



DIGITAL THEREMIN

USER MANUAL

Table of Contents

Introduction.....	1
Features.....	2
Controls.....	3
Inputs / Outputs.....	5
Main Modules.....	7
Pitch & Volume Fields.....	8
Pitch Corrector.....	11
Envelope Generator.....	13
MIDI Module.....	15
LED Visual Tuner.....	17
Pitch Preview.....	19
DSP Synthesizer.....	22
Triple Oscillator.....	23
Noise Generator.....	28
Formant Bank.....	30
Inharmonic Resonator.....	32
Miscellaneous Controls.....	35
Preset System.....	37
Pitch & Volume Field Setup.....	39
Troubleshooting.....	43
FAQ.....	45
Acknowledgments.....	47


Introduction

The D-Lev is a (mostly) digital Theremin with many unique features. The primary achievement is the linearity of the pitch and volume fields, which are much more linear than anything you've likely experienced. Field sensitivity can be configured to best suit your playing style, and the volume field has an adjustable knee. True pitch correction, full-featured pitch preview, and real-time visual pitch feedback via the LED "tuner" all provide much welcomed assists to performance.

The synth is sufficient to emulate human vocals, strings, brass, woodwinds, bells, percussion, crazy sounds, etc. The variety of voices encourages the development of different playing techniques, and helps maintain an interest in playing itself. MIDI out lets you control almost any outboard synth for further sonic explorations.

The D-Lev was designed from the ground up using FPGA technology, and incorporates my Hive soft processor core and phase locked peripherals into a highly responsive, tightly integrated whole. Boot time and null calibration are almost instantaneous, so you won't lose any ideas or inspiration waiting to play.

Exactly like a traditional analog Theremin, the D-Lev employs oscillators that vary in frequency with the capacitance of the player's hands. The exact same electrostatic interaction is at work, but in the D-Lev the oscillator frequencies are now represented by digital numbers, and the heterodyning process used to develop a difference frequency in the analog theremin is now replaced by a simple arithmetic subtraction. Musicians who are accustomed to traditional analog Theremins won't have to alter their fundamental style or techniques in order to play the D-Lev.



How a digital Theremin works.

The variable oscillators in the D-Lev are almost entirely digital, which makes them quite stable. The operating frequency and volume are therefore just a series of numbers which are fed to an internal DSP (digital signal processing) based music synthesizer in order to produce audio.

Separating the process of sensing hand capacitance (related to hand position) from the generation of sound opens up a new world of exciting possibilities for both the designer and the player. The pitch and volume numbers are extensively digitally filtered to dramatically reduce environmental noise and hum. The responses of the fields, which are key to making the D-Lev comfortable and predictable to play are linearized mathematically and stored in any of four system presets. The octave and timbre of preset voices are entirely defined by the synth, with presets to easily switch among a variety of very different voices. In terms of technology and features, the D-Lev stands out among Theremins, both past and present.

Features

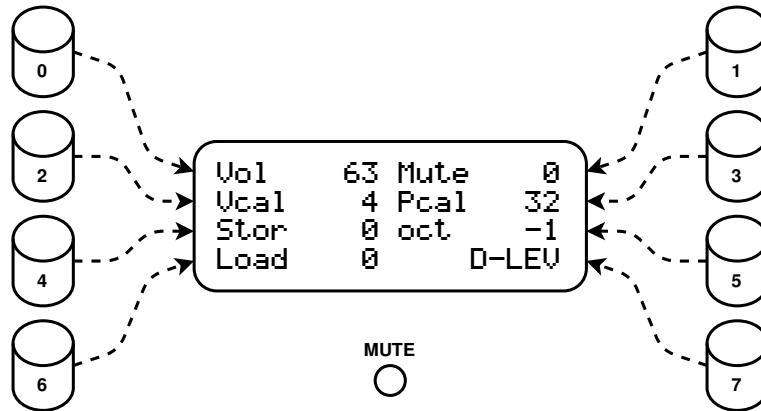
- Nearly perfect pitch & volume field linearity, tailored by the user — right up to the antennas!
- Adjustable field sensitivities and offsets to accommodate any playing style.
- Volume field has a fully configurable / defeatable knee (to reduce "hand flapping").
- Wide gestural bandwidth precisely tracks your every hand movement.
- High antenna voltage swings provide large fields immune to most common forms of interference.
- 22nd order pitch & volume field digital filtering kills mains hum dead (50/60Hz).
- Instant or delayed auto-calibration of both pitch & volume fields via foot pedal or button.
- Zero lag LED visual tuner for real-time pitch & volume feedback during play.
- Full-featured chromatic pitch correction — from none, to unobtrusive, to full hard quantization.
- Highly configurable itch preview on its own separate stereo channel.
- Subtractive / modal synthesis engine with pitch and volume modulation.
- Realistic sounding human vocals and strings — in stereo!
- Envelope generator with velocity sense for plucked and percussive effects.
- MIDI out with CC, pitch bend range, trigger location, velocity, and octave controls.
- 199 voice presets (99 read-only, 100 read/write) and 4 system presets.
- *ALL* parameters fully adjustable via the front panel interface.
- Easy to own, zero internal adjustments.

Technology

- Low-drift, high Q / high voltage, single layer, hand wound air core solenoids.
- Rock steady, high resolution, dithered phase locked digital oscillators.
- FPGA-based logic with custom 8 thread Hive soft processor core.
- Internal DSP 32 bits @ 48kHz; DACs 24 bits @ 48kHz.
- LED visual tuner employs PWM dimming between notes to intuitively display real-time pitch and volume.
- User interface: 8 rotary encoders; 20 x 4 line backlit LCD.
- Outputs: Stereo monitor / headphone out (1/4" TRS lower 1/8" TRS front panel); Stereo line out (1/4" TRS); MIDI TX (DIN).
- Inputs: RS-232 over USB; ACAL-auto calibration (1/8" TS from foot pedal).
- Preset librarian / remote control software (Windows & Linux).

Controls

The D-Lev is controlled by two columns of four rotary encoders each. These directly correspond to the two columns of text on the four line LCD display:



Each encoder controls one associated value on the LCD display.

For example, on the main **D-LEV** screen above: turning encoder 3 adjusts the **Pcal** parameter, and pressing the 6 encoder zeros out **Load**.

How the encoders and LCD work:

- Turning an encoder quickly changes the displayed value in larger steps, so you can make big edits with less work.
- Pressing an encoder zeros out or defaults the displayed value (with some exceptions). Use this to quickly lower a value or defeat a function.
- Turning encoder 7 selects one of 20 screens, or *pages*. Turning / pressing encoders 0 through 6 will adjust the associated display values on the current page:

Vol	63	Mute	0	Vol	63	Line	63	vmod	0	fmod	0	vloc	0	levl	31
Ucal	4	Pcal	32	Prev	0	Treb	0	Prev	0	treb	0	slew	20	span	31
Stor	0	oct	-1	osc	46	Bass	0	harm	0	bass	0	cntr	12	prvw	0
Load	0		D-LEV	nois	20		LEVELS	oct	0		PREVIEW	Post	0		PITCH
kloc	50	rise	0	vloc	0	velo	0	vmod	7	fmod	-10	vmod	0	fmod	0
knee	6	fall	36	cloc	0	cc	2	nois	20	treb	0	frea	1357	nois	20
velo	0	damp	36	bend	64	chan	1	puls	5	bass	0	mode	-2	xmix	31
dloc	0		VOLUME	oct	0		MIDI	duty	0		NOISE	reso	0		FLT_NOISE
vmod	20	fmod	0	offs	0	offs	0	vmod	0	fmod	0	frea	46	tap	42
harm	18	treb	0	hmul	0	hmul	0	frea	27	osc	46	hpf	659	harm	35
odd	0	bass	0	fm	0	srpd	0	mode	0	xmix	0	mode	2	xmix	11
oct	-1		0_OSC	xmix	0		1_OSC	reso	0		FLT_OSC	reso	0		RESON
vmod	0	fmod	0	vmod	0	fmod	0	vmod	0	fmod	0	vmod	0	fmod	0
frea	905	levl	58	frea	1357	levl	58	frea	3623	levl	57	frea	320	levl	48
frea	27	levl	0	frea	27	levl	0	frea	2960	levl	46	frea	27	levl	0
reso	10		0_FORM	reso	13		1_FORM	reso	11		2_FORM	reso	7		3_FORM
Ucal	4	Urev	1	Pcal	32	Cent	0	Brit	1	Cent	0	Wait	10	50Hz	0
Lin	-2	Dith	0	Lin	-2	Dith	0	LCD	31	Note	0			P<>U	0
Ofs-	30	Ofs+	63	Ofs-	32	Ofs+	100	Qant	1	Oct	0	Stor	0	Erev	0
Sens	75		U_FIELD	Sens	65		P_FIELD	Post	0		DISPLAY	Load	0		SYSTEM

The 20 D-Lev screens / pages.

- Pressing encoder 7 on any page initiates an *ACAL* (auto-calibration of the fields) which also blinks the Mute LED.
- Pressing encoder 4 (**Stor**) on the **D-LEV** or **SYSTEM** pages initiates a two step preset write. The the screen label will change to **?WR?** and the Mute LED will start blinking. Pressing it again actually performs the write and returns things back to normal. If you want to abort the write, turn any encoder, or press any encoder other than 4.

Most screen values have binary limits (such as 0 to 31) as opposed to decimal (such as 0 to 9) which may seem unusual at first. These give finer grained ranges and are a better fit to the underlying hardware and software.

In this manual these control ranges are represented by bracketed values. For example, the **Vcal** range of 0 to 255 is shown as **Vcal[0:255]**, and setting **Load** to 45 is shown as **Load[45]**.

Inputs / Outputs

Inputs

5VDC power is supplied by a “wall wart” connected to the base of the D-Lev via a 12mm 4-pin “aircraft” connector that includes a thread-on ring to prevent the power connector from falling out. The required current is 1A, and an inexpensive switching-type AC adapter should work fine. This power supply is a widely available generic type that has been modified with the improved output connector..

Earth ground is essential to proper Theremin operation. Without it, the pitch and volume fields may be small or unstable.

Unlike most other theremins the D-Lev is not automatically grounded when connected to mains power. The player has the option to connect to mains ground through a separate cable, or in cases where the mains ground has a great deal of noise it may be necessary to use another method.

The D-Lev is grounded by connecting the banana jack (located on the base of the theremin directly between the power switch and the threaded stand coupling) to either mains ground or a local ground.

See Appendix A for more information on grounding alternatives and ideas.

RS-232 over USB is used to connect the librarian software to the D-Lev for preset management.




1/4" TS phone jack for ACAL external foot switch.

ACAL is a momentary foot switch jack (1/4" TS).

Outputs

Monitor / Headphone is an unbalanced 1/4" stereo TRS phone jack on the bottom of the theremin and the left 1/8" TRS jack on the front panel under the tuner. Use the front panel jack for monitoring pitch preview through an earpiece or use the bottom jack for connection to a small powered monitor placed behind or beside the player.

The MONITOR level of the pitch preview signal in both jacks is controlled by the left knob directly under the tuner on the front panel.



1/4" stereo TRS phone plug for Monitor / Headphone and Line outputs.

Line is an unbalanced 1/4" stereo TRS phone jack. Run this into the mixing board for recording purposes, or into the amplifier / PA that your audience will hear.

The LINE level is set with the right knob under the tuner display, and this should be set to the fully clockwise position to obtain a standard +4dBu audio output to the mixer.



MIDI 5-pin DIN.

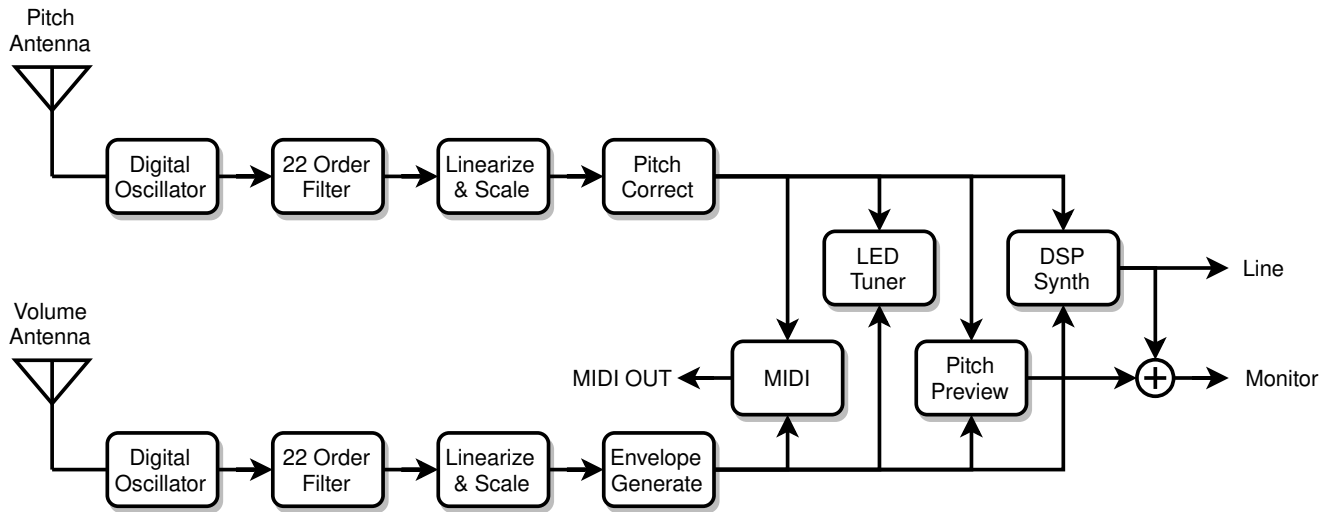
MIDI TX is a 5-pin female DIN jack. Connect this to an external MIDI music synthesizer use the D-Lev strictly as a controller with the actual sounds being generated by the external synth and not the D-Lev's internal synthesizer section.

It should be noted that the D-Lev's sound section is not disabled by the use of the MIDI output, and in some cases it may be possible to find interesting sound combinations by mixing both the external MIDI synth and internal synthesizer voices.

WARNING: The bottom 1/4" and front panel 1/8" LINE and MONITOR jacks are stereo outputs requiring stereo Tip-Ring-Sleeve (TRS) cables.

NEVER plug mono cables into any of these jacks because they will short out the right audio channels and possibly damage the driver amplifiers.

Main Modules




D-Lev main modules.

The D-Lev consists of many interconnected functional modules as shown above.

The player's hands interact with the pitch and volume fields which are generated by the digital oscillators and resonant circuits. After extensive filtering and mathematical manipulation, this pitch and volume information is fed to a digital audio synthesizer to generate sounds. The LED visual tuner displays highly responsive pitch and volume information to visually assist setup and playing. Pitch preview provides a means for the player to personally monitor pitch even when in full volume cutoff. This gives the player a little advance information to confidently start a note on the right pitch without "hunting". On top of all this, the D-Lev sound palate is easily expanded by connecting it to an external MIDI synthesizer.

The following sections will discuss the D-Lev modules in more detail, ending with a focus on the DSP music synthesizer and the modules within it.

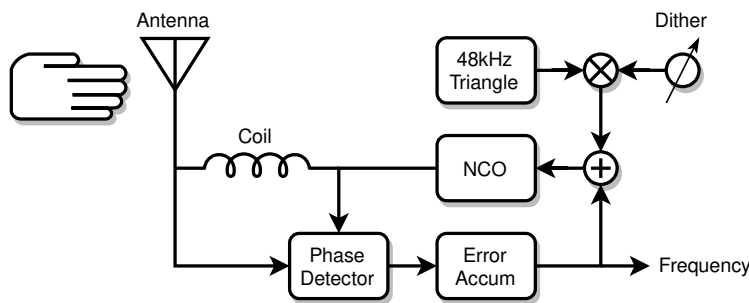
Pitch & Volume Fields



Pitch & volume field generation and processing.

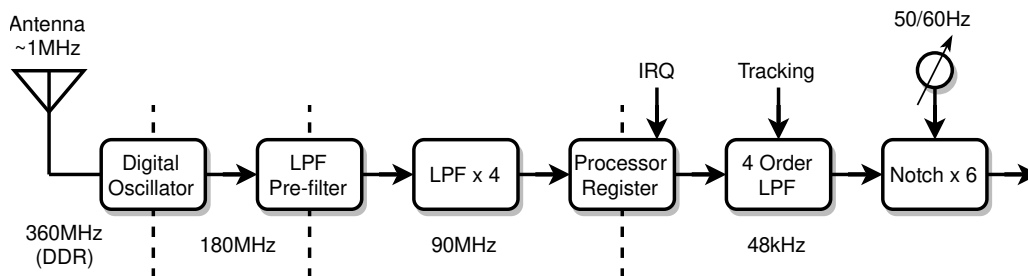
Description

As the pitch or volume hand approaches the antenna, the capacitance increases slightly, which slightly lowers the frequency of the oscillator connected to the antenna. The oscillator frequency is translated into digital data representing hand location, which is then used to control the pitch and volume of the music synthesizer.



Pitch & volume digital phase-locked loop (DPLL) oscillator.

The D-Lev oscillators are almost entirely digital, which makes them exceptionally stable. In each oscillator, a numerically controlled oscillator (NCO) drives a coil, which resonates with the antenna capacitance. The phase difference is fed back to keep the NCO frequency locked at the coil & antenna resonance frequency. This produces a large (though harmless) radio frequency voltage on the antenna, which generates a large playing field and drowns out environmental interference.



Pitch & volume field filtering (and sampling rates).

The oscillator forms a 1st order low-pass filter for (phase) noise. Following this is a 1st order low-pass pre-filter, then four more 1st order low-pass filters. These hardware filters reduce noise, increase resolution, and allow the software to more accurately sample the oscillator frequency. In software, this is passed through a 4th order low-pass tracking filter, followed by six 2nd order notch filters to filter out AC mains hum. Altogether these stages form one big 22nd order filter which removes the vast majority of interference while remaining highly responsive.

Following the oscillators and filters are mathematical stages which linearize, scale, and offset the oscillator frequencies. With minor variations, this entire oscillator / filtering / mathematical processing gauntlet is identical for the pitch and volume fields.

Controls

All controls are on the **V_FIELD**, **P_FIELD**, and **SYSTEM** pages. Note that these are all system controls, and can be saved to and recalled from any of four system presets.

Vcal	[0:255]	Vrev	[0:1]
Lin	[-15:15]	Dith	[0:7]
Ofs-	[0:127]	Ofs+	[0:127]
Sens	[0:127]		V_FIELD

V_FIELD page controls.

Pcal	[0:255]	Cent	[-99:99]
Lin	[-15:15]	Dith	[0:7]
Ofs-	[0:127]	Ofs+	[0:127]
Sens	[0:127]		P_FIELD

P_FIELD page controls.

Wait	[0:99]	50Hz	[0:1]
		P<>U	[0:1]
Stor	[0:3]	Erev	[0:1]
Load	[0:3]		SYSTEM

SYSTEM page controls.

Vcal & **Pcal** correct the far-field linearity, and are identical to the “null” type controls found on analog Theremins. Increasing the value reduces far-field sensitivity (moves the “null point” farther behind your body). These are offsets to the ACAL process, and you will need higher settings here to compensate if you hold your body or hands closer to the antennas during ACAL than during normal play. Once **Vcal** and **Pcal** are properly set, save the system preset, and the next time you do an ACAL you should have little need to adjust them. Once you develop the correct technique for holding your arms in a re-ceded position during ACAL you may find that you never need to touch Vcal or Pcal.

Lin tweaks the near-field linearity. Positive values produce a small hump of higher sensitivity, negative values produce a small trough of lower sensitivity. This is an “advanced” feature and the effect on linearity is modest, so if you are uncertain as to how to tune it simply set it to zero. **Lin** interacts strongly with **Ofs-**, **Sens**, and **Ofs+**, so expect some back and forth when adjusting it.

Ofs- moves the field around. Increase the value to move the field away from the antenna, decrease the value to move it closer. This control should be used to center up the operation of the **Sens** control, and not to position the final playing field.

Sens sets the sensitivity of the field, and acts as a sort of gain control to your hand gestures. For “traditional” analog Theremin pitch sensitivity, set this so that an open hand / closed fist produces about an octave.

Ofs+ moves the field around in a manner similar to **Ofs-**. You may be wondering why there are two controls which essentially do the same thing? **Ofs-** should be used to center up the action of **Sens** on the mid-field. **Ofs+** should be used to position the field where you want it to be when playing. This is particularly useful for bringing the entire pitch field closer to the player’s body lessen arm stress and to avoid any risk of bumping into the pitch antenna on the highest of high notes.

Dith sets the amplitude of a 48kHz triangular dither signal injected into the field oscillator. Dither is used to break up small non-responsive “sticky points” in the field. Higher values can cause instability, so use the minimum necessary to fix a sticky situation. Otherwise **Dith[0]**, which injects a nominal amount of dither, is recommended. The tuner is particularly useful for catching any pitch or volume transitions that appear to jump or stick, thereby indicating the need for a higher dither setting.

Vrev reverses the sense of the volume field. **Vrev[0]** gives a traditional “hand farther away is louder” response; **Vrev[1]** a non-traditional “hand closer is louder” response.

50Hz selects the mains hum filter frequency. Use **50Hz[0]** in areas which have 60Hz AC power, and **50Hz[1]** in those which have 50Hz AC power.

P<>V swaps the pitch & volume antennas. **P<>V[1]** performs the swap, after which you should perform an ACAL. Left handed theremins will require the setting **P<>V[1]**.

Wait sets the ACAL wait timeout in 1/10 second increments, so for example **Wait[10]** gives a timeout of one second. This gives the player time to press the calibration button and move to the receded hand position before the calibration is actually performed. Wait has no effect if pedal calibration is used; this is instantaneous.

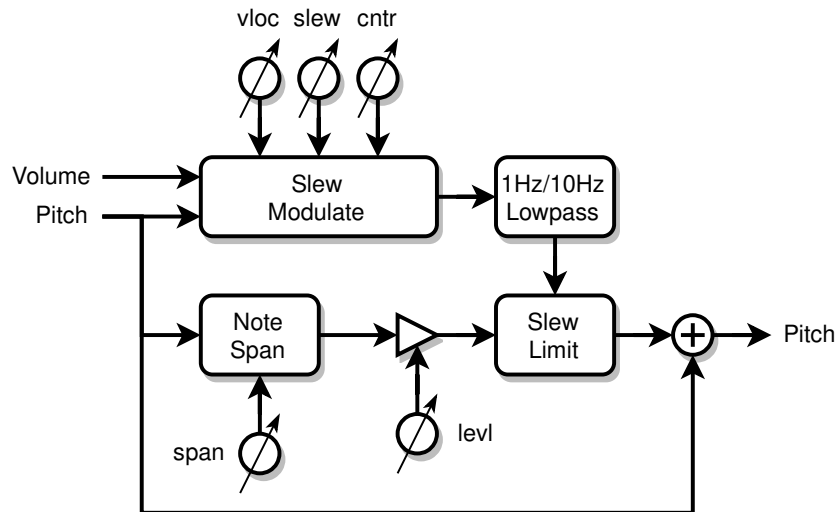
Tips & Tricks

Back up your favorite system preset to at least one other system preset slot, and preferably more if they are not otherwise used. Whenever you substantially change any system settings, save your new system preset to a preset file on your computer using the librarian. You should also occasionally write down your **V_FIELD** and **P_FIELD** settings somewhere. Theremin play relies on muscle memory, which in turn relies on a certain stable setup of the fields, so the inadvertent overwriting of your field settings could be a minor disaster if no back-up exists.

Consult the section on setting up the fields for more insight into pitch field sensitivity options and volume field reversal pros and cons.

Pitch Corrector

Description



Pitch corrector block diagram.

The pitch corrector is chromatic (all 12 notes in the octave are corrected identically) and employs multiple mechanisms to mask its audible functioning. Since it corrects to absolute pitch, it is most effective when used actively in conjunction with the tuner during play.

Controls

All controls are on the **PITCH** page.

vloc	[-63:63]	levl	[0:31]
slew	[0:31]	span	[0:31]
cntr	[0:31]	prvw	[-12:11]
Post	[0:1]		PITCH

PITCH page controls.

levl sets the pitch correction strength. To enable pitch correction, both the **span** and **levl** knobs must be non-zero. Typically **levl** is maxed out, but if you experience phasing effects with external audio sources due to overly perfect pitch correction, reduce **levl** to intentionally introduce some minor pitch tracking error. **levl[0]** will disable pitch correction, and this provides a quick and simple means of disabling correction while retaining all of the other carefully-tuned correction parameters.

span sets the note width over which pitch correction is applied. Typically **span** is maxed out, but lower values will limit correction to a narrower span located at the center of each note. **span[0]** will disable pitch correction.

slew slows down the correction rate; larger values give longer and slower correction times. For maximum correction this should be set as low as possible for the piece you are playing. Use **slew[0]** for hard quantization.

cntr sets the strength of slew rate modulation. The modulation speeds up correction when the pitch hand is near the center of a note and slows down correction near the edges. This enables slower note transitions with lower **slew** settings, imparts a “magnetic” feel to the note centers, and helps to hide the otherwise constant correction rate pulling the notes this way and that at the note boundaries.

vloc sets a volume threshold to speed up pitch correction. If the audience can’t hear anything, we may as well hard quantize and go instantly to the note center so that the next note comes in on pitch. **vloc[1:63]** adjusts this switching threshold from -48dB to 0dB. For the opposite effect (quantization at high volumes switching to pitch correction at lower volumes) use negative values of **vloc**. **vloc[0]** disables this switching to give full-time pitch correction.

Post is repeated here for convenience because moderate pitch correction settings can be quite subtle and difficult to hear – using **Post[1]** lets you see what pitch correction is up to on the tuner. Otherwise use **Post[0]** so that the pitch hand position is displayed on the tuner – this helps you hit the note centers when playing, which helps the pitch corrector operate optimally.

Tips & Tricks

vloc can be an aid to pitch preview by hard quantizing the pitch preview signal.

Some typical settings:

vloc	10	levl	31
slew	20	span	31
cntr	12	prvw	0
Post	0		PITCH

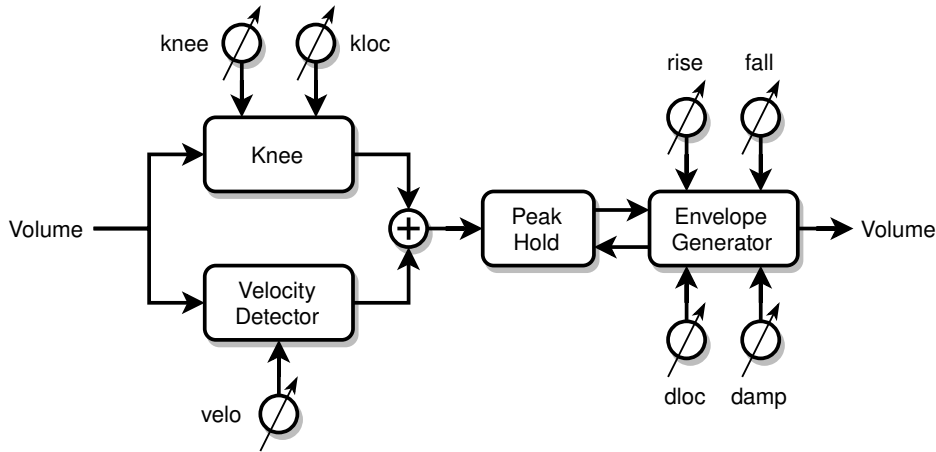
Pitch correction

vloc	0	levl	31
slew	0	span	31
cntr	0	prvw	0
Post	0		PITCH

Pitch quantization

Envelope Generator

Description



Envelope generator block diagram.

The envelope generator is used to alter the static and dynamic response of volume hand movements. It consists of a knee unit and a velocity detector unit in parallel, followed by a peak hold / envelope generator. The knee unit is used to introduce a secondary linear response to the volume hand location – which sounds bad but is actually highly desirable when done in a controlled manner. The velocity detector enables more nuanced percussive play. The peak hold and envelope generator together create rapid – or very slow and smooth – volume changes which would be difficult or impossible to accomplish with hand movements alone.

Controls

All controls are on the **VOLUME** page.

kloc	[0:63]	rise	[0:63]
knee	[0:31]	fall	[0:63]
velo	[0:31]	damp	[0:63]
dloc	[0:63]		VOLUME

VOLUME page controls.

knee causes the volume, when it is below a certain volume threshold – set by **kloc** – to respond more quickly to hand movement. This is known as *expansion*, and the volume response above **kloc** is unaffected. Most natural sounds quickly ramp up to a certain loudness, with a more expressive region above that. Moderate settings of **knee** facilitate this with smaller, slower, and more natural hand movements, while retaining good control over dynamics. Very high settings of **knee** cause the volume to instantly “pop” to the **kloc** loudness, which can be used percussively. **knee[0]** defeats the knee.

kloc sets the **knee** location in the volume field. The range is -48dB to 0dB (**kloc[0]** to **kloc[63]**, respectively). An upper expressive region of 10dB or so (**kloc[50]**) works well for most voices.

velo controls the volume hand velocity gain, where the speed of the volume hand is sensed and added to the overall volume. Velocity sensing is in one direction only, when the hand is moving in the direction from quiet to loud, and continuous over the full range of hand movement (sensing is not restricted to a threshold point). Adding velocity to a high gain knee provides nuanced control to percussive presets.

rise sets the volume rise time, with higher settings giving a slower volume ramp up rate. Rise will continue until the output matches the peak hold value, after which it falls. **rise[0]** will defeat ramping up.

fall and **damp** set the volume fall time, with higher settings giving a slower volume ramp down rate. If the volume hand is past the location given by **dloc** then the fall rate is set by **damp**, otherwise **fall** is used. Two separate decay rates provide better control over percussive presets.

dloc sets the **fall / damp** switching location in the volume field. The range is -48dB to 0dB (**dloc[1]** to **dloc[63]**, respectively). **dloc[0]** disables the switch to **damp**, which keeps you from having to set **damp = fall** in order to defeat the secondary fall rate.

Tips & Tricks

The application of some knee is also very useful if you tend to play with your body close to the D-Lev. Without it, your body capacitance can be more of a confounding factor, and larger arm movements would be required to produce the full range of audio.

High **rise** values together with **fall[0]** will generate a reverse envelope.

Envelope generation can be fully defeated by setting **rise[0]**, **fall[0]**, and **damp[0]** (and/or **dloc[0]**). Once defeated it will not interfere in any way with normal play.

Experiment with setting damp higher and lower than fall.

Some typical settings:

kloc	50	rise	0
knee	7	fall	16
velo	0	damp	0
dloc	0		VOLUME

Human vocal

kloc	52	rise	0
knee	14	fall	38
velo	6	damp	0
dloc	0		VOLUME

Cello

kloc	28	rise	0
knee	31	fall	51
velo	12	damp	38
dloc	1		VOLUME


Percussive with damping

kloc	28	rise	50
knee	31	fall	0
velo	12	damp	0
dloc	0		VOLUME

Reverse envelope

MIDI Module

Description



MIDI module block diagram.

The MIDI module takes the pitch and volume hand information and sends serial MIDI OUT data to the MIDI IN of an external music synthesizer. This is strictly a MIDI transmit module, there are no MIDI IN or THRU connections nor processing.

Controls

Most controls are on the **MIDI** page.

vloc	[0:63]	velo	[0:31]
cloc	[0:63]	cc	[0:31]
bend	[0:127]	chan	[0:16]
oct	[-7:7]		MIDI

MIDI page controls.

chan selects the MIDI channel. MIDI channels are sort of sub-synthesizers with their own independent controls. Valid channels are 1 through 16. **chan[0]** disables the MIDI module, which can be handy when changing **cc**.

cc selects the MIDI control change manipulated by the volume hand. Here are some common MIDI CC's:

CC	Control
0	Bank Select
1	Modulation Wheel
2	Breath controller
3	- Undefined -
4	Foot Pedal
5	Portamento Time
6	Data Entry
7	Volume
8	Balance
9	- Undefined -
10	Pan
11	Expression
12	Effect Control 1
13	Effect Control 2

Common MIDI CC values.

cloc sets the **cc** start location in the volume field. The range is -48dB to 0dB (**cloc[0]** to **cloc[63]**, respectively). The full range of control is scaled to fit into the selected volume range. For example: **cloc[32]** will set the cc start point to -24dB; when the volume is below this the CC data is zero; above this the CC data goes smoothly from 0 to a maximum at 0dB.

velo controls the note-on velocity gain. The speed of the volume hand as it crosses the **vloc** threshold (to play a new note) is converted into MIDI note-on velocity. Higher **velo** settings make it easier to play higher velocities. Some special modes here are **velo[0]** which transmits a fixed velocity of 64 (the MIDI standard for non-velocity instruments) and **velo[31]** which transmits a fixed velocity of 127 (this may work better for some synths / patches).

vloc sets the note-on velocity sense point and note on / off trigger location in the volume field. The range is -48dB to 0dB (**vloc[0]** to **vloc[63]**, respectively). Setting **vloc** below that of **cloc** can help hide note on/off trigger events, which can make the voice being played sound less percussive and more Theremin-y.

bend sets the MIDI pitch bend span in terms of semitones. **bend[12]** will produce a +/- 1 octave pitch bend range, and so on. Pitch bend is re-centered whenever a note is played (when the vloc threshold is exceeded), and when playing exceeds the pitch bend range the note will be re-triggered with pitch bend re-centered on the new note. In order to function correctly, the value of **bend** needs to match the pitch bend span setting in the external synthesizer. **bend[0]** defeats pitch bend and note re-triggering, so notes just stick at the pitch they were played at until done.

oct sets the relative octave of the signal; this is an offset to the oscillator octave setting.

Mute (located on the **D-LEV** page), when switched from 0 to 1, sends MIDI note-off (note-on with velocity=0), as well as MIDI Channel Mode Message 120 (*All Sound Off*) for the current MIDI channel (set by **chan**).

Tips & Tricks

Set **bend** as high as possible to get the largest smooth and continuous pitch range. This value will likely be limited more by the external synth than by the range of **bend** (which can accommodate a pitch bend range of up to +/- 10.5 octaves).

MIDI controls usually remember where you last set them, which isn't always desirable. To get around this use **Mute[1]** or **chan[0]** to disable MIDI transmission when changing **cc**.


Radically mismatched pitch bend ranges can be used as a special effect.

MIDI CC data can be 7 or 14 bits. 7 bits are usually sufficient for volume type control, but 14 bits provide increased resolution for pitch type control. The MIDI module sends 14 bits, which is transmitted MSB first (as per the MIDI standard) in a contiguous running status block to minimize LSB glitching.

The D-Lev's internal sound generator still functions even if an external MIDI synthesizer is used, which may be useful for creating blended voices.

LED Visual Tuner

Description



LED tuner display.

The LED visual tuner provides zero lag real-time pitch and volume feedback during play. The pitch circle indicates the note within the octave being played, the octave is shown on a 7-segment display, and the volume level is indicated with a 4-LED thermometer type bargraph with 12dB graduations. Pulse-width modulation (PWM) dimming is employed to dramatically increase the resolution by using peak brightness to indicate the exact note centers.

Controls

All controls are on the **DISPLAY** page.

Brit	[0:31]	Cent	[-99:99]
LCD	[0:31]	Note	[-11:11]
Qant	[0:4]	Oct	[-15:15]
Post	[0:1]		DISPLAY

DISPLAY page controls.

Brit sets the brightness of all of the tuner LEDs. If you find the tuner distracting, **Brit[0]** turns all the LEDs off.

LCD controls the LCD backlight brightness. **LCD[0]** turns off the backlight.

Qant controls the tuner pitch brightness quantization. Higher quantization settings cause the brightness to be more centered on one LED, rather than smeared out among two or three LEDs, which makes it easier to see at a glance the note being played in the presence of pitch inaccuracies and/or vibrato. I use **Qant[1]**.

Post selects the tuner pitch input. **Post[0]** displays the pitch pre-correction; **Post[1]** displays the pitch post-correction. **Post[1]** is useful for visualizing and adjusting the often quite subtle sounding pitch correction controls. I recommend **Post[0]** for normal playing as this shows you exactly where your pitch hand is.

Cent offsets the global frequency reference in cents (1/100 note). This is a system-wide offset for the entire D-Lev including all oscillators and filters, and not just the tuner.

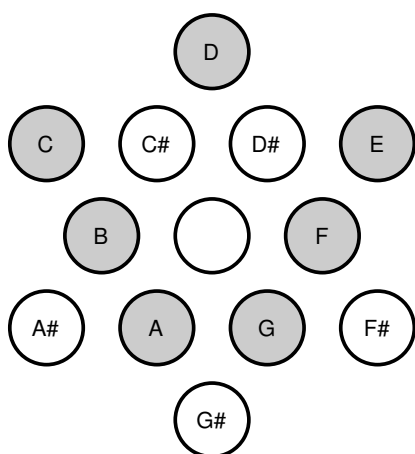
Note offsets the tuner pitch display only, and is handy if you prefer to view, say, the key of C pattern while actually playing in another key.

Oct offsets the tuner octave display only.


Tips & Tricks

The pitch circle is composed of staggered inner and outer rings. Within a ring the notes are spaced two semitones apart, and the rings are offset from each other by one semitone. So sequential notes within a ring are whole steps, and switching rings produces a half step.

The pitch circle LEDs form two distinct patterns when playing in major or natural minor keys, dubbed the “Crown” and “House” patterns:



C / Am “Crown” pattern.



C# / A#m “House” pattern.

It’s easy to infer the key of the music you are playing by recognizing these patterns and their rotation. Once recognized, the patterns can be a powerful aid to improvisation.

The location of C on the pitch circle was chosen mainly to give a vertically symmetric pattern on the pitch circle with C lying on the outer ring. Conveniently, this also places C near the octave display, the value of which morphs in brightness over the B/C interval. Clockwise rotation is used to indicate increasing pitch, which follows the convention of the controls.

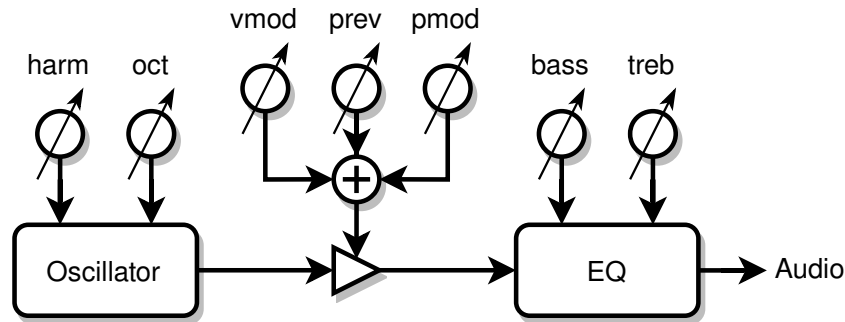
The point at which all volume bargraph LEDs just turn off is -48dB. Similarly, the bottom LED fully on with the others fully off is -36dB; two bottom LEDs fully on is -24dB; three LEDs on is -12dB, and all four LEDs on is 0dB, or the loudest possible audio.

-48dB (1/256) is right at the edge of audibility, so when all the volume bargraph LEDs are off the volume is almost certainly low enough that the audience won’t hear anything. Actual volume control continues on to -96dB to ensure inaudibility, but this isn’t displayed.

The D-Lev tuner is sort of the inverse of a guitar tuner (where the note is displayed as a character and the sharp/flat information is displayed graphically). Guitar tuners are often quite sluggish, and their displays aren’t really designed for real-time intuitive feedback during play.

Pitch Preview

Description



Pitch Preview block diagram.

Pitch preview is a separate oscillator which provides a way for the player to more effectively monitor the pitch during play, particularly during passages of low volume and silence. This is a feature commonly found on professional Theremins, but this implementation offers more control so you can tailor it to your needs.

Controls

Most controls are on the **PREVIEW** page.

vmod	[-63:63]	fmod	[-63:63]
prev	[0:63]	treb	[-31:31]
harm	[-31:31]	bass	[-31:31]
oct	[-7:7]		PREVIEW

PREVIEW page controls.

prev sets the maximum volume of the preview signal. **prev[0]** disables the pitch preview section.

vmod modulates pitch preview volume with the volume hand. Positive values of **vmod** increase pitch preview volume with increasing volume; **vmod[32]** makes the pitch preview volume perfectly track the volume of the preset voice. Negative values increase pitch preview volume with decreasing volume and produce a “ducking” effect where the pitch preview signal fades out as the preset voice is fading in, and vice-versa. Some find this useful, and others find it confusing. **vmod[0]** is the safest starting value and it replicates the behavior of more basic pitch preview implementations on other theremins.

pmod modulates pitch preview volume with the pitch hand. Positive values of **pmod** increase pitch preview volume with increasing pitch. Negative values increase pitch preview volume with decreasing pitch, which can help when higher notes seem much louder than bass notes. Depending on the timbre of the preset, **pmod[-15:-25]** may give the perception of more balanced pitch preview loudness over the full pitch range. The modulation pivot point is A440.

harm sets the fixed harmonic level; negative values here give odd rather than all harmonics and have a hollow sound, which may make pitch identification a little more difficult. Use

higher harmonic levels (positive or negative) to better hear the pitch, and consider using odd harmonics to make the pitch preview timbre more distinct.

oct sets the relative octave of the signal; this is an offset to the global octave setting.

oct[1] will shift the octave up one and may enhance pitch perception of low notes by making the pitch preview signal more distinct.

bass and **treb** EQ the lows and highs. Increasing **treb** may enhance pitch perception. Always be conscious that adding bass or treble may push a preset into digital distortion at high volumes, and it may be necessary to decrease the oscillator amplitude parameter to compensate and stay safely below any distortion level.

prvw controls the mixing of the pitch preview signal, as well as the pitch correction and quantization of it, and is located on the **PITCH** page. The pitch preview signal is injected into the L&R (left & right) **monitor** output only (not the L&R **line** output). Any applied pitch correction is the same as that applied to the preset voice; any hard pitch quantization (note stepping) is entirely separate from pitch correction:

prvw	L	R	L PP	R PP	PC	HQ	Osc
11	✓	✓	✓	✓		✓	
10	✓	✓		✓		✓	
9	✓	✓	✓			✓	
8	✓	✓				✓	✓
7	✓	✓	✓	✓			
6	✓	✓		✓			
5	✓	✓	✓				
4	✓	✓					✓
3	✓	✓	✓	✓	✓		
2	✓	✓		✓	✓		
1	✓	✓	✓		✓		
0	✓	✓			✓		✓
-1			✓	✓	✓		
-2	✓		✓	✓	✓		
-3		✓	✓	✓	✓		
-4	✓	✓	✓	✓	✓		
-5			✓	✓			
-6	✓		✓	✓			
-7		✓	✓	✓			
-8	✓	✓	✓	✓			
-9			✓	✓		✓	
-10	✓		✓	✓		✓	
-11		✓	✓	✓		✓	
-12	✓	✓	✓	✓		✓	

prvw settings (PP=pitch preview, PC=pitch correction, HQ=hard quantization).

Applying either no pitch correction (in the case where the preset voice uses pitch correction) or hard quantization to the pitch preview signal will produce phasing between the pitch preview and preset voices, providing further pitch cues.

Tips & Tricks


If you don't use pitch preview as an aid to playing on pitch, its output can be employed as a fourth independent oscillator (use **prvw[0]**, **prvw[4]**, or **prvw[8]** and **vmod[32]**) to add timbre, subsonics, or phase / beating effects to the preset voice.

You may hear quantization switching in at low volume if the **vloc** control on the **PITCH** page is positive and you are using a pitch preview mode with pitch correction / quantization different than the synthesized voice. This is normal behavior and may actually be useful.

Different preset voices benefit from different pitch preview settings, and this is why pitch preview settings are saved with each voice preset, rather in the system preset.

Some players are able to use a single earpiece to monitor the pitch preview signal, but for others using a single ear makes accurate pitch identification difficult. Placing a small powered monitor behind or to one side of the playing position and aimed away from the audience allows the player to hear the pitch preview audio without the audience hearing it. This eliminates the need use an earpiece and can be less distracting.

DSP Synthesizer



DSP synthesizer modules.

The synthesizer section uses digital signal processing (DSP) to produce sounds via a combination of FM, subtractive, and modal synthesis. A group of digital oscillators and a digital noise source feed a variety of digital filters.

The oscillators and noise source each have their own bass and treble controls, as well as dedicated multi-mode filters. Following this are global bass and treble controls, a formant filter bank, and the inharmonic resonator.

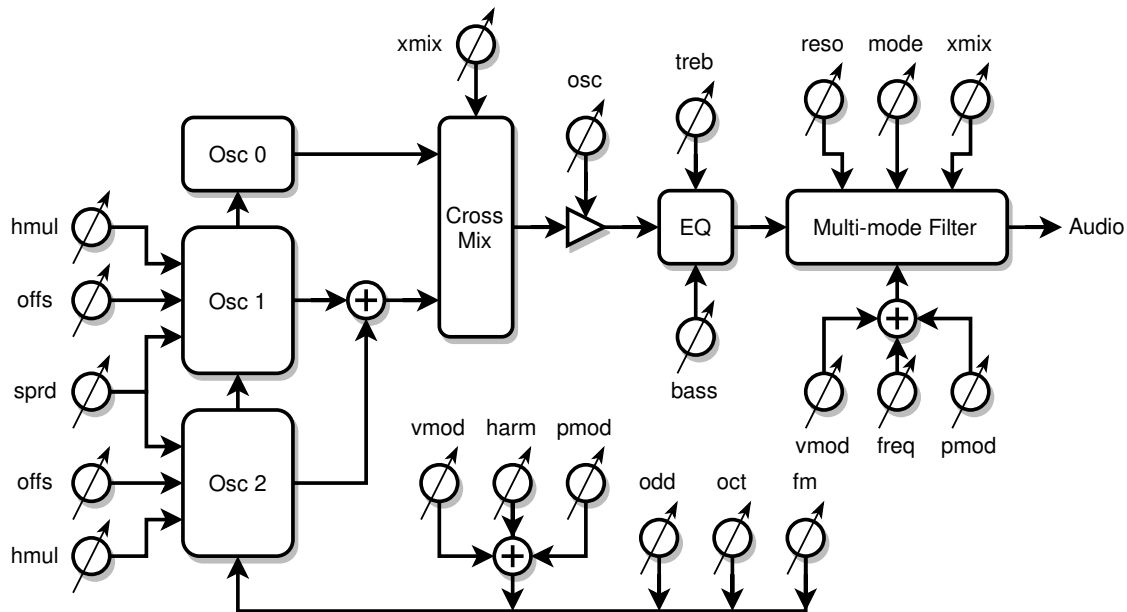
The formant bank is used to mimic human vocals, but it's equally useful at providing the resonances found in wind and stringed instruments, as well as bells and other percussion.

The inharmonic resonator mainly provides resonance fill, pseudo stereo, and ambiance effects, but it can also do passable drums and gongs, and can even human vocals all by itself.

The following sections discuss these synth modules in more detail.

Triple Oscillator

Description



Triple oscillator block diagram.

The triple oscillator consists of a bank of three oscillators, an EQ section, and a multi-mode filter. The oscillators generate non-aliased harmonic content via phase modulation, with a blend of *all* and *odd* harmonics on tap. The harmonic level can be modulated via the pitch or volume hand, and the oscillator phase modulation inputs and outputs are interconnected for more complex timbres. Oscillator 0 is the “center” and always on the pitch being played, 1 & 2 can be offset from it in multiple ways to give a thicker sound.

Controls

All controls are on the `0_OSC`, `1_OSC`, and `FLT_OSC` pages.

vmod	[-63:63]	pmod	[-63:63]
harm	[0:31]	treb	[-31:31]
odd	[0:31]	bass	[-31:31]
oct	[-7:7]		0_OSC

`0_OSC` page controls (common to all oscillators).

harm sets the static harmonic level. **vmod** and **pmod** can be used to modulate this level up or down. Unmodulated **harm[0]** gives a sine wave.

vmod modulates the harmonic level with the volume hand. Positive values of **vmod** increase harmonics with increasing volume; negative values increase harmonics with decreasing volume. Harmonics for most natural sounds increase with increasing volume (i.e. positive **vmod**). Setting **harm[31]** and **vmod[31]** will give sine to full harmonics over the -24dB to 0dB volume range; **vmod[42]** reduces this range to -12dB to 0dB. **harm[18]** and **vmod[16]** are typical for human vocals.

pmod modulates the harmonic level with the pitch hand. Positive values of **pmod** increase harmonics with increasing pitch; negative values increase harmonics with decreasing pitch. The modulation pivot point is A440.

odd crossfades between all and odd harmonic content. **odd[0]** gives all harmonics, **odd[31]** gives only odd harmonics.

oct sets the octave of the oscillators; this is a master setting which also affects pitch preview, but does not affect any filter cutoff frequencies, nor the tuner.

bass and **treb** EQ the lows and highs.


```

offs[-127:127]  offs[-127:127]
hmul[-127:127]  hmul[-127:127]
fm      [0:31]  sprd      [0:63]
xmix    [0:31]

```

1_OSC page controls (mostly oscillator 1 & 2).

offs are used to offset the pitches of oscillators 1 & 2 (+/- 1 octave). These are multiplicative offsets, so pitch ratios are maintained. They are scaled to increase the resolution around the unison (**offs[0]**) and octave (**offs[-127, 127]**) regions:



*Cubic plateau scaling of **offs** and **hmul**.*

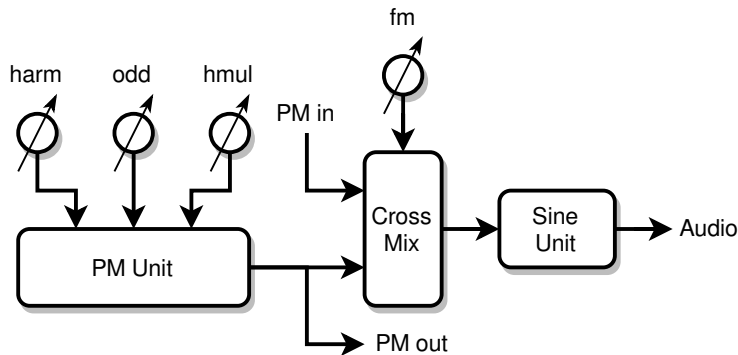
Steps	C scale	Ratio	Interval	offs / hmul
0	C	1	unison	0
1	C#	1.0595	minor second	35
2	D	1.1225	major second	44
3	D#	1.1892	minor third	53
4	E	1.2599	major third	55
5	F	1.3348	perfect fourth	60
6	F#	1.4142	diminished fifth	64
7	G	1.4983	perfect fifth	68
8	G#	1.5874	minor sixth	72
9	A	1.6818	major sixth	77
10	A#	1.7818	minor seventh	84
11	B	1.8877	major seventh	93
12	C	2	octave	127

*Some useful **offs** and **hmul** values for various musical intervals.*

hmul are multiplicative pitch offsets of the phase modulation units within oscillators 1 & 2 (+/- 1 octave) and are relative to their associated **offs** values (**offs** offsets **hmul**). They are scaled the same as **offs**, and are useful for introducing pitch-dependent beating and ring-modulation type harmonics.


sprd introduces an additive pitch offset to oscillator 1 and a subtractive pitch offset to oscillator 2; the pitch of oscillator 0 is unaffected. It's useful for unison and beating type effects when you want the beating to be independent of pitch.

fm cross-fades the phase modulation input with the output of the phase modulation (PM) unit within each oscillator:



Single oscillator block diagram.

The phase modulation outputs of oscillator 1 & 2 are combined and sent to the phase modulation input of oscillator 0, and the phase modulation output of oscillator 0 is sent to the phase modulation inputs of oscillators 1 & 2. With **fm[0]** all oscillators only use their internal phase modulation units to create harmonics; **fm[31]** makes them fully cross-coupled.



Triple oscillator phase modulation cross-coupling.

fm can radically transform the timbre and introduce some fairly wild harmonics, particularly when **offs** and/or **hmul** are non-zero. Use it to produce scratchy records, trumpeting elephants, UFO invasions, and other sonic mayhem.

xmix cross-fades the audio output of oscillators 1 & 2 with that of oscillator 0. **xmix[0]** gives only oscillator 0; **xmix[31]** only 1 & 2. This control is scaled so that **xmix[16]** produces equal volume for all three oscillators.

vmod	[-63:63]	pmod	[-63:63]
freq	[27:7040]	osc	[0:63]
mode	[-4:4]	xmix	[0:31]
reso	[0:63]		FLT_OSC

FLT_OSC page controls (multi-mode filter).

osc sets the oscillator mix audio level. Reduce this if the filter is overloading.

freq sets the center or corner frequency of the multi-mode filter. The frequency is displayed in Hz, and the increments are 1/2 semitone.

vmod modulates the filter frequency with the volume hand. Positive values of **vmod** increase the frequency with increasing volume, and negative values increase filter frequency with decreasing volume.

pmod modulates the multi-mode filter frequency with the pitch hand. Positive values of **pmod** increase the filter frequency with increasing pitch, and negative values increase filter frequency with decreasing pitch. The modulation pivot point is where the pitch hand normally plays an A440, so setting **freq[440]** and **pmod[32]** will make the filter frequency track the oscillator frequency (this is only true if **oct[0]**; if **oct[1]** then use **freq[880]**, etc.).

reso sets the resonance of the multi-mode filter. This is analogous to Q (or the inverse of damping) but is scaled to give useful range of resonance behavior, from highly damped to infinite sustain. Higher **reso** increases peak gain, so you may have to reduce **osc** to keep the filter from overloading.

mode selects the filtering mode and order of the multi-mode filter:

mode	function	order
-4	notch	2
-3	high-pass	2
-2	band-pass	2
-1	low-pass	2
0	bypass	n/a
1	low-pass	4
2	band-pass	4
3	high-pass	4
4	notch	4

Filter mode functionality.

Higher orders are more effective at filtering in general and tend to sound “synthy.” When constructing a more natural sounding preset a second order filter is usually what you’re looking for.

The hollow “phasey” sounding notch filter modes benefit from lower reso settings, otherwise they can be hard to hear.

xmix cross-fades the filter input and output. **xmix[0]** completely bypasses the filter (as does **mode[0]**); **xmix[31]** gives full filter output. Mixing some input into the output can be useful when you need milder emphasis, but you usually want **xmix[31]** when using the filter.

Tips & Tricks

A little **odd** can often add some realism to a preset.

Here are some basic starting points (with all **1_OSC** controls zeroed out):

vmod	16	fmod	0
harm	18	treb	0
odd	0	bass	0
oct	-1	0_OSC	

Vocal

vmod	0	fmod	0
harm	31	treb	14
odd	3	bass	0
oct	-2	0_OSC	

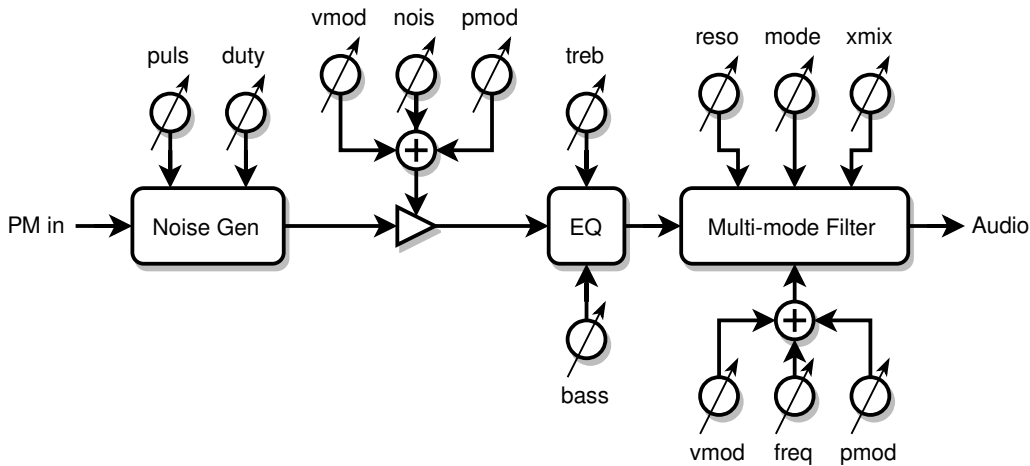
Cello

To reset the relative oscillator phases, save and reload the preset after messing with any controls that might influence multiple fixed phases.

The settings **offs[-127,0,127]**, **hmul[-127,0,127]**, and **sprd[0]** are guaranteed to have *no* phase drift, so any preset voice that relies on multiple oscillators using these settings should sound the same even over long periods of time.

Noise Generator

Description



Noise generator block diagram.

The noise generator consists of a noise source that can be modulated (turned on and off) by the phase modulation output of oscillator 0, followed by an EQ section and a multi-mode filter.

Controls

All controls are on the **NOISE** and **FLT_NOISE** pages.

vmod	[-63:63]	fmod	[-63:63]
nois	[0:63]	treb	[-31:31]
puls	[0:31]	bass	[-31:31]
duty	[0:31]		NOISE

NOISE page controls.

nois sets the noise audio level. **nois[0]** turns off the noise section. Reduce this if the noise multi-mode filter is overloading.

vmod introduces an upward hump or downward dip contour centered on -24dB to the volume response. **vmod[0]** defeats the contour, **vmod[24]** makes the noise constant above -12dB, and negative values of **vmod** cause the noise to ramp up more quickly to maximum. Extreme settings of this control interact strongly with **nois**.

pmod modulates the noise level with the pitch hand. Positive values of **pmod** increase the noise level with increasing pitch; negative values increase the noise level with decreasing pitch. The modulation pivot point is where the pitch hand normally plays an A440 (with oscillator **oct[0]**).

puls and **duty** modulate the noise level with the phase modulation output of oscillator 0. **puls** controls the depth of modulation: **puls[0]** turns modulation entirely off giving you constant noise; below **puls[8]** or so the noise doesn't turn fully off; above this the noise is fully modulated and higher **puls** values will tend to decrease the on time. Whereas increasing **duty** will increase the on time. These controls interact quite a bit, so you may

have to play with them to find a setting you like. The oscillator harm level will influence modulation depth too.

bass and **treb** EQ the lows and highs. Very often you want to roll off the highs.

```
vmod [-63:63]  fmod [-63:63]
freq [27:7040] nois  [0:63]
mode  [-4:4]   xmix  [0:31]
reso  [0:63]   FLT_NOISE
```

FLT_NOISE page controls (multi-mode filter).

nois is a copy of the same control on the **NOISE** page (for your convenience).

freq, **vmod**, **pmode**, **reso**, **mode**, **xmix** – same as the triple oscillator multi-mode filter (see that section for descriptions).

Tips & Tricks

When exploring the effects of **puls** and **duty** it can be helpful to turn the oscillator volume off and maximize the harmonic level while listening to subsonic pulses: **osc[0]**, **harm[31]**, and **oct[-7]**.

A good place to start with human breath and vocal noise is to roll off the highs and set noise filter cutoff to 2nd order band-pass at a mouth resonance:

```
vmod  12  fmod  -10
nois  16  treb  -17
puls   2  bass   0
duty  14          NOISE

vmod   4  fmod   0
freq 1281 nois  16
mode  -2  xmix  31
reso   0  FLT_NOISE
```

Female vocal noise settings.

The settings are similar for human humming, but with the noise filter cutoff set to a nasal resonance:

```
vmod  14  fmod  -16
nois  25  treb  -4
puls  12  bass   0
duty   8          NOISE

vmod   4  fmod  17
freq 1077 nois  30
mode  -2  xmix  31
reso   0  FLT_NOISE
```

Male humming noise settings.

This makes for a good cello bowing noise:


```
vmod  14  fmod  -20
nois  50  treb  -31
puls  17  bass   0
duty   0          NOISE

vmod   4  fmod   0
freq 2960 nois  48
mode   2  xmix  31
reso   0  FLT_NOISE
```

Cello bowing noise settings.

Formant Bank

Description



Formant bank block diagram.

The formant bank is a parallel combination of eight second order band-pass filters. Each filter has its own center frequency and level control, but the filters are paired up in order to share common center frequency modulation and resonance controls.

Controls

All controls are on the **0_FORM**, **1_FORM**, **2_FORM**, and **3_FORM** pages.

vmod	[-63:63]	fmod	[-63:63]
freq	[27:7040]	levl	[0:63]
freq	[27:7040]	levl	[0:63]
reso	[0:63]		0_FORM

0_FORM page controls (other pages are the same).

freq sets the center frequency of the filter. The frequency is displayed in Hz, and the increments are 1/2 semitone. Each **freq** is associated with the **levl** to the right.

levl sets the audio input level. **levl[0]** turns off the filter. Reduce this if the formant filter is overloading. Setting all eight **levl[0]** will completely bypass the formant filter bank.

vmod modulates the filter center frequency with the volume hand. Positive values of **vmod** increase the frequency with increasing volume, and negative values increase filter frequency with decreasing volume.

pmod modulates the filter center frequency with the pitch hand. Positive values of **pmod** increase the filter frequency with increasing pitch, and negative values increase filter frequency with decreasing pitch. The modulation pivot point is where the pitch hand normally plays an A440, so setting **freq[440]** and **pmod[32]** will make the filter frequency track the oscillator frequency (this is only true if **oct[0]**; if **oct[1]** then use **freq[880]**, etc.).

reso sets the resonance of the filter. This is analogous to Q (or the inverse of damping) but is scaled to give useful range of resonance behavior, from highly damped to infinite sus-

tain. Higher **reso** increases peak gain, so you may have to reduce **levl** to keep the filter from overloading.

Tips & Tricks

You must set all eight **levl[0]** in order to bypass the formant filter bank. If any **levl** isn't zero then the bank will not be bypassed.

A spooky effect is to space out all eight high **reso** formants and **pmod** them so that the lowest formant modulates in one direction, the next modulates in the opposite direction, etc.

You can create resonant beating effects (like real bells do) by setting up two identical formants on two different pages, and applying a small amount of **pmod** or **vmod** to only one of them.

Strings and bells have lots of resonances, certainly more than 8! Use the formant bank for the lowest prominent resonances, then fill in the top resonances with the inharmonic resonator.

vmod	0	pmod	0	vmod	0	pmod	0
frea	174	levl	46	frea	359	levl	42
frea	246	levl	46	frea	452	levl	42
reso	23		0_FORM	reso	23		1_FORM

vmod	0	pmod	0	vmod	0	pmod	0
frea	678	levl	42	frea	1397	levl	42
frea	830	levl	42	frea	1710	levl	45
reso	23		2_FORM	reso	23		3_FORM

A set of (lower) cello formants.

Human vocals consist of (in order of increasing center frequency): throat radiation, two mouth resonances, and several nasal resonances. All have similar **reso** settings except for the throat radiation, which should have a lower **reso** setting to spread it out. Here too, the inharmonic resonator can be used to back-fill additional nasal resonances.

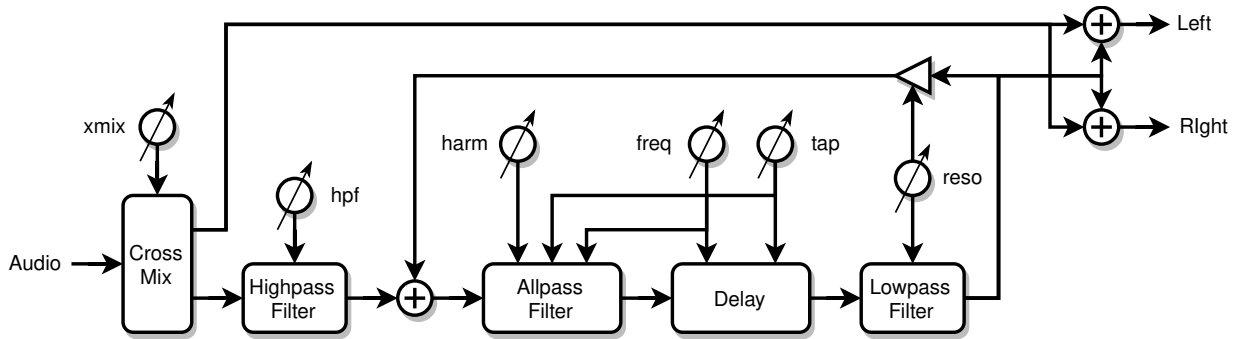
vmod	0	pmod	0	vmod	0	pmod	0
frea	905	levl	58	frea	1357	levl	58
frea	27	levl	0	frea	27	levl	0
reso	10		0_FORM	reso	13		1_FORM

vmod	0	pmod	0	vmod	0	pmod	0
frea	3623	levl	57	frea	320	levl	48
frea	2960	levl	46	frea	27	levl	0
reso	11		2_FORM	reso	7		3_FORM

A set of female vocal formants.

Inharmonic Resonator

Description



Inharmonic resonator block diagram.

The inharmonic resonator is a delay-based filter that consists of an all-pass filter in series with a simple delay. The output is recirculated to increase the Q / decay time and accentuate the resonant peaks. The all-pass filter breaks up the strict harmonic resonance locations, which enables it to produce realistic string body resonances, drums, gongs, and even human vocals. With longer delay times it provides room ambiance and a pseudo stereo effect which can considerably pep up a preset.

Controls

All controls are on the **RESON** page.

freq	[46:9600]	tap	[-63:63]
hpf	[27:7040]	harm	[-63:63]
mode	[-2:2]	xmix	[-31:31]
reso	[-63:63]		RESON

RESON page controls.

freq sets the inverse of the total delay time, where the lowest resonance roughly corresponds to the displayed frequency (Hz) if **reso** is positive. If **reso** is negative the frequency of the lowest resonance drops by an octave.

tap sets the ratio of all-pass delay time to simple delay time. **tap[0]** divides the total delay time equally between them; positive values make the all-pass delay longer than the simple delay; negative values make it shorter. Non-zero values can help jumble the resonance peaks.

harm sets the all-pass internal feedback, and is used to jumble the otherwise evenly spaced resonances. **harm[0]** disables all-pass feedback and converts the all-pass section into a simple delay. Some settings of **harm** drive the lowest resonance even lower, which can dramatically boost the bass.

reso sets the overall feedback. High positive and negative values will cause the resonator to “ring” like a high Q filter. Negative values invert the feedback and provide a more hollow sound. **reso** also controls a mild low-pass filter which tames “zinging” with low to moderate feedback levels.

hpf sets the cutoff frequency of the 4th order high pass filter at the input of the resonator. Quite often you don't want the resonator processing lower frequencies, such as when blending it with the formant bank or generating pseudo stereo.

xmix cross-fades the resonator input and output. **xmix[0]** completely bypasses the resonator; **xmix[-31,31]** give full resonator output. The input / output mix directly sets the height of the resonance peaks, and is also critical to getting a good pseudo stereo effect. Negative values of **xmix** invert the output, which may help when blending with the formant bank.

mode selects a number of useful input/output mixing modes:

mode	I/O mix
-2	Parallel pseudo stereo
-1	Parallel (formant fill)
0	Bypass (with hpf, variable with xmix)
1	Series (ambience / reverb)
2	Series pseudo stereo

*Resonator **mode** functionality.*

Series modes place the resonator in series with the formant bank, which is useful for ambience and reverb. Parallel modes place the resonator in parallel with the formant bank, which combines all the resonances together.

mode[-2,2] are pseudo stereo modes where the resonator output is added and subtracted from the input to produce L & R; the rest of the modes are mono (where the output is only added to the input so L & R are the same).

mode[0] bypasses the resonator, but does not bypass the high pass filter at the input nor **xmix**, so you can still use these elements if desired.

Tips & Tricks

Changing **freq** or **tap** during play may produce clicks, pops, or “boing” sounds.

The lower you set **freq** the more resonances you will get, and they will be more tightly packed together.

When using parallel modes for high frequency resonance fill, set **hpf** at or somewhat above the highest formant frequency. Experiment with positive and negative **xmix** to get the proper blend. **freq**, **tap**, and **harm** can also strongly affect the blend.

mode[-2] high frequency fill can simultaneously provide pseudo stereo, though the resonance frequencies and blend region will be different for L & R.

Some useful starting points:

freq	774	tap	0
hpf	116	harm	-63
mode	-1	xmix	31
reso	-11		RESON

Male vocal

freq	539	tap	43
hpf	1244	harm	-63
mode	-2	xmix	-14
reso	6		RESON

Cello fill

freq	1333	tap	-32
hpf	1281	harm	63
mode	2	xmix	31
reso	59		RESON

Bell

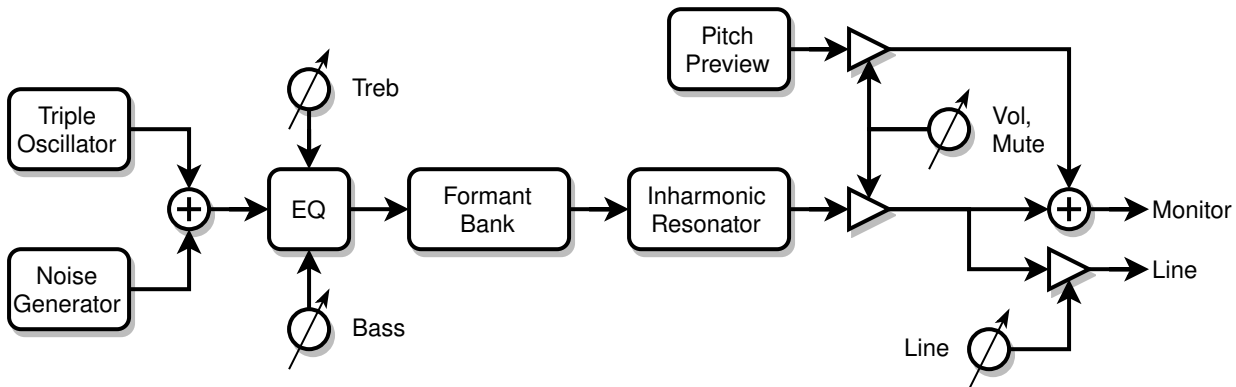
freq	46	tap	42
hpf	659	harm	35
mode	2	xmix	11
reso	0		RESON

Room ambience

To better understand exactly what the inharmonic resonator is doing to the signal and how the various parameters behave, choose a preset or create one that has a harmonic-rich single oscillator. Set **Pcal** to lower the pitch until it is about 1 or 2Hz with your arm in the rest position, so that you hear a tick, tick, tick impulse sound. Now you can experiment with the various parameters on the **RESON** screen and easily hear the resulting changes. You may note that having the inharmonic resonator in the signal path can make the impulses sound like the beat of a drum with adjustable tension, brightness, and sustain. As you play short musical phrases you may begin to correlate the resultant timbre with the settings.

Miscellaneous Controls

Description



These are a few miscellaneous *system* controls which weren't covered in the preceding sections.

Controls

Most controls are on the **LEVELS** page.

Vol	[0:63]	Line	[0:63]
Prev	[0:63]	Treb	[-31:31]
osc	[0:63]	Bass	[-31:31]
nois	[0:63]		LEVELS

LEVELS page controls.

Vol is a master audio output control over the volume of *everything* except MIDI. This is a digital control, so reducing it also reduces the resolution of the audio output. For your convenience this control also appears on the **D-Lev** page.

Line controls the volume of the LINE output. This is a digital control, so reducing it also reduces the resolution of the audio output.

Bass and **Treb** EQ the lows and highs. Use these to compensate for sonic deficiencies of your PA / amplifier / speaker setup, *not* to EQ any specific voice.

Mute (located on the **D-LEV** page) when set mutes all audio and MIDI output.

Erev (located on the **SYSTEM** page) reverses the rotational sense of all encoders. This is set in the "factory" and you'll never have reason touch it unless you replace the encoders. (There unfortunately doesn't seem to be an industry standard here).

Tips & Tricks

Vol and **Line** control the digital levels presented to the digital-to-analog converters right at the very end, so you can't use them to address filter overloading.

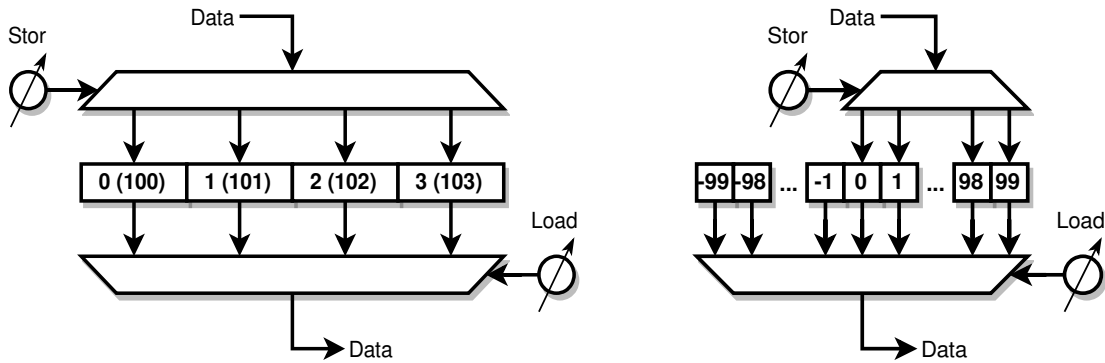
Please note that the D-Lev has both digital and analog level controls that at first glance appear to perform the same function. However the **Vol** and **Line** controls described above are digital level controls and the rightmost potentiometer knob below the tuner on the D-Lev's front panel is an analog line level control.

Even though they reduce the audio resolution, it's acceptable to use **Vol** and **Line** to control the volume levels. Indeed, the **osc**, **nois**, **prev**, **levl**, **xmix**, etc. controls do this too. The D-Lev internal resolution is 32 bits, with 24 bits presented to the digital-to-analog converters, so there are plenty of bits to go around. And any reduction in resolution is accompanied by a reduction in volume, so the quantization noise will remain relatively constant. In general, best practice dictates that you should keep the digital level as high as possible and use the analog volume knob on the front panel to control the theremin listening or blending volume.

If you want the D-Lev muted at power-up, set **Mute[1]** and save your *system* preset.

Mute can also be activated without touching the theremin by waving your volume hand over the infrared sensor located above and behind the volume loop. With traditional volume sensing (where the theremin is quiet with the hand close to the volume loop) this provides a graceful means of muting and unmuting the audio.

Preset System



System presets (left) – Synth voice presets (right).

Description

A preset is a group of control settings stored in a memory slot. The preset system manages this preset memory, and enables you to quickly and easily store and recall presets. So if you want to switch to a clarinet sound you can simply load its preset rather than twiddle all the controls until you finally get a clarinet sound. If you've created a violin sound you really like, you can save all the control settings at once as a preset. If someone else is using your D-Lev and prefers different pitch / volume field settings than you do, those can be stored in a preset too.

There are two types of presets, *system* and *voice*, which are managed by two separate preset systems. *System* presets store basic setup settings that can be considered global player preferences and once set they usually don't change very often, and *voice* presets store settings that tend to change with the chosen synthesizer sound. To tell them apart on the screen, *system* controls begin with a capital letter (title case) and *voice* controls are all lower case.

There are 4 *system* preset slots, and the *system* preset in slot 0 is loaded at power-up. There are 199 *voice* preset slots: 0 through 99 are read & write, and the negative preset slots -1 thru -99 are read-only. The *voice* preset in slot 0 is loaded at power-up.

All *system* and *voice* preset slots (including the negative voice slots) can be read, written, backed up, and restored via the USB serial port and the D-Lev preset librarian software (d-lib) running on a Windows or Linux PC.

Controls

All controls are on the **D-LEV** and **SYSTEM** pages.

Vol	[0:63]	Mute	[0:1]
Vcal	[0:255]	Pcal	[0:255]
Stor	[0:99]	oct	[-7:7]
Load	[-99:99]		D-LEV

D-LEV page controls.

Wait	[0:99]	50Hz	[0:1]
		P<>U	[0:1]
Stor	[0:3]	Erev	[0:1]
Load	[0:3]		SYSTEM

SYSTEM page controls.

Load, Stor - All *voice* presets are managed via the **Load** and **Stor** controls on the **D-LEV** page. All *system* presets are managed via the **Load** and **Stor** controls on the **SYSTEM** page.

Turning the **Load** control immediately loads the preset, and automatically sets the **Stor** control value to the **Load** value (if **Load** is set to a negative value then **Stor** is set to 0).

Pressing the **Stor** control “arms” the write process, blinks the mute LED, and displays **! WR!** instead of **Stor**. Pressing the **Stor** control a second time performs the actual pre-set write, which automatically sets the **Load** slot to match the **Stor** slot.

While the write process is armed, turning any control, or pressing any control other than **Stor**, will disarm and cancel the write process. So **Stor** follows **Load** when loading slots, and **Load** follows **Stor** when writing slots.

Tips & Tricks

Whenever you save your *system* preset, save it to at least one other additional slot for safe-keeping.

Periodically back up all of your *voice* and *system* presets using the librarian to prevent the loss of those you really care about.

The safest procedure if you wish to experiment and develop your own voices is to load a desired starting preset and immediately store it to either a blank slot or a used slot that you don't mind overwriting. Now you can safely modify the new preset without the risk of altering the original.

When developing a new voice or modifying an existing voice, consider using successive voice preset slots to store intermediate variations as you work. Not all changes are improvements, and the ability to easily go back makes the editing process less risky.

You can't write to a negative *voice* preset slot without the librarian, but you can always copy it to a non-negative slot and do whatever you want with it there.

Write your favorite voice to slot 0 so it's ready for you immediately after at power-up.

The value of **Load** following **Stor** when writing to a slot is confirmation that the write actually took place.

Use the librarian to trade preset files with other D-Lev owners!

If you backup all presets with the librarian when you first receive your D-Lev, you'll later be able to restore them all to the “factory” settings.

Pitch & Volume Field Setup

In order to get the most out of your D-Lev, the fields must be properly aligned. The D-Lev is uniquely capable of being calibrated to exhibit nearly ideal playing characteristics for almost any playing style, but with this capability there are necessarily a good number of parameters to adjust and knobs to twiddle. But once set, these parameters are saved to a system preset and rarely if ever need to be touched again.

The main goal here is to linearize the fields, which is accomplished by adjusting the field controls until you get the same pitch / volume change with an open & closed hand gesture anywhere in the playing zone. Aligning the fields doesn't require any special equipment, just the built-in LED tuner and your hand. And it isn't particularly difficult, though it can be a bit confusing the first few times through. Fortunately, once the fields are aligned, the settings can be saved to your system preset, with only occasional minor touch-up after that.


ACAL (auto-calibration) Process

ACAL is a process which simultaneously auto-calibrates both the pitch and volume fields. With your body in the playing position, initiate an *ACAL* by pressing and releasing the lower right encoder on any page. Use the **Wait** time to position your arms and hands (but keep your body in the playing position) into an easily repeatable orientation away from the antennas (I use an arms-crossed cadaver pose).

The **Wait** timer starts when you press the encoder, not when you release it. The internal *ACAL* event happens in an instant at the very end of the **Wait** timeout, so feel free to squirm around until then. The mute LED blinks at a 2Hz rate during the **Wait** interval, so you can count the blinks to know exactly when the time period will end.

This isn't a comprehensive auto-calibration as it only offsets the **Pcal** and **Vcal** values, but these are the primary environmental adjustments.

Pitch Field Setup Process



Pitch field setup open & closed hand gesture.

Steps for setting up the pitch field:

1. Turn off pitch correction: **levl[0]**.
2. Stand in your normal playing position and keep your body stationary.
3. Perform an *ACAL* like you normally do.

4. Fully open your hand and extend your arm, your fingers pointing at the center of the antenna, your fingertips very near but not touching the antenna.
5. Watch the tuner and try not to move your arm as you fully close your hand to make a fist, and then fully open your hand, doing this several times. It may help to open and close your hand quickly and repeatedly rather than slowly. The change in semitones between the two hand gestures is the *near-field sensitivity*. Adjust **sens** until the near-field sensitivity is what you desire (usually one octave).
6. Fully open your hand and retract your arm so your hand is back near your body, with fingers still pointing at the antenna.
7. Watch the tuner and perform the same close / open hand gesture. The change in semitones here is the *far-field sensitivity*. Adjust **Pcal** until this is the same as the near-field sensitivity, then save your system preset.
8. Because **Pcal** interacts a little with the other controls, you may have to repeat steps 5 through 7 until the near-field and far-field sensitivities are the same and what you desire.
9. Now we'll adjust **Ofs-**. Note your **sens** setting, then hold your hand in a fixed position in the *mid-field* (around 25cm or 10" away from the antenna) and note the octave number you see on the tuner. With your hand still in place, turn **sens** all the way up and adjust **Ofs-** until you get the same octave number. Set **sens** back to what it was originally and save your system preset.
10. Play some tunes with various voices and adjust **Ofs+** until the pitch field is comfortably positioned to where you like it. Save your system preset.
11. Do the open / closed hand gesture at various distances from the antenna. Some adjustment of **Lin** may improve any near-field and mid-field inconsistencies, but **Pcal** should be your go-to control for balancing near-field and far-field sensitivity. **Lin** strongly interacts with the other controls, so this is the most tedious adjustment. If it all seems too confusing just set **Lin[0]** for now. Don't forget to save your system preset!

Volume Field Setup Process

The procedure for setting up the volume field is mostly the same as setting up the pitch field, though the open / closed hand gesture must be performed vertically, which can be a bit awkward, and the steps are a little different for **Vrev[0]** and **Vrev[1]**.



Volume field setup open and closed hand gesture.

Steps for setting up the volume field *with Vrev[0]*:

1. Disable knee, velocity, and envelope: **knee[0]**, **velo[0]**, **rise[0]**, **fall[0]**, **dloc[0]**.


2. Linearity of the volume field is much less critical than that of the pitch field. So it's fine to simply set the volume field **Lin** to whatever the pitch field **Lin** setting is, or to just set **Lin[0]**. You can experiment around with **Lin** if you like, but it certainly isn't necessary.
3. Stand in your normal playing position and keep your body stationary.
4. Perform an ACAL like you normally do.
5. Fully open your hand and extend your arm, your fingers pointing down at the center of the antenna, your fingertips very near but not touching the antenna. Adjust **Ofs+** until all of the volume LEDs just go out.
6. Watch the tuner and try not to move your arm as you slowly fully close and open your hand. Adjust **sens** until this causes only the bottom LED to turn fully on and off.
7. Fully open your hand and raise it up until the bottom 3 LEDs are lit, with fingers still pointing down at the antenna.
8. Watch the tuner and perform the same close / open hand gesture. Adjust **Vcal** until this causes all four LEDs to just fully light with your closed hand.
9. Because **Vcal** interacts a little with the other controls, you may have to repeat steps 5 through 8 until your open / closed hand gesture causes exactly a one full LED change anywhere in the volume field.
10. Now we'll adjust **Ofs-**. Note your **sens** setting, then hold your hand in a fixed position from the antenna so that the two bottom LEDs are fully on and the top two LEDs are fully off. Without moving your hand, turn **sens** all the way up and adjust **Ofs-** until you get the same two LEDs on and two LEDs off. Set **sens** back to what it was originally and save your system preset.
11. Turn any knee back on (reload your voice preset) and play some tunes with various voices. Adjust **Ofs+** until the volume field is comfortably positioned where you like it. Don't forget to save your system preset!

Differences for Vrev[1]:

If you prefer to play with reversed volume field sense **Vrev[1]**, set **Vrev[0]** and follow the above steps. Then switch to **Vrev[1]** and adjust **Ofs+** to position the volume field to your liking. Alternatively, it's easy to reinterpret the **Vrev[0]** steps for **Vrev[1]**.

Volume Knee Setup

Believe it or not, you really don't want a perfectly linear volume field. It sounds desirable, but you would be flapping your arm like a bird to achieve even mild dynamics. No, what you want is a field with higher sensitivity at lower volumes, a field that quickly ramps up to a certain loudness, with a more expressive region above that. You want to use small, slow, natural hand movements, yet still retain good control where it really matters.



The influence of **knee** and **kloc** on the volume field response; **Vrev[0]** left, **Vrev[1]** right.

So why go to all the trouble to linearize the volume field? To make it non-linear in a very controlled and predictable way – which is precisely what **knee** does. **kloc** sets the **knee** location in the volume field. For volume levels louder than **kloc** there is no change to the response, but for volume levels below **kloc** a higher sensitivity kicks in, with the increase set by **knee**. The volume field in both sections of the field is still linear, but hand movements at lower volumes will cause larger volume changes.

Moderate **knee** settings can also minimize the influence of your body near the volume antenna. Without the **knee** it may be difficult to reach the extremes of full volume or silence due to the nearby bulk capacitance of your body.

Analog Theremins usually have a volume response that is quite non-linear near the antenna and “baked in” to the design. The **kloc** and **knee** controls set this to wherever and however strongly you like, allowing you to keep your hand farther away from the volume antenna at all times, making accidental contact during play much less likely.

Troubleshooting

Software Errors

Like anything containing a processor, the D-Lev is susceptible to software errors, which can be caused by power drooping, an electrostatic discharge event, rogue preset data, certain commands from the librarian, the phase of the moon, etc. These errors don't happen often, but when they do:

**THE FIX FOR MOST ERROR CONDITIONS IS:
POWER CYCLE THE D-LEV**

Some typical symptoms of software trouble:

- No audio or strange audio issues.
- Unresponsive pitch / volume fields.
- Unresponsive encoders / LCD / tuner.
- Weird characters / odd data on the LCD.
- **ER!** or **SW!** error indicators on the main **D-LEV** screen:

Vol	64	Mute	0	Vol	64	Mute	0
Ucal	4	Pcal	32	Ucal	4	Pcal	32
Stor	0	oct	-1	Stor	0	oct	-1
Load	0	ER! D_LEV		Load	0	SW! D_LEV	

*Error indicators on the main **D-LEV** screen.*

The **ER!** error indicator means the software has gone into the weeds. The processor has either experienced a stack fault, or an interrupt routine has exceeded its real-time allotment. Time to power-cycle!

The **SW!** error indicator means the EEPROM software image failed the error check, and the default FPGA software image is being used instead. If a power cycle doesn't clear this error then you will have to upload new software via the librarian. You probably shouldn't operate the D-Lev in this state for too long as it is likely running an older software version (in particular avoid writing presets via **Stor**).

Having a default software image in the FPGA means it's impossible to permanently "brick" the D-Lev by uploading bad software. But there are rare situations (caused by boot code / software mismatches) where the D-Lev just won't boot. In this scenario you can force the D-Lev to bypass the EEPROM image testing and use the default FPGA software image instead.

1. Press and hold down encoder 7 (located at lower right).
2. Do a power cycle.
3. Release encoder 7.

After this you should see **SW!** on the main **D-LEV** screen. Now use the librarian to upload new software.

Pilot Errors

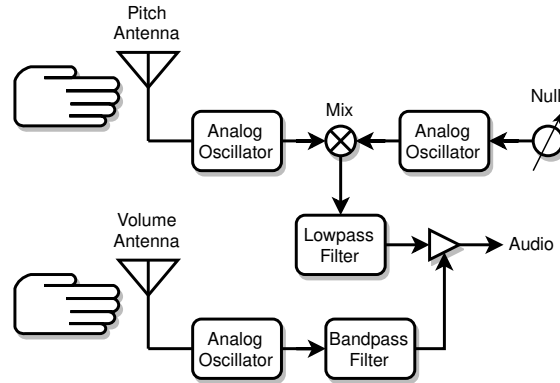
Certain control settings can produce what might initially seem like incorrect behavior, but they're actually doing what they're supposed to do. Some rascals to watch out for:

Symptom	Cause / Solution
Incessant screeching sound.	Filter oscillation: lower any high reso settings & lower osc / nois to not overload filter input; reload the preset.
Sound won't stop.	Lower high fall / damp on VOLUME page; lower any high reso settings.
Tuner dark.	Brit[0] on DISPLAY page.
LCD dark.	LCD[0] on DISPLAY page.
Reversed volume.	Vrev[1] on the V_FIELD page.
Sticky pitch / volume fields.	Bad ground; Dith on V_FIELD / P_FIELD pages too low;
Bobbling pitch / volume fields.	Bad ground; Dith on V_FIELD / P_FIELD pages too high; 50Hz on SYSTEM page.
Tuner acting sticky.	Qant, Post[1] on DISPLAY page.
Out of tune.	Cent on DISPLAY page.
ACAL too short / long.	Wait on SYSTEM page.
Swapped fields.	P<>V[1] on SYSTEM page.
Reversed encoder rotation.	Erev on the SYSTEM page.

FAQ

Is the D-Lev a TRADITIONAL Theremin?

No. A traditional Theremin heterodynes (non-linearly mixes) the outputs of two radio frequency analog oscillators. One oscillator operates at a fixed frequency, and the frequency of the other oscillator varies with the capacitance of the pitch hand. Mixing generates sum and difference frequencies at F_1+F_2 and $|F_1-F_2|$, and low-pass filtering removes the upper sum frequency, leaving only the lower difference frequency, which is within the range of human audibility. Even though the frequency of the variable oscillator only changes a few percent over the full range of pitch hand movement, the difference frequency varies over many octaves.



How a traditional analog Theremin works.

Traditional volume control can take many forms. Similar to the pitch side, the frequency of the volume oscillator varies with the capacitance of the volume hand. One common technique to produce a large change in signal over a relatively short range of hand movement is to sharply filter the oscillator output, rectify and filter it (not shown), and use the result to control the audio volume.

Note that the heterodyning process links the pitch and timbre of the resulting audio, which ties the Theremin designer's hands when it comes to altering the basic (non)linearity, sensitivity, timbre, octave, etc. Pitch field sensitivity can't be altered without negatively impacting overall linearity. Coupling the fixed and variable oscillators can enrich low end timbre, but this also negatively impacts far-field linearity.

It's pretty amazing that you can make a traditional Theremin which works pretty well with only a handful of parts, but the engineering constraints present at the lowest level will remain almost no matter how fancy any additional circuitry gets. The near and far fields will always be somewhat cramped, and the mid field will always have the same sensitivity.

With the various pitch and volume field options, what should I be considering?

It's easy to produce D-Lev fields that simply can't exist on any analog Theremin. This creates a rather dangerous situation for the budding Thereminist because, if you put in a lot of practice time with non-traditional fields, you might discover that you can't play any other type of Theremin that you encounter. Then again, maybe traditional fields for one reason or another turn you off.

If there is a particular player whose technique you find attractive and want to emulate, then you should set up the fields to be like their instrument, which probably means like a traditional analog Theremin.

Can you give me an example of traditional analog field settings?

A traditional one octave open / closed pitch hand, and farther is louder volume hand setup:

Vcal	4	Vrev	0	Vcal	32	Cent	0
Lin	-2	Dith	0	Lin	-2	Dith	0
Ofs-	30	Ofs+	63	Ofs-	32	Ofs+	100
Sens	75	U_FIELD		Sens	65	P_FIELD	

Traditional field settings example (FIXME!).

These should only be taken as a rough guide. Your own settings will likely vary somewhat, and ideally would be the result of you personally performing the field setup procedures given in this manual.

How do you set up the fields on your own D-Lev?

I use a reversed volume field and 1/3 octave (4 semitones) open / closed hand pitch field sensitivity:

Vcal	4	Vrev	1	Vcal	32	Cent	0
Lin	-2	Dith	0	Lin	-2	Dith	0
Ofs-	30	Ofs+	63	Ofs-	32	Ofs+	100
Sens	75	U_FIELD		Sens	65	P_FIELD	

My current field settings.

I don't want to give anyone the impression that non-traditional field settings are a cure-all, but I *would* encourage experimentation here.

Acknowledgments

Huge thanks to Roger Hess for throwing his entire life into this labor of love for an extended period of time, taking on the task of alpha / beta tester, spearheading the PCB initiative, exploring exotic cabinetry design, making the first demo units, editing this manual, and being just an all around incredibly good friend and collaborator. Many features in the D-Lev owe their existence / refinement to his keen musical and engineering insights and observations.

Special thanks to Jason Barile for hosting the fantastic Theremin World web site (www.thereminworld.com) where the D-Lev was born. Many super nice and interesting folks there helped guide the journey.

Everyone who took the time and effort to openly document and post their research into Theremin technology and sound synthesis. This very long list includes:

- Norio Tomisawa's patent US4249447 *Tone Production Method for an Electronic Musical Instrument* (D-Lev audio oscillators).
- Jon Dattorro's *Effect Design* papers (D-Lev noise source).
- Hal Chamberlin's book *Musical Applications for Microprocessors* (D-Lev filters).
- *Musical Applications Of Nested Comb Filters for Inharmonic Resonator Effects* by Jae hyun Ahn and Richard Dudas (D-Lev inharmonic resonator).
- *DAFX - Digital Audio Effects* edited by Udo Zölzer (D-Lev bass & treble EQ).
- Ilya Mamontov's characterization of hand & antenna capacitance.
- The *musicdsp.org* source code archive (via web.archive.org).
- Chet Singer's *Physical Modeling on the Nord Modular G2*.
- Julius O. Smith's many papers and web references.

The *LibreOffice* development team. This manual was written in *Writer*, and many critical engineering tasks were made possible by *Calc* – you've come a long way, *OpenOffice!*

Most of the illustrations in this manual were generated by the excellent *diagrams.net* (aka *draw.io*) program.

The font *hd44789-regular.ttf* from fontsov.com was used for the LCD illustrations.

The ASL font *GallaudetRegular.ttf* by David Rakowski was used for the hand illustrations, specifically the characters 'S' and '5', which were edited slightly in *InkScape*.

Finally, a very special thanks to Fred Mundel, who was always up for a new circuit design. Fred inspired everyone with his passion, infectious curiosity, and sincere kindness, and he left us way too soon.