where it is wanted.
Strikes strikes	2 ... 32	10	Strikes in one roll. Measured 4 to 21, median 10.
Strike Rate strike_rate	2.00 /s ... 50.0 /s	15.0 /s	How fast they come. Measured 8.1 to 20.8 a second, median 15.
Accelerate drum_accel	-100.0 % ... 100.0 %	23.0 %	Whether the roll speeds up as it goes. A woodpecker's roll is usually described as slowing down; the library says the opposite, and clearly – 22 of 35 rolls accelerate, and the last third of a roll runs at 0.79 of the interval of the first. Positive is faster.
Knock knock	150 Hz ... 8.00 kHz	1.00 kHz	The wood's first mode. There is a second above it, so the spectral centroid of a strike lands about a quarter higher than this: the library measures centroids of 654 Hz to 5.9 kHz with a median of 1251 over a bandwidth of 935, and 1 kHz here is what produces that.
Ring ring	0.50 ms ... 300 ms	13.0 ms	How long the contact lasts. The body itself is broad – 935 Hz of bandwidth is a Q near 1.3, which has rung out in a fifth of a millisecond – so what makes a strike last is the contact, not the resonance. Measured strikes fall 20 dB in 2 to 51 ms with a median of 8, which is 13 ms here.

PLACE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Distance	0.0 % ... 100.0 %	35.0 %	How far away the birds are. Distance takes the top off a call and drops its level, and it is the

<code>distance</code>				single most effective control here for putting birds behind a mix rather than in it.
Depth <code>distance_spread</code>	0.0 % ... 100.0 %	40.0 %		How much the birds differ in distance. A wood has near birds and far ones, and that spread is most of its depth.
Air <code>air</code>	0.0 % ... 100.0 %	50.0 %		How much of the high end survives the distance. Cold dry air keeps more.
Width <code>width</code>	0.0 % ... 100.0 %	60.0 %		Stereo spread of the flock. The library's median L/R correlation is 0.71, which is noticeably more correlated than wind or surf: birds are point sources, so they sit somewhere rather than being everywhere.
Space Amount <code>space_amount</code>	0.0 % ... 100.0 %	25.0 %		How much of the surroundings you hear: a wood, a courtyard, a hall.
Space Size <code>space_size</code>	0.0 % ... 100.0 %	55.0 %		Size of that space, in metres of it.
Space Damping <code>space_damping</code>	0.0 % ... 100.0 %	60.0 %		How absorbent its surfaces are. A wood full of leaves is very absorbent; a stone courtyard is not.

FILTER

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Filter Type <code>filter_type</code>	Lowpass, Bandpass, Highpass, Notch	Lowpass	Shape of the output filter.
Highpass <code>highpass</code>	20.0 Hz ... 2.00 kHz	20.0 Hz	Permanent highpass. Birds have almost nothing below 250 Hz, so this can go high without losing any of them.
Filter Cutoff <code>filter_cutoff</code>	30.0 Hz ... 20.00 kHz	20.00 kHz	Corner of the output filter.
Filter Resonance <code>filter_reso</code>	0.0 % ... 100.0 %	10.0 %	Resonance of the output filter.
Filter Key Track <code>filter_keytrack</code>	0.0 % ... 100.0 %	0.0 %	How far the filter follows the played note.

ENVELOPE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Attack <code>attack</code>	0.50 ms ... 10.00 s	2.00 ms	How long the flock takes to come up when a note starts. Short, unlike the rest of the suite: a bird does not fade in, and the syllable envelopes are inside the syllables.
Decay <code>decay</code>	5.00 ms ... 20.00 s	400 ms	Fall from the initial level to the sustain.
Sustain <code>sustain</code>	0.0 % ... 100.0 %	100.0 %	Level held while the note is down.
Release <code>release</code>	5.00 ms ... 30.00 s	600 ms	How long the flock takes to stop after the note is let go. Syllables already in flight always finish, so a short release thins the flock rather than cutting anything off.
Velocity To Level <code>vel_to_level</code>	0.0 % ... 100.0 %	60.0 %	How much velocity sets the level.
Velocity To Pitch <code>vel_to_pitch</code>	0.0 % ... 100.0 %	20.0 %	How much velocity raises the pitch, which is not the same thing: a bird calling harder pushes more air past a tighter syrinx, so it goes up as well as getting louder.

OUTPUT

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Output Gain <code>gain</code>	-inf dB ... 12.00 dB	0.00 dB	Output level.
Random Seed	0 ... 999	0	Non-zero gives the same birds every time. Zero is always different.

The preset library

20 factory presets ship with the plugin. They are installed beside the binary and exposed through CLAP preset discovery, so they appear in the host's own browser as well as in the plugin window's.

Alarm Flock

Short, high and fast, which is what an alarm call is for: hard to locate and impossible to ignore. 45 ms syllables at sixteen a second, seven birds, and Restless well up so the density lurches rather than holding steady.

Blackbird Song

A phrase that goes somewhere. Five syllables with three semitones between them, drawn from the warbler contours and exaggerated to 140 % of their measured depth, repeating after a gap four times the one inside it – which is the ratio the library measures between the two.

Budgie Chatter

Four birds talking over each other. The budgie recordings are the busiest in the library – 293 syllables a minute, and the shortest rise of any group at 17 ms – and they carry three harmonics, so the valve closes a little and the sine grows a stack.

Crane Bugle

A bugle carrying across water. Cranes are the second slowest group measured at 140 syllables a minute, with four harmonics on a one kilohertz fundamental. Distance is well up and the air is dry, so what arrives is the middle of the call and none of its top.

Cuckoo Call

The common cuckoo, from four references of the bird calling in a wood: 38 whole calls, each a short “cu” at 665 Hz and a longer “coo” 4.2 semitones below it – a major third – one call every 1.2 s. One harmonic and -37 dB of spectral flatness, the purest voice in the library. Cuckoo plays its two measured notes in turn, Motif is the measured fall, and a note starts a run of eight calls. The references' echo of each note, 15 to 30 dB down for 170 ms, is the Space here.

Cuckoo Valley

Two cuckoos a long way off, answering each other across open ground. The same measured call as Cuckoo Call, further away and wetter: distance takes the level and the top, and the space carries the call down the valley. No note-fired call; hold a note and they keep it up on their own.

Dawn Chorus

Twelve birds over two octaves, answering each other at half again the library's densest measured rate, each drawing from the eight measured sparrow contours with Variation letting them wander between. Nothing is played: Shot Level is off and the whole thing runs from the flock, so a held note is a morning rather than a bird.

Far Treeline

Birds a field away. Distance takes the level down by its inverse and the top off by air absorption, and the highpass and the output filter finish what the air started. Detail is pulled back a little: fine contour motion is the first thing distance takes, so smoothing it is what makes far birds sound far rather than quiet.

Forest Morning

A wood with a woodpecker in it. The drum sits 14 dB under the birds and arrives four times a minute, which is roughly what a walk through a wood in spring gives you, and the space is damped the way somewhere full of leaves is.

Garden Sparrows

The library's middle, made audible. Everything here is a median of the 4268 syllables measured across the reference recordings – 96 ms, 2.6 kHz, three to a phrase at seven and a half a second, five birds spread over 0.8 of an octave – and the contour is a real sparrow syllable played at its measured depth. The starting point rather than a destination.

Goose Flight

Geese going over. Only one goose recording survived the contour gate, so this species has a single archetype and Contour does nothing on it – the variety has to come from Sweep, Skew and the flock instead. Spread over six birds at different distances, which is most of what makes a skein sound like a skein.

Green Woodpecker

The yaffle: a falling laugh at ten notes a second. Both woodpecker recordings arch rather than fall, and the archetypes carry that – this sits at the top of their set, where the widest of them are. The phrase slows as it goes.

Jungle Screech

The one species with no reference of its own. Screech borrows the eight most extreme contours in the whole library – the ones whose pitch travels furthest, 7.5 to 10.3 octaves of path – which is what ‘every range at once’ honestly means when there is nothing to measure. Kept because a plugin for birds should be able to make a noise no bird makes.

Melody Bird

For playing tunes with. One syllable a note and hardly any variation, on the flattest of the whistler contours with Sweep pulled down, so the pitch that comes out is the pitch of the key – the contour is measured about its own centre and Pitch transposes it, which is what makes that true. Velocity moves the pitch as well as the level, the way a bird calling harder goes up.

Nightingale Song

A nightingale does not repeat itself. Across 45 s of the reference it sings 254 syllables from 21 to 1296 ms, over 1644 to 5731 Hz, whose shapes differ from each other by a median pairwise distance of 5.22 – so this preset runs Variation wide enough to cross the whole archetype set rather than settling on one. Phrases are short, one to sixteen syllables, with the syllables inside a phrase running together and real silence between them.

Parakeet Flock

The same voice driven hard and made irregular. Rasp is the closure varying from one cycle to the next – a contact that is never the same twice – which is broadband in a way no harmonic stack is. Ten of them at once is a street full of parakeets rather than a cage.

Robin Whistle

The purest voice in the library. The robin recordings measure one harmonic and a spectral flatness of -34 dB, the cleanest of any group, so Voice sits at the bottom where the valve never closes and what comes out is a sine. Four notes with a slight fall between them.

Single Chirp

One syllable, one bird, no flock: the plugin reduced to what a note does. The contour is one measured sparrow chirp and nothing is smoothing it, so this is the preset to sweep Contour on and hear what the archetypes actually are. Dry on purpose – the room is the next thing to add, not part of the syllable.

Tui Bell

What the tract does on its own. A short trachea and a wide beak put the resonance high and broad, and Formant well up makes it most of what is heard – the source is the same valve as everything else. In the library’s harmonic syllables the loudest harmonic is not the first in 80 % of cases, and this is that effect used deliberately.

Woodpecker Drum

Drumming only, and only on a note: Drum Rate is zero, so every note is one roll and nothing arrives unasked. Fourteen strikes accelerating by a quarter across the roll, which is what the library measures – 22 of its 35 rolls speed up, against the usual claim that they slow down.

Using it in a host

Notes

A note does two things: it fires one phrase at *Shot Level*, and it opens the gate on the flock at *Flock Level*. Which of those you want is a decision, and the presets are split roughly half and half between them.

The note is a **pitch**. Middle C leaves *Pitch* alone and every other key transposes it, so ChirpParade can be played as an instrument rather than triggered as a sound effect — **Melody Bird** is set up for exactly that. *Velocity To Pitch* is worth knowing about: velocity moves the pitch as well as the level, because a bird calling harder pushes more air past a tighter syrinx. Both are referenced to full velocity, so *Pitch* means the pitch at velocity 1.

The ADSR gates the **flock** rather than shaping individual syllables — the syllable envelopes are inside the syllables. A short *Release* thins the flock out rather than cutting anything off mid-chirp, because syllables already in flight always finish.

The phrase a note fires is not gated at all. It completes however short the note was, which is what makes a note usable as a one shot: a sixteenth-note trigger on *Green Woodpecker* still gets all fourteen syllables of the laugh. The flock around it stops with the envelope.

Attack is short by the standards of the rest of the suite, and deliberately: a bird does not fade in.

There are 16 voices, with 64 syllables, 64 strikes and 64 phrases in shared pools, so the number of held notes cannot multiply the CPU cost without bound.

Automation and modulation

Every parameter is automatable, and host parameter modulation (`CLAP_EVENT_PARAM_MOD`) is supported globally. Per-note modulation and note expressions are ignored. Parameter changes are sample accurate.

The ones worth automating are *Flock Rate* and *Restless* for a chorus that builds, and *Pitch* for a bird that moves. Changes to the syllable controls take effect on the **next** syllable spawned, not the one already sounding — a syllable is committed when it starts, which is what stops an automation ramp smearing across a 96 ms chirp.

Reproducible renders

Random Seed at 0 never repeats. Any other value renders the same birds for the same notes every time — verified byte-identical across two runs at 44.1, 48 and 96 kHz for all of the factory presets.

The flock's layout is deferred to the first scheduling tick rather than being drawn at note-on, because a note can arrive before the parameter events in its own block have been applied. That is what makes a fixed seed promise the same flock even when the host sends the note first.

CPU cost

Cost is bounded by the pools. Scheduling runs once every 32 samples, which is 0.67 ms at 48 kHz — fine enough for syllables at forty a second — and everything audible runs per sample.

The expensive control is *Voice*. A hard-driven low voice takes up to eight integration substeps a sample, because the nonlinear term has to stay inside the step; a whistle takes one. *Max Voices* is the ceiling on syllables and strikes sounding at once. *Space Amount* is the room model, and *Birds* costs nothing by itself — a bird is an identity, not a synthesiser.

Rendering and tails

A phrase that has started always finishes, so the tail is the release plus the phrase that was under way when the note went plus whatever the space is still doing with it. When bouncing, leave a couple of seconds past the end of the note, and more if *Repeats* or *Phrase Gap* are up.

What is still being fitted

ChirpParade 0.9.0 is a first fit of a new model, and the measurements say where it is not there yet. All of this is in the plugin's `TODO.md`.

The valve saturates around five harmonics and the noisiest references measure twelve. Closing it further aliases. The likely answer is a second syringeal side — a bird has two, controlled independently, which is how some species produce two notes at once.

Roughness cannot be met for the roughest voices. A harmonic stack is spectrally flat whether or not any noise is present, and the flatness measure cannot tell the two apart — so this may be a limit of the measurement rather than of the engine.

There is no crow and no raven. Both were modelled and both were removed. A corvid's call is not a shaped tone with noise on it; it is rough in a way this engine reproduces as a rough tone rather than as a croak, and no amount of fitting closed that gap. The reference recordings are still in the library and still part of its census — they just do not become an instrument.

Nothing morphs between contours. *Contour* steps from one measured curve to the next. Interpolating them naively would give smooth glides, since two curves that zig-zag out of phase average to a straight line — the exact failure this version was built to avoid — so a continuous control needs the curves aligned first.

Screech does not fit and cannot. It is the one species with no reference, borrowing the library's most extreme contours, and its measured pitch runs 54 % low because those curves swing so far that a fundamental estimate is meaningless. It is an effect, not a bird.

The quality gate still throws away most of the library. Relaxing it in this version widened every species' set — Goose went from one contour to three — but most of the dense multi-bird recordings are still rejected, because there the tracker follows two birds at once.

Warbler and Goose come out about half the length their table says. Both have archetypes of very mixed duration, and stretching a short curve to a long *Length* turns its own amplitude modulation into separate notes.

Reference and limits

Where the model comes from

- **Adams, V. H.**, *Birdsong synthesis* and the Cornell ECE 4760 birdsong lab — the method this plugin uses: read the frequency contour off a spectrogram, fit a formula, drive an oscillator with it. $f(x) = -260 \sin(-\pi x / 5200) + 1740$ is his cardinal swoop, and *Contour* is that idea automated over 58 recordings with forty terms instead of one.
- **Gardner, T., Cecchi, G., Magnasco, M., Laje, R. and Mindlin, G. B.**, *Simple motor gestures for birdsongs*, Phys. Rev. Lett. **87**, 208101 (2001) — that a syllable is two coupled gestures with a phase between them. The reference library confirms the ordering that predicts, which is why forty cosine terms describe a syllable so well; it is *not* enough to make two sinusoids draw one.
- **Laje, R., Gardner, T. and Mindlin, G. B.**, Phys. Rev. E **65**, 051921 (2002) — the reduced two-equation syrinx model.
- **Zysman, D., Méndez, J. M., Pando, B., Aliaga, J., Goller, F. and Mindlin, G. B.**, *Synthesizing bird song*, Phys. Rev. E **72**, 051926 (2005) — that air sac pressure tracks the sound envelope and syringeal tension the pitch, so both can be recovered from a recording. Also the observation, quoted in the source, that a richer model is needed for species with a wide timbre — which is the one-sided airflow above.
- **Mindlin, G. B. and Laje, R.**, *The Physics of Birdsong*, Springer (2005).
- **Sabine, W. C.** — the decay time of the space model, as elsewhere in the suite.

The papers themselves are not part of the release; they are not ours to redistribute. `tools/analysis/README.md` in the source tree records every measurement, the eleven bugs the measurements found in the engine, and the two they eventually found in themselves.

Notes and limits

- Linux and Windows, x86-64. Plugin state is stored little-endian.
- 16 voices, with 64 syllables, 64 drum strikes, 64 phrases and 16 birds per voice in shared pools.
- 71 measured contours, 9656 coefficients, 38 KB. No recorded audio of any kind.
- 60 s of *Dawn Chorus* — twelve birds, 48 voices — renders in 1.07 s, about 56 times realtime.
- Scheduling runs at a control rate of one update per 32 samples. Everything audible runs per sample.
- Host parameter modulation is supported globally. Per-note modulation and note expressions are ignored.
- Deliberately built without `-ffast-math`: on x86 GCC that links `crtfastmath.o`, which would flip the FTZ and DAZ flags for the entire host process.
- Birds only. No insects, no frogs, no wing beats. Several of the references also contain wind, traffic and water; those parts were subtracted as a noise floor and excluded from the measurements.
- The syllable controls are committed when a syllable starts, so automating them changes the next syllable rather than the current one.
- The contour tables are rendered once when the plugin is activated, not per note, and shared between instances.

Versioning

Semantic versioning. **MAJOR** is an overhaul: a rewrite of the synthesis model, or a change that breaks existing presets or saved host state. **MINOR** adds something — a parameter, a layer, presets. **PATCH** is fixes that add nothing new.

While the version is below 1.0 the model itself is still settling, and 0.2.0 replaced 0.1.0's synthesis outright. Presets and saved state are not carried across a 0.x bump.

License

MIT. The only external dependencies are the CLAP headers, which are MIT licensed, plus X11 and Cairo for the plugin window.

Everything ChirpParade produces is computed at run time. It contains no recorded audio of any kind — the measured contours are coefficients describing two curves, in the same sense that $f(x) = -260 \sin(-\pi x / 5200) + 1740$ is.

Reference recordings used during development are not part of the release and are not redistributable.