Theta
Tontechnik

Bass Trickkiste

User Manual

Version 1.2.0

Build fd2ed71 — 2026-05-20

Contents

Chapter 1

Overview

Bass Trickkiste is a bass decorrelation plugin designed for multi-channel and multi-speaker audio systems. It splits the incoming signal at a configurable crossover frequency, applies one of seven decorrelation algorithms to the low band, and recombines the result with the phase-matched high band—all while preserving flat frequency response across the crossover. Each of the 16 output channels can be independently configured as Full, Sub, LFE, PassThrough, or Off, making it suitable for everything from simple stereo setups to complex immersive speaker arrays.

1.1 Key Features

- **Seven decorrelation modes**—Allpass, Delay, Hilbert, Schroeder, Complementary AP, DiSP (dynamic), and DiSP (static)—each with distinct character and CPU profile.
- **16-channel output routing**—every channel independently assignable to Full, Sub, LFE, PassThrough, or Off.
- **Per-channel bass decorrelation**—Full outputs decorrelate each channel's own low band, preserving the original stereo bass image.
- **True-peak limiter**—configurable threshold from -12 to 0 dBTP, with brickwall limiting in real-time mode and lookahead limiting in buffered mode.
- **Output buffer**—real-time mode or configurable buffer (0.5 ms to 10 ms) with a background DSP thread for reduced CPU spikes on the audio thread.
- **Standalone application**—file import/export, offline processing, waveform display with transport controls.
- **48 dB/oct Linkwitz–Riley crossover**—phase-matched high band ensures flat reconstruction with no audible dip at the crossover frequency.

1.2 Main Parameters

The four parameters you will use most often:

Crossover Frequency Sets the split point between the low band (which gets decorrelated) and the high band (which passes through with phase matching). Typical values range from 80 Hz

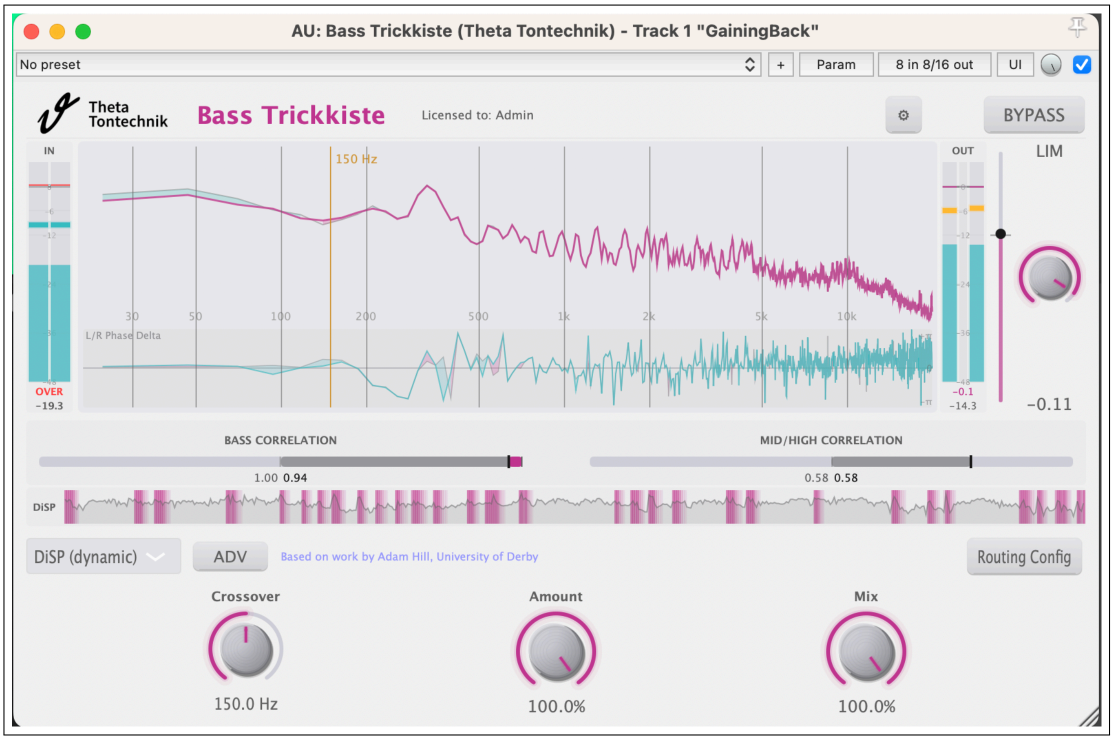

Figure 1.1: The Bass Trickkiste plugin window.

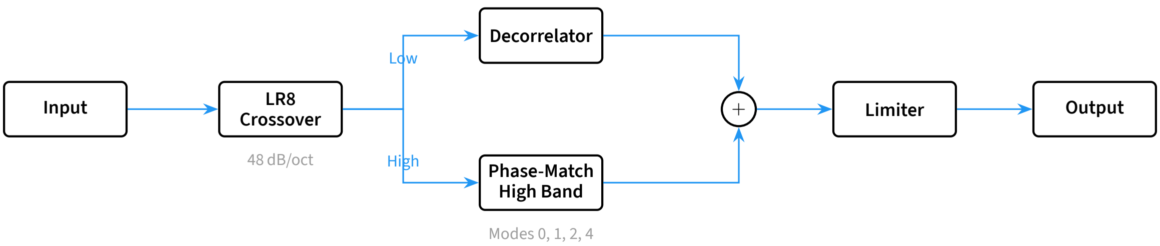

Figure 1.2: High-level signal flow. The input is split by a Linkwitz–Riley crossover into low and high bands. The low band is decorrelated; the high band is phase-matched to preserve flat crossover reconstruction.

to 200 Hz depending on your subwoofer configuration and room.

Mode Selects one of the seven decorrelation algorithms. Each mode has a different sonic character, latency, and CPU cost. See Chapter 3 for a detailed comparison.

Amount Controls how much decorrelation is applied to the low band. At 0 % the bass is untouched; at 100 % it is fully decorrelated. Intermediate values blend between the dry and decorrelated bass signals.

Output Level A final gain trim applied after decorrelation and before the limiter. Range -12 to $+6$ dB; double-click the slider to reset to 0 dB. Use this to compensate for any perceived level changes introduced by the processing.

Tip

Start with **Amount** at 100 %, listen, and then back it off until you find the balance between bass tightness and punch retention for your material. Bass Solo (see section 1.3) lets you audition just the low-band processing in isolation while you tune.

1.3 Bass Solo

The **Bass Solo** button (top-level toggle, next to the bypass button) mutes the high-band recombine step so only the decorrelated low band reaches the output. Use it to audition exactly what the decorrelator is doing without the full-range signal masking it: switch Bass Solo on while you adjust the mode, **Crossover Frequency**, and **Amount**, then switch it off to hear the result in context.

Bass Solo affects every output channel that would normally carry a high band — **Full** outputs become bass-only, while **Sub**, **LFE**, **PassThrough**, and **Off** outputs are unchanged. It has no effect on the limiter, output gain, or routing, so the output level you hear in solo matches what you would get summed with the high band (just without the high content). Bass Solo is not stored as a project preset parameter that you should leave on for printing — it is purely a monitoring tool.

Chapter 2

Signal Flow

This chapter walks through the processing chain from input to output. Understanding the signal flow will help you make informed decisions about mode selection, crossover frequency, and output routing.

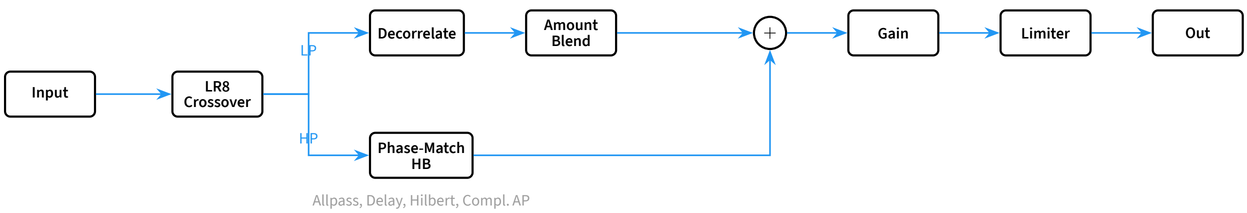

Figure 2.1: Detailed signal flow for Full-type outputs. The low band is decorrelated and blended via **Amount**. The high band passes through the same phase filter to maintain crossover flatness.

2.1 Crossover Filter

The input signal is split into a low band and a high band by a Linkwitz–Riley 8th-order crossover at 48 dB/oct. Everything below the **Crossover Frequency** is routed to the decorrelator; everything above is routed to the high-band path. The steep slope ensures a clean separation: the decorrelator sees very little high-frequency content, and the high band carries very little bass.

Tip

For most subwoofer-based systems, a crossover between 80 Hz and 120 Hz works well. If you are decorrelating a full-range speaker pair without a sub, you may want to go higher—up to 200 Hz—to cover the fundamental range where bass buildup is most problematic.

2.2 High-Band Phase Tracking

Every decorrelation mode changes the phase of the low band. To prevent a notch at the crossover frequency when the two bands are recombined, Bass Trickkiste applies phase tracking to the high band where possible, so the high band's phase follows the low band's phase through the crossover region. For modes where the time-varying or diffuse nature of the algorithm prevents a clean phase-tracking step, the 48 dB/oct slope keeps the overlap region narrow enough that any residual mismatch is not audible.

2.3 Per-Channel Bass

Bass Trickkiste offers two distinct approaches to bass decorrelation, depending on the output type:

Full outputs Each channel's own low band is decorrelated independently. If the input has a stereo bass image (for example, a bass guitar panned slightly left), that spatial information is preserved in the decorrelated output. This is the default and recommended mode for full-range speakers.

Sub outputs The low bands of all input channels are summed to mono before decorrelation. This is standard practice for subwoofer feeds, where mono bass is desirable for consistent low-frequency coverage regardless of listener position.

> **Note**
>
> The per-channel vs. mono-sum distinction only affects the low band. The high band is always processed per-channel, regardless of output type.

2.4 Output Assembly

Each of the 16 output channels is assigned one of five output types via the Routing Configuration panel (see Chapter 4). Here is what each type produces:

Full Decorrelated low band (per-channel) recombined with the phase-matched high band. This is a full-range signal suitable for main speakers.

Sub Decorrelated low band only (mono sum of all inputs). No high-band content. Intended for subwoofer outputs.

LFE An independently decorrelated version of the raw input signal (no crossover split). This is designed for dedicated LFE channels that carry their own content, separate from the bass-managed signal.

PassThrough The original dry input signal, delay-compensated to align with the processed channels. Useful for reference channels or surrounds that should not be decorrelated but need to stay time-aligned with the rest of the system.

Off Silence. The channel produces no output. Use this for unused channels to avoid sending noise or stale signal to disconnected outputs.

After output assembly, the **Output Gain** is applied, followed by the true-peak limiter (see section 2.7).

2.5 Amount

The **Amount** parameter blends between the dry low band and the decorrelated low band, before the low and high bands are recombined. At 0 %, the bass passes through untouched. At 100 %, the bass is fully decorrelated. Values in between produce a partial decorrelation effect.

This is the primary creative control. Reducing **Amount** from 100 % lets you dial in just enough decorrelation to reduce bass buildup without completely altering the bass character. For A/B comparison and bypass-style checks, use the **Bypass** button rather than dialling **Amount** to zero — bypass keeps the latency compensation correct so the comparison is sample-aligned.

2.6 Output Buffer

By default, Bass Trickkiste runs in real-time (RT) mode: all DSP processing happens directly on the audio thread within the host's **processBlock** callback. This gives the lowest possible latency but means the audio thread must complete all processing within one buffer period.

For sessions with high channel counts or CPU-intensive modes (especially DiSP), you can enable the output buffer. This moves all DSP processing to a dedicated background thread, while the audio thread only copies samples in and out of lock-free ring buffers. Available buffer sizes range from 0.5 ms to 10 ms.

The trade-off is additional latency equal to the buffer size. The host is informed of the total latency (DSP algorithm latency plus buffer latency) so that automatic delay compensation works correctly.

> **Note**
>
> The output buffer also affects the limiter: in RT mode, the limiter operates as a brickwall clipper with zero lookahead. In buffered mode, the limiter uses the buffer latency as lookahead time, producing smoother and more transparent limiting.

2.7 True-Peak Limiter

The very last stage in the signal chain is a true-peak limiter. It is on by default and sits there even when no gain is being applied — it is not a creative loudness tool, it exists to catch inter-sample peaks (ISPs) that the decorrelation processing itself can introduce.

Every decorrelation mode in this plugin works by rearranging the phase of the low band. Once the phase relationship between frequency components changes, the resulting time-domain samples can crest between two samples at a level higher than either sample. A signal that was -1 dBFS at the input can therefore reach 0 dBTP at the output even though no sample value exceeds -1 dBFS. This is invisible to a sample-domain meter but causes audible clipping in any downstream DA converter or sample-rate converter that reconstructs the continuous-time signal.

The limiter measures the true (inter-sample) peak and applies just enough gain reduction to keep it below the configured **Limiter** threshold (default -0.1 dBTP). Two flavours are available depending on the **Output Buffer** setting:

Real-time (RT) mode Zero-lookahead brickwall limiting — instant attack, fast release. Adds no latency. Slightly less transparent than the lookahead version on rare hard transients but introduces no audible artefacts at the default threshold.

Buffered mode Uses the output-buffer latency as a lookahead window, smoothing the gain reduction over the lookahead period. This produces a more transparent limiter response and is the recommended setting if you are using a buffered output mode anyway.

You can bypass the limiter via the **Limiter Bypass** parameter, but doing so means the plugin can produce signals that exceed 0 dBTP. Leave it on unless you have a specific reason to disable it.

Tip

The default -0.1 dBTP threshold is conservative — enough margin to safely survive any downstream resampling or DA conversion without ever clipping. Lower it further (e.g. -1 dBTP) if you are feeding into a chain that includes lossy encoding, where bit-perfect headroom is more valuable than absolute loudness.

Chapter 3

Decorrelation Modes

Decorrelation reduces the similarity between channels so that each speaker reproduces a slightly different version of the bass signal. This spreads low-frequency energy across the listening area instead of concentrating it in a single pressure zone. Below roughly 150 Hz, wavelengths are so long that all speakers in a typical room reproduce nearly identical waveforms—decorrelation breaks that uniformity, giving the bass more definition, better seat-to-seat consistency, and a wider perceived sound field.

Recommended workflow — read this before tuning

If you have a multitrack mix with the kick drum on a separate channel from the bass instrument: process the bass instrument only and leave the kick drum untouched. Then turn the transient detector off (applies to DiSP modes only).

The transient detector exists to preserve punch on material where bass and kick share a channel. If you can give the plugin a kick-free bass signal in the first place, your kick keeps its full impact because it never goes through the bass-band processing at all, and the bass-only signal can be decorrelated without the detector intervening. This is the configuration the plugin sounds best in end-to-end.

For full-mix or stem material where the kick and bass share a channel, leave transient detection on (the default) — it is the right tool for that case, just not as good as separating the two sources upstream.

Bass Trickkiste offers seven decorrelation modes, numbered 0 through 6. The table below summarises their character and use case.

3.1 Mode 0 — Allpass

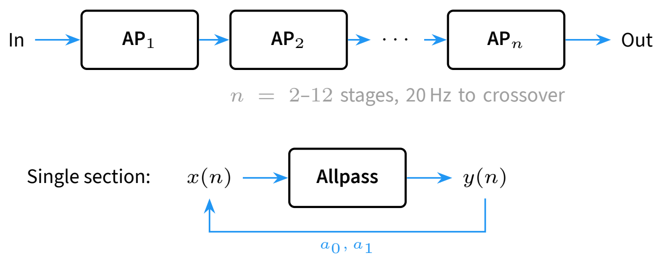

Mode	Character	Best for	CPU
0 – Allpass	Transparent, frequency-based	Default choice; subtle widening	Low
1 – Delay	Time-based, subtle	Reducing seat-to-seat cancellation	Low
2 – Hilbert	Warm, phase-based	Sustained bass material	Low
3 – Schroeder	Diffuse, reverberant, aggressive	Enveloping, room-like character	Low
4 – Complementary AP	Smooth, wide	Between Allpass and Schroeder	Low
5 – DiSP (Dynamic)	Most natural; time-varying	Reflective rooms, cinemas, live venues	High
6 – DiSP (Static)	Most natural; steady	Dry rooms, near-field, CPU-tight setups	Medium

Table 3.1: Summary of decorrelation modes.

Spread	0–100 %. Controls how far apart the L and R allpass centre frequencies are placed. Higher values produce stronger decorrelation.
Stages	2–12. Number of cascaded second-order allpass sections per channel. More stages give denser phase variation across the bass range.

Allpass mode passes each channel through a cascade of second-order allpass filter sections. Every section shifts phase without changing the magnitude spectrum, so the frequency response stays perfectly flat. The left and right channels receive different centre frequencies—spread apart according to the **Spread** parameter and speaker position—which creates a frequency-dependent phase difference between them.

Because this is a full-band mode, the allpass chain is also applied to the high band for phase matching. You will not hear any timbral change; only the spatial impression of the bass widens.

Tip

For the most transparent decorrelation, start with Allpass mode at moderate spread (50–70 %).

3.2 Mode 1 — Delay

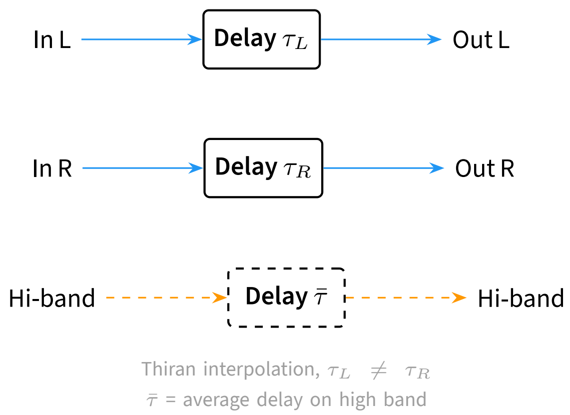

Delay Time	0.5–5 ms. Maximum delay applied between the two channels. L and R receive complementary delays based on speaker position.

Delay mode applies Thiran-interpolated micro-delays to each channel. Thiran allpass interpolation gives sub-sample accuracy, avoiding the comb-filter artefacts that naive integer-sample delays would introduce at bass frequencies. The left and right channels receive different delay amounts derived from their speaker positions, while the average delay is applied to the high band to maintain time alignment at the crossover.

The effect is subtle and time-based: you hear the bass arriving at slightly different moments from each speaker, which reduces the standing-wave pattern in the room. This mode works particularly well in rooms where bass cancellation at specific seats is the primary problem.

3.3 Mode 2 — Hilbert

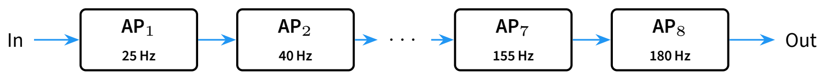

8 first-order allpass sections, 25–180 Hz

Amount	Controls the frequency spread between the L and R allpass chains. Higher values push the centre frequencies further apart.

Hilbert mode uses eight cascaded first-order allpass sections with centre frequencies distributed across the bass range (25–180 Hz). Each section places its $-90°$ phase point at a specific frequency; the cascade accumulates phase differences that grow with frequency across the entire bass band.

The left and right channels receive different centre-frequency sets, spread apart by the **Amount** parameter. A matching allpass chain is applied to the high band to preserve phase coherence at the crossover.

The result is a warm, phase-based decorrelation that sounds natural on sustained bass notes. Because the phase difference varies smoothly with frequency, the effect avoids the “phasey” character that abrupt phase jumps can produce.

3.4 Mode 3 — Schroeder

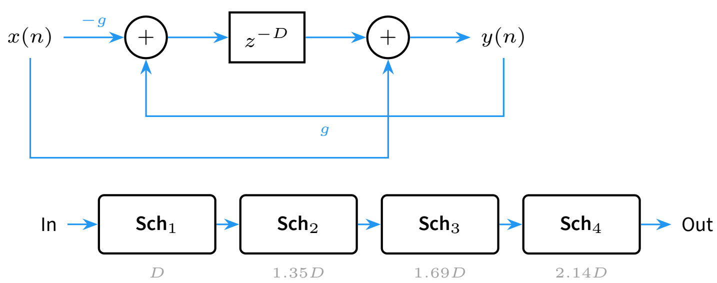

Delay	1–50 ms. Base delay time for the diffuser sections. L and R channels receive different prime-ratio multiples of this value.
Amount	Controls the L/R delay spread.

Schroeder mode cascades four Schroeder allpass diffuser sections per channel. Each section uses the classic structure $y(n) = -gx(n) + x(n-D) + g\,y(n-D)$ with a fixed feedback gain of $g = 0.7$. The four sections use prime-ratio delay spacing—a different set of ratios for L and R—to avoid resonant build-up at any single frequency.

This mode produces a diffuse, reverberant character that smears the bass energy over time. It is the most aggressive decorrelation mode and works well when you want the bass to feel enveloping rather than precise. High-band phase matching is not applied because the comb-filter delay lines (441–942 samples at typical settings) would require an impractically long settling time.

Note

Because Schroeder mode does not phase-match the high band, you may notice a slight timbral shift at the crossover frequency. Keep **Delay** moderate (5–15 ms) to minimise this effect.

3.5 Mode 4 — Complementary AP

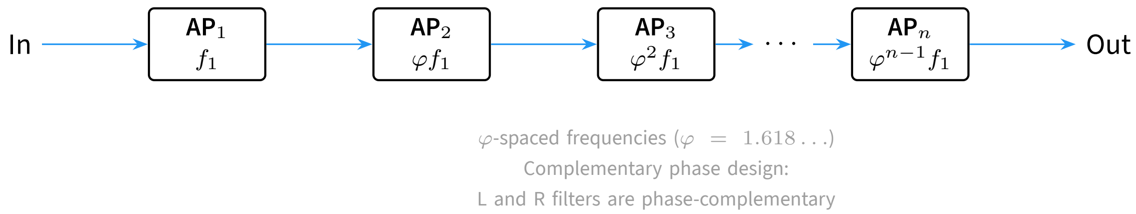

Spread	0–100 %. Controls how far apart the L and R centre frequencies are placed.
Stages	2–12. Number of cascaded second-order allpass sections.

Complementary AP mode is structurally similar to Allpass mode (both use cascaded second-order allpass sections), but differs in two important ways. First, the centre frequencies are spaced using the golden ratio ($\varphi \approx 1.618$) rather than linear spacing. This irrational spacing prevents the phase-shift peaks of different stages from aligning at harmonic frequencies, producing a smoother overall phase spread.

Second, the L and R coefficients are designed as a complementary pair: their combined phase response covers the bass band more uniformly than two independently designed chains.

Like Allpass mode, this is a full-band mode with high-band phase matching. The character is smooth and wide—a good choice when Allpass mode feels too subtle but Schroeder mode is too aggressive.

3.6 Mode 5 — DiSP (Dynamic)

Parameters	See chapter 5 for the DiSP Advanced Settings panel.

Dynamic DiSP[1] is the most sophisticated decorrelation mode and the most natural-sounding of the seven. The decorrelation continuously varies over time, which avoids the static colouration that any single fixed-character decorrelator would impose on a long passage of bass.

[1] The DiSP modes are based on Moore & Hill, "Decorrelation of multi-loudspeaker reproduced audio using Temporally Diffused Impulses," Journal of the Audio Engineering Society, 2018.

A transient detector keeps percussive elements (kicks, bass attacks) sharp by routing them around the bass-band processing during the attack phase, then smoothly returning to fully decorrelated steady-state once the transient passes. This makes Dynamic DiSP suitable for full-mix material where bass and kick share a channel, as well as for clean bass-only stems.

> **Note**
>
> DiSP modes prepare their filter set asynchronously on a background thread when the mode is engaged or parameters are changed. The decorrelation effect fades in once the set is ready — typically a fraction of a second to a few seconds depending on the **Duration** and **Interp Factor** settings.

The trade-off is higher CPU usage and the brief warm-up period. For material where bass quality is paramount — cinema, live events, immersive installations — this is the recommended mode.

3.7 Mode 6 — DiSP (Static)

Parameters	Shares the DiSP Advanced Settings with Dynamic mode (chapter 5); the cycling-related parameters do not apply.

Static DiSP uses the same flat-magnitude, random-phase decorrelation as Dynamic DiSP, but the decorrelation is fixed in time rather than continuously varying. The transient detector is still active — attacks pass through cleanly while the sustained portion is decorrelated.

Static DiSP uses noticeably less CPU than Dynamic mode and has no warm-up. The decorrelation quality is still excellent and the bass is uncoloured, but it lacks the continuously varying character that makes Dynamic DiSP particularly effective at avoiding room-mode excitation. Choose Static DiSP when CPU budget is tight or when the bass content is relatively steady (pads, drones, sustained tones) rather than rhythmically complex.

3.8 Dynamic vs. Static DiSP — choosing the right one

Both DiSP variants produce flat-magnitude, randomised-phase decorrelation. The difference is whether the decorrelation varies over time (Dynamic) or stays fixed (Static).

What each mode is solving acoustically

The listener hears the direct wavefront from each speaker plus a sequence of reflections off walls, ceiling, floor and other surfaces. Static and Dynamic differ in which of these the decorrelation acts on.

Static Decorrelates the direct sound. The direct wavefronts arriving at the listener from different speakers no longer carry the same phase pattern and do not reinforce coherently. Reflections, however, are delayed copies of the already-decorrelated signal — because the decorrelation is fixed in time, each reflected arrival was emitted with the same decorrelation as the direct sound and recombines with the direct using the same phase relationship. The reflected field therefore retains the original inter-speaker correlation structure even though the direct field does not.

Dynamic Decorrelates both direct and reflected sound. Because the decorrelation continuously changes over time, a reflection arriving at the listener at time T was emitted earlier than the direct sound that arrives at T — so it was emitted under a different decorrelation. The result is that reflections are decorrelated from the direct wavefront (not just between speakers), giving a more uniform bass field across the room and over time. This is what makes Dynamic particularly valuable on long-held bass in reflective rooms — the reverberant build-up that Static cannot address is what Dynamic was designed for.

Pick Dynamic when

- The room contributes a substantial reflected field — large untreated rooms, live venues, cinemas, immersive installations. Decorrelating the reflections is where Dynamic earns its keep over Static.
- The bass material is long-form (synth pads, organ pedals, drones, sustained bass notes), where there is time for the reverberant field to build up and for the time-varying decorrelation to act on it.
- You can spend the extra CPU.

Pick Static when

- The room is acoustically dry, the system is near-field enough that the reflected field is small, or there is no room at all — studio mix positions, near-field monitoring, well-treated rooms, open-air festivals and outdoor events. With little reflected energy, the direct-sound decorrelation that Static provides is most of what is acoustically available to gain.
- The bass is rhythmically dense and short notes dominate, so the reverberant tail never has time to build up between notes and Dynamic's reflected-sound benefit is not in play.
- CPU is tight (large channel counts, low-power host, many plugin instances).
- You want a time-invariant character within a session, so the bass colouration does not drift over the course of a long passage and A/B comparisons of a parameter toggle land on the same decorrelation on both sides of the comparison.

Latency

Both Static and Dynamic DiSP report the same latency to the host. Switching between them does not change the reported plugin latency.

Tip

A reasonable starting point: pick **Dynamic** if you are tuning a system that plays into a real room (PA, cinema, install, untreated mix room) and bass content sustains long enough for reflections to matter. Pick **Static** for near-field monitoring, dry rooms, open-air festivals / outdoor events, very rhythmic bass, or when CPU is tight.

Note

In every mode, multi-speaker installations need each speaker / cluster to receive its own plugin output channel so that each one carries a different decorrelated signal — see section 4.3 in the Routing chapter.

Chapter 4

Routing Configuration

Bass Trickkiste outputs 16 independent channels, each of which can be configured to one of six output types. This flexibility allows a single plugin instance to drive main speakers, subwoofers, LFE channels, and monitoring outputs simultaneously.

4.1 Output Types

Type	Source	Use Case
Full	Decorrelated bass + original high band	Main speakers
Full (Mono Bass)	Decorrelated mono-sum bass (shared decorrelation filter) + per-channel high band	Full-range output fed to an L+R-summing amp rack
Sub	Decorrelated bass only (mono sum)	Subwoofers
LFE	Same-index input decorrelated full-band	LFE channel in surround
PassThrough	Dry input with delay matching	Reference / monitoring
Off	Silence	Unused channels

Table 4.1: The six output types available for each channel.

Full The primary output type. The low band is decorrelated using the selected algorithm, then recombined with the phase-matched high band. Each Full channel decorrelates its own input channel's bass independently.

Full (Mono Bass) Like **Full**, but the low band is the mono sum of all input low bands and a single shared decorrelation filter is applied across every **Full (Mono Bass)** output. All channels of this type therefore carry identical decorrelated bass content; only the per-channel high band differs. Use when downstream amplifiers internally sum the stereo input's L+R for a mono sub feed — see section 4.4.

Sub Outputs only the decorrelated low band as a mono signal (sum of all input low bands). No high-band content is included. Ideal for dedicated subwoofer feeds where you want the decorrelated bass without the full-range signal. Each Sub-configured output channel carries its own independently decorrelated version of the same mono bass — see the per-subwoofer-routing note below.

LFE Takes the input on the same channel index as the output (e.g. the LFE bus in a 5.1/7.1/Atmos layout, fed into the plugin on channel 3) and runs it through a dedicated decorrelator instance with its own random seed. The signal does not pass through the LR8 crossover — the full bandwidth of the input is sent into the decorrelator, since the LFE feed is already band-limited at the source. This is the distinction from Sub: Sub is the mono sum of all input low bands, intended for bass-managed subwoofer feeds; LFE is the surround-format LFE channel's own content, decorrelated independently and routed back out.

PassThrough Passes the dry input signal through unchanged, but applies delay compensation to match the latency of the decorrelator. This ensures time-alignment with the processed channels, making it useful for reference monitoring or channels that should not be decorrelated.

Off Outputs silence. Use this for any channels that are not connected to physical outputs.

4.2 Per-Channel Output Delay

Each channel has an independent output delay parameter, adjustable from 0 to 10 ms. This allows fine-grained time alignment between speakers at different physical distances. The delay is applied after all other processing (decorrelation, limiter, gain) and is reported as part of the plugin's total latency.

4.3 One Output Channel per Physical Speaker Location

Each of the plugin's 16 output channels is decorrelated independently — internally, every channel index has its own random seed and speaker position, so each output carries a unique decorrelation filter / phase offset / delay tap, not a copy of channel 0's processing. To get the decorrelation benefit you have to preserve that one-to-one mapping between plugin output channels and physical speaker locations (not individual speakers — co-located stacks share a single channel; see below).

Per-location routing — one plugin output channel per speaker location / cluster

Send each physical speaker location (cluster) its own plugin output channel. Do not split one channel to feed speakers at different locations, and do not feed several plugin channels into co-located speakers at one location.

If your installation has, say, four subwoofer stacks distributed around a room, configure four **Sub**-typed output channels (one per stack) and route each plugin output to its own subwoofer. Splitting a single **Sub** output to all four subs in the DAW means every sub plays the exact same samples and they reinforce coherently in the room — no decorrelation benefit between them. The same logic applies to **Full**-typed outputs for multiple main-speaker pairs.

The requirement is the same in every decorrelation mode — the decorrelation filters (or equivalent per-channel randomisation in non-DiSP modes) need to be different across all clusters at every instant in time. This is what produces the spatial averaging across the array.

Same location, multiple subs — they should share the same decorrelation filter

When a single cluster contains several subwoofers physically stacked or co-located, those subs are not separate radiators — acoustically they form one larger one. Feeding co-located subs different decorrelation filters makes

Figure 4.1: The output routing configuration panel.

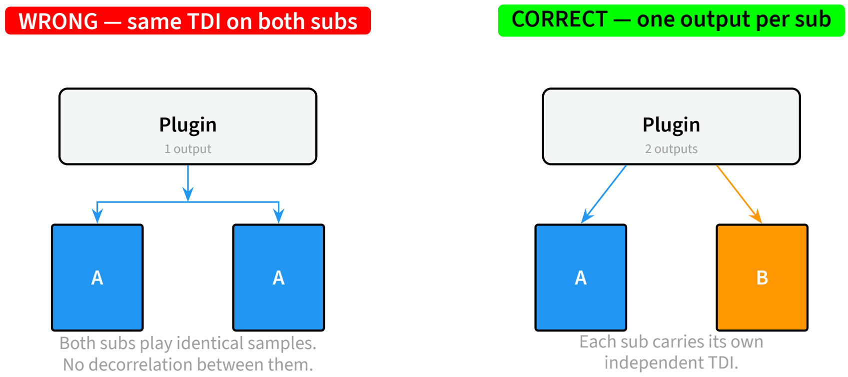

Figure 4.2: Stereo (two-subwoofer) routing. Box colour and letter indicate which TDI that sub plays. **Wrong:** both subs are fed from a single plugin output and therefore receive the same TDI, so they reinforce coherently and contribute no decorrelation. **Correct:** each sub has its own plugin output channel, producing two independent TDIs.

them play different phase-scrambled versions of the same bass at the same point in space, so they interfere destructively with each other and lose in-band energy. The correct configuration is: every sub at one location shares the same decorrelation filter; every location gets a different decorrelation filter. Figure 4.3 illustrates the two failure modes and the correct configuration for a four-cluster setup.

Mono setups — a single decorrelation filter is correct

A mono installation — one subwoofer, or several subwoofers stacked at a single location — has only one cluster, so the "different decorrelation filters across locations" rule is trivially satisfied. There is nothing to differentiate from. Configure one **Sub** output channel, feed all the co-located subs from it, and the single decorrelation filter assigned to that channel is the correct answer.

The plugin still applies its decorrelation filter to that one output, so the emitted signal is phase-scrambled in the usual way — the per-location decorrelation benefit just isn't in play when there is only one location.

4.4 When downstream amps sum L+R bass

Some live-sound amplifier racks may receive a stereo input pair (L, R) and internally sum the two low bands to drive a mono subwoofer feed, while keeping the high bands per-channel. Two amps configured this way typically feed one co-located sub cluster. This is only one of many possible amp topologies — the section below applies when that internal L+R bass summing is in play.

With the plugin upstream in regular **Full** mode, each plugin output already carries an independently decorrelated low band: channel 0 carries the L-input bass decorrelated with one decorrelation filter, channel 1 carries the R-input bass decorrelated with a different decorrelation filter. In real music the L and R low bands are nearly identical, so the amp ends up summing two phase-scrambled versions of the same signal. The sum is incoherent and the sub feed loses roughly -3 dB of in-band bass energy compared to what a single un-decorrelated mono bass channel into the same amp would produce.

Full (Mono Bass) fixes this. The plugin first sums the input L+R into a mono bass bus and decorrelates it once, then sends the same decorrelated bass on every **Full (Mono Bass)** output, with each output's matching input-channel high band added on top. The amp's internal L+R sum is now coherent: the sub-feed level matches the un-decorrelated reference, recovering the ~ 3 dB shortfall from **Full** mode. The high band stays per-channel,

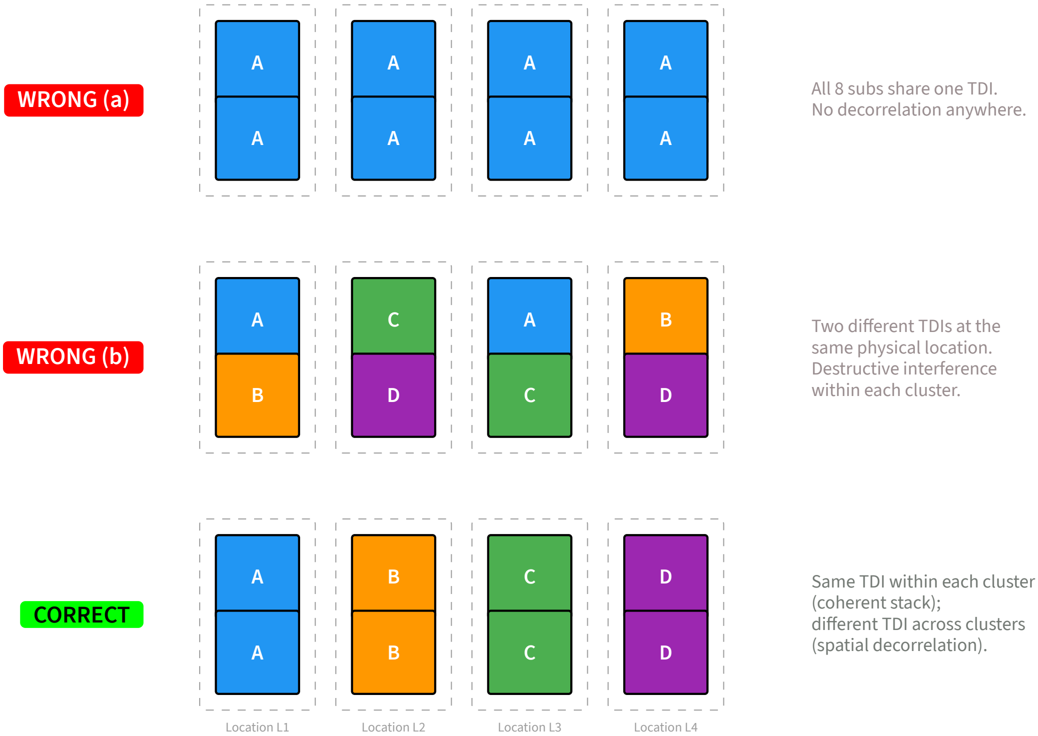

Figure 4.3: Four-cluster routing (two co-located subs per cluster, four cluster locations). Box colour and letter indicate the TDI each sub plays; dashed outlines group subs that share a physical location. **Wrong (a):** a single plugin output feeds every sub — no decorrelation between any of them. **Wrong (b):** eight distinct TDIs, but configured so the two subs at the same location receive different TDIs — they interfere destructively within the cluster, perforating the in-band response. **Correct:** all subs at one location share the same TDI (so they stack coherently as one radiator), and each location's TDI is different from every other location's (so the array is decorrelated across the room).

so the stereo image above the crossover is preserved. Figure 4.4 traces the signal flow for both modes against this amplifier topology.

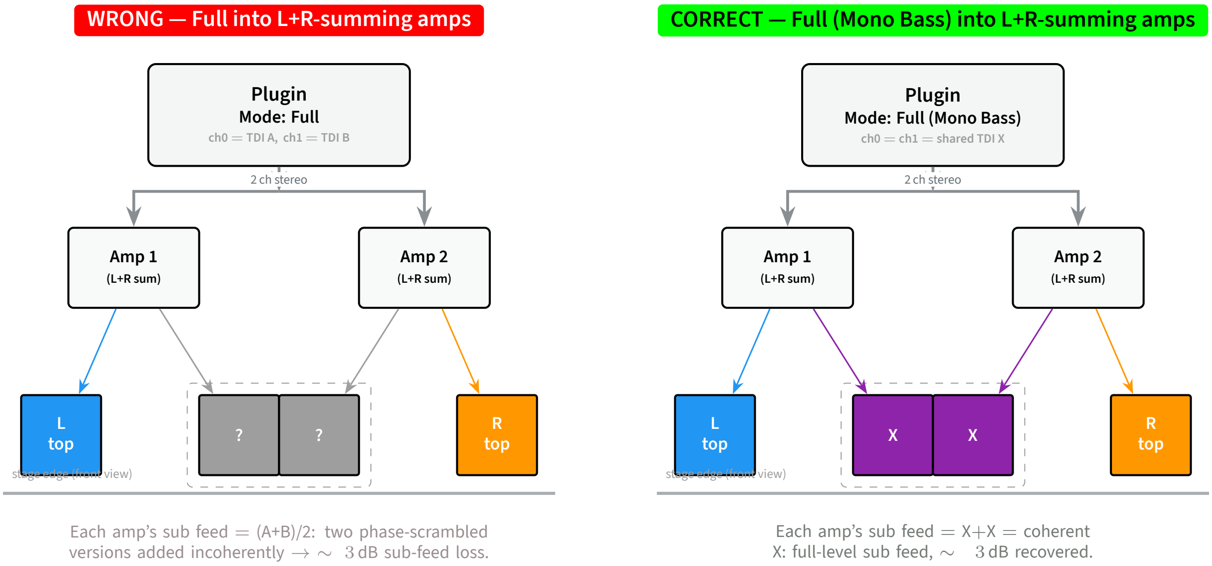

Figure 4.4: L+R-summing amplifier topology, viewed from the audience facing the stage. Two amps ingest the plugin's full stereo output pair; each derives a mono sub feed by summing the two low bands internally and routes the corresponding high band to its top cabinet. The two amps drive one co-located sub cluster (dashed). Box colour and letter indicate the bass TDI carried by each speaker. **Wrong:** in **Full** mode the two plugin outputs carry different TDIs (A, B); the amps' L+R sum is then two phase-scrambled versions of the same signal added incoherently, costing ~ 3 dB at the sub feed. **Correct:** in **Full (Mono Bass)** both outputs carry the same shared TDI X, the amp sum is coherent, and the sub feed is restored. High-band stereo is preserved in both cases.

Pick the mode that matches the amp topology

Full (Mono Bass) **is only correct when each amp ingests the full stereo pair and internally sums L+R for its sub feed.** For physically separated subwoofers — one sub per side, each driven by an amp that does not sum L+R — use **Full** instead. In that wiring the amps preserve the two independently decorrelated low bands on their own outputs, and per-location decorrelation across the left and right subs is exactly what you want (see section 4.3). Using **Full (Mono Bass)** there would feed both physically separated subs the same decorrelated bass signal — two spatially distinct sources radiating identical content into one room. The path-length difference from each sub to any listening position then collapses into a comb-filtered bass response (notches and peaks that shift with seat), which is the exact spectral artifact per-location decorrelation is designed to prevent. Figure 4.5 shows this inverse case alongside the correct **Full**-mode configuration for it.

Tip

Use **Full** for direct-wired stereo mains where the plugin's output pair feeds two separate speakers (or two amps that do not sum L+R for bass). Use **Full (Mono Bass)** when each amp ingests the full stereo pair and derives its own mono sub feed.

4.5 Multi-Speaker Setup

A single Bass Trickkiste instance handles up to 16 output channels and 12 input channels (enough for Atmos 7.1.4). Multi-speaker installations do not require multiple plugin instances coordinated via shared parameters — ev-

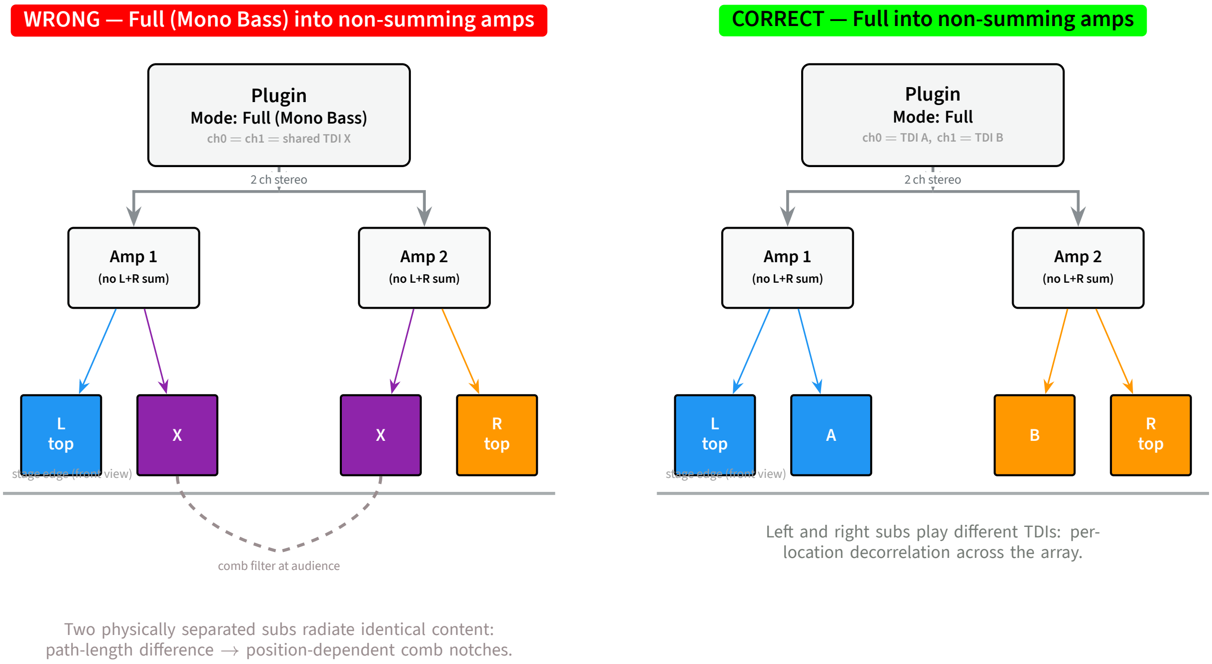

Figure 4.5: Non-summing amplifier topology, viewed from the audience facing the stage. Each amp passes its assigned channel through to one top cabinet plus one physically separated subwoofer. **Wrong:** in **Full (Mono Bass)** both physically separated subs receive the same shared TDI X. Two spatially distinct sources radiating identical content combine at each listening position by simple summation with a path-length delay, producing a comb-filtered bass response that shifts with seat — the exact artefact per-location decorrelation is designed to prevent. **Correct:** in **Full** mode the two subs receive distinct TDIs (A and B); per-location decorrelation does its job.

erything is configured inside one instance using the **Input Mode** parameter and the per-channel Routing Configuration panel.

Input Mode Tells the plugin which input channel layout to expect. Available options: Mono, Stereo, 5.1, 7.1, Atmos 7.1.2, Atmos 7.1.4. The plugin uses this to map the host's input channels onto its internal processing slots so the right input feeds the right decorrelator pair.

Per-channel routing Each of the 16 output channels is independently assigned an output type (Full, Sub, LFE, PassThrough, Off) and an output delay in the Routing Configuration panel. Each output channel runs its own decorrelator instance with a unique random seed and speaker position, so all 16 outputs are independently decorrelated — no coordination between plugin instances required.

Internally, every output channel index has its own seed and speaker position. Each decorrelation algorithm uses these to differentiate its output: Allpass varies its frequency offsets, Delay distributes its tap times, DiSP uses different random seeds, Schroeder uses different prime-ratio spacings, and so on. The result is that even when multiple output channels share the same input (e.g. all Sub outputs derived from the mono input low-band sum), each one carries a uniquely decorrelated version.

Tip

For a 5.1 setup, use one instance with **Input Mode** set to "5.1". Configure channels 0–1 (L/R) and 4–5 (Ls/Rs) as Full, channel 2 (C) as Full, channel 3 (LFE) as LFE-typed. Add additional Sub-typed outputs on free channels (6, 7, …) — one per subwoofer location (cluster), not per individual subwoofer — see section 4.3.

Chapter 5

DiSP Advanced Settings

The DiSP modes (Dynamic and Static) expose an advanced settings panel for tuning the decorrelation character. **The defaults are calibrated to sound good on most material — start there and only reach for these controls if you have a specific reason.**

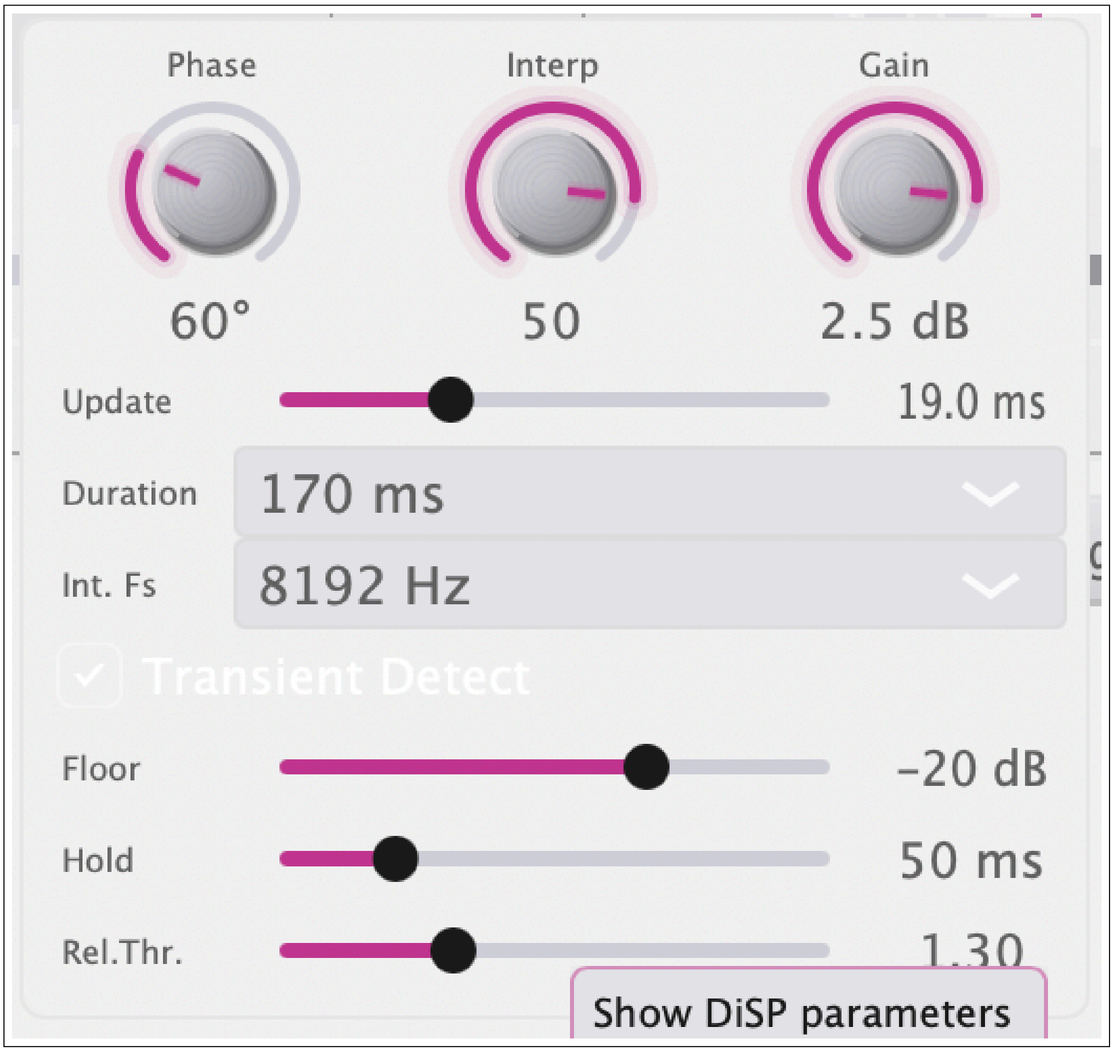

Figure 5.1: The DiSP advanced settings panel.

5.1 Parameters

Internal Fs Default: **8192 Hz**

Trade quality against CPU. Lower it if you need to claw back CPU on large channel counts.

Duration Default: **170 ms**

Filter length. Longer is more decorrelated but uses more CPU and adds latency.

Interp Factor Default: **50** (Dynamic only)

Smoothness of the time-varying decorrelation. Lower it together with **Internal Fs** when reducing CPU.

Update Rate Default: **19 ms** (Dynamic only)

How often the decorrelation changes. Raise it if you hear a subtle tremolo or pulsing on long sustained bass.

Phase Step Default: **60°** (Dynamic only)

Aggressiveness of the time variation. Reduce to 30–45° if you hear artifacts on sustained tones, at the cost of slightly less decorrelation.

Gain Boost Default: **0 dB (auto)**

Compensation gain. At 0 dB the plugin auto-computes the right value for the current **Internal Fs** and **Duration**.

Decay Envelope Default: **on**

Shapes the decorrelation over time. With it on, percussive bass keeps its punch; with it off, sustained bass holds a steadier stereo image.

Transient Detection Default: **on**

Keeps kick-drum and bass attacks sharp when bass and kick share a channel. Turn it off when feeding the plugin a kick-free bass signal (see the workflow recommendation at the start of chapter 3).

Transient Floor Default: -20 **dBFS**

Level below which the transient detector ignores the signal. Raise it if you want the detector to fire only on loud transients.

Transient Hold Default: **50 ms**

How long the transient pass-through stays active after an attack is detected.

Transient Rel. Threshold Default: **1.3**

Fast/slow envelope ratio at which the transient hold ends early. Lower values release back to fully decorrelated steady-state faster.

Regenerate TDIs Forces an immediate rebuild of the filter set. Useful after changing several parameters at once.

Note

Changing DiSP parameters triggers an asynchronous rebuild. Audio continues with the previous filter set until the new one is ready — no dropout or glitch during regeneration.

Chapter 6

Settings

The settings menu is accessed via the gear icon in the top-right corner of the plugin window. These options control the plugin's appearance, behaviour, and resource usage.

6.1 Appearance

Theme Cycles through three display themes: Light, Dark, and System (follows the operating system's appearance setting). The selected theme is persisted as part of the plugin state and restored when the session is reopened.

Tooltips Enables or disables parameter tooltips that appear on hover. When enabled, hovering over any knob, slider, or button displays a brief description of the parameter and its current value.

6.2 Latency Mode

Minimal Each decorrelation mode reports its own native latency. This gives the lowest possible latency for the currently selected mode, but switching modes may cause a latency change that the host DAW needs to re-compensate.

Compensated All modes are padded to match the latency of the highest-latency mode. This allows seamless mode switching without latency jumps or delay compensation changes in the host. The trade-off is slightly higher latency when using low-latency modes such as Allpass or Delay.

6.3 Output Buffer

RT (Real-Time) All DSP processing runs directly on the audio callback thread with zero additional latency or overhead. This is the default and is suitable for most systems.

Buffered (0.5 ms–10 ms) DSP processing moves to a dedicated background thread. The audio callback thread only performs lock-free ring buffer reads and writes. This reduces CPU spikes on the audio thread at the cost of additional latency equal to the selected buffer size. Use this if you experience audio dropouts with CPU-intensive modes such as DiSP on large channel counts.

The Buffer Health panel (visible when a buffered mode is active) shows the current ring buffer fill level and an underrun counter. Click the counter to reset it.

6.4 Reset Plugin State

Restores all parameters to their factory defaults. This includes the decorrelation mode, crossover frequency, amount, mix, output routing, and all DiSP advanced settings. Theme and tooltip preferences are also reset.

Chapter 7

Standalone Mode

Bass Trickkiste can run as a standalone application without a DAW. The standalone build provides its own audio I/O, transport controls, and file processing capabilities, making it useful for quick tests, demonstrations, and batch processing workflows.

7.1 Features

File Import Load WAV or AIFF files via the file menu or by dragging and dropping onto the waveform display. Multi-channel files are supported up to the plugin's 16-channel limit.

File Export Export the processed output to a new WAV file. The export captures the full signal chain including crossover, decorrelation, routing, limiter, and output gain—exactly as it would sound through the plugin in a DAW.

Offline Processing Files are processed faster than real-time when exporting. The progress bar shows the current position during offline rendering.

Waveform Display A scrollable waveform view of the loaded file, with a playhead indicator. Click anywhere on the waveform to seek to that position.

Transport Controls Play, pause, and stop buttons, plus a time display showing the current playback position.

> **Tip**
>
> The standalone application uses the same parameter state as the plugin. You can set up your routing and mode in the standalone, save the preset, and load it later in your DAW session.

Figure 7.1: The standalone application with a loaded audio file.

Chapter 8

System Requirements

8.1 Operating Systems

- **macOS** — 10.13 (High Sierra) or later, Intel and Apple Silicon.
- **Windows** — Windows 10 or later, 64-bit.
- **Linux** — Ubuntu 20.04 or later (or equivalent). ALSA or JACK required for standalone audio I/O.

8.2 Plugin Formats

Bass Trickkiste is available in five formats:

- **AU** (Audio Unit) — macOS only. Installed to **~/Library/Audio/Plug-Ins/Components/**.
- **VST3** — macOS, Windows, Linux.
- **CLAP** — macOS, Windows, Linux.
- **LV2** — Linux only. (LV2's helper toolchain needs to **dlopen** the just-built shared library, which only works on native Linux; cross-compiled and macOS/Windows builds therefore omit this format.)
- **Standalone** — all platforms. No DAW required.

8.3 Hardware

CPU Any modern x86_64 or ARM64 processor. DiSP modes are more CPU-intensive than the simpler algorithms (Allpass, Delay). For 16-channel setups using DiSP, a processor with 4 or more cores is recommended.

RAM Typical memory usage is under 100 MB, including the DiSP filter library. Memory usage scales with channel count and DiSP duration settings.

8.4 Credits & Acknowledgements

JUCE Framework Bass Trickkiste is built with the JUCE framework (**juce.com**).

DiSP Algorithm The DiSP decorrelation modes are based on the work of Moore & Hill, “Decorrelation of multi-loudspeaker reproduced audio using Temporally Diffused Impulses,” Journal of the Audio Engineering Society, 2018.

Catch2 The test suite uses the Catch2 testing framework (v3.5.2).

Source Sans 3 The user interface uses the Source Sans 3 typeface by Paul D. Hunt (Adobe).