L-ISA Studio: Black-Box Identification of Object Rendering

A computational model of output gain, delay, and frequency response, with predictions fixed before measurement

Taebin Yoo | 2026-10-21 | Revision 1.1

Tested components: L-ISA Controller v3.2.2.12429 / L-ISA Processor (Desktop) v3.2.2.5865, 48 kHz.

English translation prepared on 2026-10-25 from the Korean Revision 1.1. L-ISA Studio denotes the software suite; component names identify the software used in the experiment.

Abstract

Object rendering in L-ISA Studio's Processor Desktop was identified through digital electrical loopback measurements. Pan follows the tangent law of two-dimensional vector base amplitude panning (VBAP), using directions relative to the Soundscape origin. A second-order shelf (fc = 517.1 Hz, Q = 0.593) transitions from the geometric mean of amplitude and power normalization at low frequencies to power normalization at high frequencies. Distance separates into a cubic gain curve, G(d) = -(15.88d - 21d^2 + 25d^3) dB, and a first-order digital high-frequency filter. Elevation follows convex-hull-based three-dimensional VBAP.

With these modules fixed before the respective measurements, predictions were within a maximum of 0.03 dB for a new 16-speaker arc, order and ID changes, group-property changes, 9.1.6, held-out Island A, an asymmetric six-speaker layout, and 120 randomly selected output values with w = 0. The initial Elevation prediction for held-out Island B failed by up to 54.4 dB. The cause was tie-breaking in the triangulation of a coplanar quadrilateral face. Subsequent tests with predictions fixed in advance rejected hwOut, group order, names, position, and aiming angle as the deciding factors, and demonstrated dependence on speaker UUIDs. The resulting descending-UUID incremental-insertion rule predicted 144 output values in a UUID-modified Island copy with a maximum error of 0.027 dB.

Width was identified only partly: a two-tap decorrelation structure with a 24-sample delay and an angular support of +/-30 degrees times w were found, but its distribution kernel and normalization remain unresolved. Input delay was not applied to the audio output in this configuration. Post-hoc fits (1,936 output values at w = 0, excluding one mirrored session; maximum error 0.027 dB) are distinguished from predictions fixed before measurement. Unresolved modules are reported with numerical bounds.

Keywords: L-ISA Studio; object-based audio; VBAP; black-box measurement; system identification; prediction before measurement.

1 Background and research questions

The preceding Sound xR Image study (`reference/AFC_SoundxR_report_ko.pdf`) showed that a post-hoc fit of 0.006 dB could coexist with a failed prediction of 9.29 dB on a new layout. This study applies the same principle to L-ISA Studio's Processor Desktop and asks:

1. How do Pan, Width, Distance, and Elevation change each speaker's output gain, delay, and frequency response?
2. Do group properties (VirtualDistance and MaxGainGradient), Usage, processing mode (2D/3D), and speaker order or ID affect rendering?
3. Do predictions fixed before measurement work on templates and layouts not used to develop the model?

The objective is an input-output computational model (`model.py`), rather than recovery of the manufacturer's internal implementation.

2 Experimental method and data processing

Signal path. Python (`tools/rig.py`) -> L-ISA Audio Bridge input -> Processor Desktop -> BlackHole 64ch outputs 1-16 -> Python recording. The Processor output window was fixed at 16 channels. The 22-speaker 18.1.4 Island layout was therefore split into overlapping sessions A (speakers 1-16) and B (speakers 7-22). Early, cluster, and late reverberation were off, and FX was disabled for all sources.

Test material. The study scope was restricted to three templates: 7.1 Red Rocks (2D frontal) and 9.1.6 Atmos (360-degree ring plus Tops, 3D) for model development, and 18.1.4 Bluspace Island as a holdout. Derived sessions were copies of Red Rocks, 9.1.6, or Island B. Changes included order permutations, VD/MGG changes, an asymmetric six-speaker layout, dense arcs at 5- and 7.5-degree spacing, Tops order/name/mirror/UUID/aiming-angle changes, and an Island B UUID exchange. Original templates were not modified.

Time reference. Playback and recording used different CoreAudio devices and were not sample-synchronized. Within a single recording stream, ch2 (Source 2: center, d = 0.15, reverb and delay off) was played first as a reference sweep, followed 2.6 s later by ch1, the test source. Delay is defined relative to that reference.

Deconvolution and features. A 2 s logarithmic sweep at -20 dBFS was processed using regularized spectral division by the playback signal, with cosine tapers at the band edges. Leakage between the two sweeps was about -156 dB. Gain L is the mean power in the one-third-octave band centered at 1 kHz, expressed in dB. Delay is evaluated from the 200-4000 Hz phase slope in samples and from a parabolically interpolated peak. One-third-octave responses were also saved.

Residuals. Errors are defined as follows:

```
e_j = L_meas,j - L_model,j                              (R)
E_RMS = sqrt(mean(e_j^2))
```

Digitally silent outputs and outputs below -120 dB, the detection threshold, were recorded as observed and excluded from comparison. They were not replaced with symmetric values or zero. Complete settings for all outputs are in `data/measurements.csv`.

Measurement quality. The path is deterministic. Repeated identical settings differed by 0.0000 dB, and measured gain differed by 0.0000 dB across -30, -20, and -10 dBFS inputs, indicating no dynamic processing in the tested path. Settling after an OSC change was 13 ms; the measurement wait was

set to 0.25 s. Forward/reverse sweep differences were 0.000 dB for Pan and 0.0005 dB for Width. The approximately 0.006 dB residual floor is a numerical consequence of the measurement window and band averaging, not an estimate of equipment uncertainty.

Measurement-path error and correction (Cycle 1). Phase 0 opened a single-channel output stream. CoreAudio duplicated it onto Audio Bridge ch1 and ch2, mixing in a fixed central source with reverb enabled. The Phase 0 component on outputs 4-7 that appeared independent of p was this reverberation tail. Subsequent measurements used an eight-channel stream with only ch1 driven. The Phase 0 raw-value table was invalidated; only the settling, linearity, and repeatability conclusions were retained.

3 Model and equations

Notation: speaker i has azimuth phi_i relative to the Soundscape origin (0 degrees = +Y, positive toward +X), elevation epsilon_i, and unit vector u_i. The object parameters p, w, d, and e are OSC-normalized values in [0, 1]. ASCII notation below uses `norm1` and `norm2` for the L1 and L2 norms, `^T` for transpose, and `^-1` for inversion.

Pan angle (prediction fixed before c02; successful holdout tests)

```
theta(p) = theta_min + p * (theta_max - theta_min)       (1)
Auto (frontal): theta_min = min_i(phi_i),
                theta_max = max_i(phi_i)               (1b)
360-degree range: theta_min = -180, theta_max = 180
```

Height/Bottom groups are excluded from the Auto range, which is recalculated when the session loads.

2D pair panning (prediction fixed before c02). Select the adjacent pair (i, j) bracketing theta in azimuth order; the ring wraps in a 360-degree range.

```
[g_i, g_j]^T = [u_i, u_j]^-1 * u(theta)                (2)
```

Negative gains are clipped to zero; gains for all other speakers are zero. Scene, Extension, Surround, and 360 groups are sorted together in a single planar ring regardless of Z (the ExL-Sc1 pair in Red Rocks and Rear-Side pairs in 9.1.6).

Frequency-dependent normalization (post-hoc c02 fit, two shared parameters; subsequently held fixed for predictions)

```
n_H,i = g_i / norm2(g)
n_L,i = g_i / sqrt(norm1(g) * norm2(g))
N_i(s) = (n_L,i + sqrt(n_L,i*n_H,i)*s/Q + n_H,i*s^2)
         / (1 + s/Q + s^2),   s = j*f/f_c               (3)
f_c = 517.08 Hz,   Q = 0.5932
```

For a single active speaker, N = 1. At the center of a pair, the low-frequency level is up to 1.5 dB below the high-frequency level, equal to 0.5 * 20log10(norm1(g)/norm2(g)).

Distance gain (post-hoc c04; residual below 0.0001 dB at 21 points)

```
G(d) = -(15.88*d - 21*d^2 + 25*d^3) dB                 (4)
```

This applies when `distatt/gain = 1`; otherwise G = 0.

Distance HF attenuation (post-hoc c04; empirical table)

```
H_HF(z) = (1-a)/(1-b) * (1 - b*z^-1)/(1 - a*z^-1)      (5)
```

The coefficients a(d) and b(d) are spline-interpolated from `data/c04_hf_fit.json`. At d = 1, the Nyquist gain is -20.0 dB and the effect at 1 kHz is -0.31 dB. Closed forms for a(d) and b(d) remain unresolved.

3D Elevation (prediction fixed before c08). A Height/Bottom group enables 3D mode. Object elevation is psi = 90 degrees * e. Only Height-group speakers use atan(z / horizontal distance) for speaker elevation; the other speakers are assigned zero elevation in this model.

```
[g_a, g_b, g_c]^T = [u_a, u_b, u_c]^-1 * u(theta, psi) (6)
```

Select a convex-hull triangle with all gains nonnegative; exclude lower faces. For a coplanar polygonal face, triangulate by inserting vertices in descending UUID-string order:

```
First three vertices -> triangle.                    (7)
Each subsequent vertex -> triangle with its two
neighbors on the current polygon boundary.
For a quadrilateral: the diagonal avoids the minimum UUID.
```

Rule (7) was derived after c18 and successfully predicted c20. For e = 0, equation (6) reduces to equation (2).

Output response and delay

```
H_i(f) = N_i(f) * 10^(G(d)/20) * H_HF(f; d)
L_i = 10*log10(mean_1kHz_third_octave(abs(H_i(f))^2))    (O)
tau_i = phase_slope[H_i] - phase_slope[H_HF(.; 0.15)]   (8)
```

Input delay was not applied to the output in this configuration (Section 4.4). Width (w > 0) is excluded from equation (O); `model.py` raises an exception for w > 0 (Section 6).

4 Results by parameter

4.1 Pan: Red Rocks, 9.1.6, and dense arcs

The c02 prediction H1 (VBAP tangent law, power normalization, and a constant) matched the active-speaker set in all 51 conditions and 99 output values (RMS 0.139 dB, maximum 0.199 dB). DBAP candidate H2 was rejected because its active set was wrong in all 51 conditions. Direction inferred from the measured gain ratio agreed with theta(p) within 1e-4 degrees; the remaining 0.2 dB corresponded to the power deficit of 0.955 at the pair center.

One-third-octave responses showed that both channels in a pair shared the same filter. At high frequencies it agreed with power normalization within 0.01 dB, while at low frequencies it moved toward amplitude normalization. Equation (3), with two shared parameters, fit 96 two-speaker response curves from 30 Hz to 16 kHz within 0.012 dB and 0.10 degrees.

Table 1. Pan residuals at 1 kHz. "Before" denotes a prediction fixed before measurement; "after" denotes a post-hoc fit.

Condition / raw data	Type	Records	Outputs n	RMS dB	Max dB
Red Rocks, 51 p points / `c02_pan.jsonl`	Before, H1	51	99	0.139	0.199
Red Rocks, 51 p points x 2 / `c02_pan.jsonl`	After, (1)-(3)	102	198	0.0032	0.0061
Dense 16-speaker arc, 5 degrees / `c05_dense.jsonl`	Before	23	40	0.0031	0.0058
9.1.6, 360-degree ring / `c08_916.jsonl`	Before	41	77	0.0035	0.0061
Asymmetric six-speaker layout / `h02_asym6.jsonl`	Before, (1b)	41	80	0.0032	0.0062
Island A/B, 360 degrees / `h01_island_*.jsonl`	Before	82	118	0.0034	0.0062

Alternative Auto-range hypotheses for the asymmetric layout were also recorded before measurement. Both a symmetric +/-maximum range (RMS 6.13 dB) and the range stored in the session (RMS 9.23 dB) were rejected. On loading, the Controller recalculates Auto as [minimum azimuth, maximum azimuth].

4.2 Distance

The test used four gain/HF on-off combinations, p in {0.5, 0.44}, and 21 d points. The gain and HF paths separate multiplicatively, with a combined error of 0.0004 dB. The panning ratio is independent of d. The gain curve is the cubic in equation (4), independent of p: p = 0.5, 0.44, and 0.1 all produced -20.188 dB at d = 1, including the HF contribution at 1 kHz. The initially observed -1.995 dB constant is G(0.15) plus the HF contribution at d = 0.15. The first-order pole-zero filter in equation (5) fits all 21 points within 0.014 dB.

Table 2. Distance residuals: post-hoc model at 1 kHz, except rows explicitly marked before measurement.

Condition / raw data	Records	Outputs n	RMS dB	Max dB
Four switches x 21 d points / `c04_distance.jsonl`	168	252	0.0049	0.0061
UI delay ON; 21 d points x 3 p values / `c04b_delay.jsonl`	63	105	0.0054	0.0061
9.1.6 and Island A/B (before) / `c08`, `h01_*`	9	9	0.0004	0.0004
Asymmetric six-speaker layout (before) / `h02_asym6.jsonl`	6	12	0.0045	0.0062

4.3 Elevation: 3D

In 9.1.6, prediction E1, equation (6), was fixed before measuring four p values and ten e steps. It gave n = 106, RMS 0.0145 dB, maximum 0.0271 dB, with the active set correct in 40/40 cases. Candidate E2, a 1/d DBAP model motivated by the documentation's description of distance interpolation between the three nearest speakers, was rejected (RMS 5.87 dB, maximum 29.0 dB). Holdout results are reported in Section 7.

4.4 Delay

The OSC address `/ext/config/src/(k)/delay/enable` did not change the state; `delay/mode` did. Enabling Delay in the UI displayed Total Delay = 30.0 ms in dynamic mode at d = 1, but the audio output was not delayed. A static 12 ms setting behaved the same way. In this Processor Desktop configuration, therefore, the applied output delay was zero; observed timing came from the filter phase in equations (3) and (5). Across all 1,712 output values at w = 0, equation (8) had RMS and maximum residuals of 0.020 and 0.136 samples. The dynamic mapping (`DynamicDelayMax = 30 ms`) could not be observed (Section 9).

4.5 Group properties

In copies with both groups' VirtualDistance changed from 7 to 14, or MaxGainGradient from 6 to 12, 29 conditions covering Pan, Distance (d = 0..1), and Width were measured. The maximum difference from the original was 0.0000 dB. Neither property affected the tested 2D direct-sound path (gain, HF, and Width). Their roles in reverb, Spatial Fill, and 3D were outside the measurement scope.

5 Order and identity controls

2D: Red Rocks, c06. At the same coordinates, UUIDs, names, XML order, within-group order, group order, and output numbers were all changed. Predictions fixed before measurement gave n = 54, RMS 0.0037 dB, and maximum 0.0061 dB. After remapping the outputs, direct comparison with the original showed a maximum difference of 0.0000 dB across 29 conditions including Width. The tested 2D rendering did not depend on speaker identity or order.

3D: derived 9.1.6 and Island sessions, c11-c20. Copies were changed one factor at a time to identify what determines coplanar-face triangulation (Section 7). Table 3 shows that speaker identity, specifically UUID, can alter the 3D output.

Table 3. Tests of 3D tie-breaking hypotheses with predictions fixed before measurement: Elevation at 1 kHz. Times are prediction-fix times from the experiment log. "Mismatch" counts active-set disagreements as recorded in the source report.

Copy / change / time	Hypothesis	Records	Outputs n	RMS dB	Max dB	Mismatch	Result
topsorder: shuffle Tops hwOut + reverse groups / 02:59	Minimum hwOut	48	70	5.68	13.01	36	Rejected
topsorder / 02:59	Original triangulation	48	122	0.0137	0.0272	0	Fits
topsrot: rotate group list / 03:06	First in group list	24	46	9.42	27.45	24	Rejected
topsrot / 03:06	Original triangulation	24	71	0.0159	0.0270	0	Fits
rename: LTF to Z-... / 03:10	First name alphabetically	24	46	9.42	27.45	24	Rejected
rename / 03:10	Original triangulation	24	71	0.0159	0.0270	0	Fits
mirror: reverse all X / 03:15	Geometric left-front rule	24	47	10.58	21.56	23	Rejected
mirror / 03:15	Keep original speaker pairs	24	71	0.082	0.488	1	Near fit
uuidswap: swap left/right Tops UUIDs only / 03:42	Keep diagonal	20	40	8.94	15.95	20	Rejected
uuidswap / 03:42	Diagonal follows UUIDs	20	60	0.0144	0.0253	0	Fits
aimflip: reverse Tops aiming angles only / 03:42	Keep diagonal	20	60	0.0144	0.0253	0	Fits; aiming irrelevant
Island B uuidperm: exchange H1/H4 UUIDs / 03:49	UUID insertion, (7)	48	144	0.0158	0.0269	0	Success
Island B uuidperm / 03:49	UUID-independent original diagonal	48	97	13.02	54.47	47	Rejected

An XML-position criterion had already been rejected on existing 9.1.6 and Island B data (RMS 7.0 and 10.5 dB). Equation (7) explains eight of nine 3D sessions with a maximum error of 0.027 dB. One condition, e = 0.6 and p = 0, on the rear mixed quadrilateral (two Tops plus two Rears) of the mirrored copy is an exception, with a maximum error of 0.49 dB. Its cause remains unresolved (Section 9).

6 Width: partial identification

All three predictions fixed before the Width measurements failed: continuous MDAP with full width, continuous MDAP with half width, and an unnormalized sum (Table 4). Post-hoc analysis established the following:

1. Each output has the structure N(f) cascaded with (D_i + E_i*z^-24). The 0.5 ms, 24-sample tap creates notches or peaks at odd multiples of 1 kHz, directly in the 1 kHz measurement band. A per-condition fit using one shelf ratio plus D_i and E_i was within 0.15 dB across all 55 conditions.
2. The angular support is [theta - 30 degrees*w, theta + 30 degrees*w]. It matches all measured speaker-onset thresholds in Red Rocks, the dense arc, and the 7.5-degree arc.
3. The sign of E follows even/odd parity in azimuth-sorted speaker order. Its magnitude follows s = 0.70*abs(a_i - a_j), where a is the amplitude-normalized VBAP vector for the theta pair. The parity group with fewer speakers has E/D = +/-s; the larger group is scaled down so that sum(E) = 0. At

Red Rocks p = 0.44/0.3/0.1, predicted values -0.0396/-0.0416/-0.0164 compare with measured -0.0396/-0.0417/-0.0163. At the pair center, E = 0.

4. In the dense 5-degree arc with half-width hw >= 12 degrees, D can be fitted within 0.001 by virtual sources at theta, a theta +/-10-degree*k grid, and the +/-hw edges, with weights 1/0.615/0.318/0.241. The same weights do not fit Red Rocks (c20 = 0.25-0.37, as labeled in the analysis).
5. The shelf-normalization ratio is not exactly given by norm1(D)/norm2(D), which differs by up to 0.53 dB, or by the power of D +/- E, which differs by 0.86-0.91 dB.

Table 4. Width tests with predictions fixed before measurement. D and E/D errors are not dB unless explicitly marked.

Prediction / raw data	Target	Records	Outputs n	RMS	Max	Result
W1 full-width MDAP / `c03_width.jsonl`	1 kHz dB	55	216	4.62 dB	22.29 dB	Failed
W2 half-width MDAP / same file	1 kHz dB	55	216	10.63 dB	43.63 dB	Failed
W3 unnormalized sum / same file	1 kHz dB	55	216	4.47 dB	21.70 dB	Failed
W*, derived after c05 / `c14_width_dense10.jsonl`	Normalized D	30	-	-	0.096	Partial
W* / same file	E_i/D_i	30	-	-	0.50 (theta = 0: 0.005)	Partial
Ignore Width: random 50 / `r50_*.jsonl`	1 kHz dB	50	120	7.4-10.6 dB	42.9 dB	Failed

7 Predictions fixed before measurement, including failures

Table 5. Prediction results. The full list is Table B in `data/validation_tables.md`. Times are the experiment-log times at which predictions were fixed.

Test	Time	Quantity	Outputs n	RMS dB	Max dB	Outcome
c02 Red Rocks Pan H1	00:38	1 kHz	99	0.139	0.199	Structure correct; 0.1 dB target missed
c03 Width W1-W3	01:02	1 kHz	216	>=4.47	>=21.7	Failed
c04 dynamic delay, 30 ms*d	01:02	Delay	-	-	-	Failed: delay not applied
c05 dense 16-speaker Pan	01:24	1 kHz	40	0.0031	0.0058	Success
c06 order/ID changes	02:01	1 kHz	54	0.0037	0.0061	Success
c07 VD 14 / MGG 12	02:01	1 kHz	108	0.0037	0.0061	Success
c08 9.1.6 Pan / Elev E1 / Dist	02:17	1 kHz	186	<=0.0145	0.0271	Success; E2 rejected
H1 Island A: holdout	02:28	Pan / Elev / Dist	85	<=0.0162	0.0255	Success
H1 Island B: holdout	02:28	Pan / Dist	58	0.0034	0.0062	Success
H1 Island B: holdout	02:28	Elevation	86	10.53	54.44	Failed: 24/40 active-set mismatches
H2 asymmetric six-speaker layout	02:28	Pan / Dist	92	0.0034	0.0062	Success
Random w = 0: RR 30 + 9.1.6 20	02:41	1 kHz	120	0.0081	0.0228	Success
c11-c13 3D tie-breaking	02:59-03:15	Elevation	Table 3	-	-	Four candidates rejected
c14 Width W*	03:34	D, E/D	-	-	0.50	Partial
c18 UUID swap / c19 aiming flip	03:42	Elevation	120	0.0144	0.0253	UUID dependence; aiming irrelevant
c20 Island B UUID exchange, (7)	03:49	Elevation	144	0.0158	0.0269	Success

Island B failure. The four height speakers form a left-right-symmetric isosceles trapezoid, producing an exactly coplanar quadrilateral on the convex hull (zero volume). The prediction used diagonal H2-H4, whereas the measured renderer used H1-H3. At e >= 0.5, measured values were a left-right mirror of the prediction. The initial geometric rule, a left-front fan, was rejected by the mirrored-session test with a prediction fixed in advance.

The UUID exchange in c18 then demonstrated UUID dependence. Equation (7), developed from that result, successfully predicted a copy of the held-out geometry with only the UUIDs changed (c20: 144 output values, maximum 0.027 dB). The original Island B holdout remains recorded as a failure. Reproducing it afterwards with equation (7) gives 110 output values with a maximum error of 0.026 dB.

Operational record. At 02:41, an import side effect rewrote the holdout and asymmetric-layout prediction files. Their inputs (`model.py`, 02:17; prediction script, 02:28:39; HF table, 01:52) had not changed, so the values were identical. This is documented in each file's `audit` field.

8 Discussion and model limits

- The direct-sound core, equations (1)-(6) and (8), is determined by geometry. The absence of order/ID dependence in the tested 2D cases contrasts with the AFC study, where reassigning Speaker IDs changed the distribution. Predictions on five classes of layouts not used for fitting (dense arc, order permutation, 9.1.6, Island A, and asymmetric layout) were within 0.03 dB.
- Successful predictions are restricted to w = 0. For 3D layouts with coplanar polygonal faces, the output depends on speaker UUIDs. Recreating a session with new UUIDs can therefore change a diagonal even when geometry is unchanged. Equation (7) was tested prospectively on quadrilateral faces; evidence is weaker for coplanar faces with five or more vertices and for mixed-group faces, where one mirrored exception remains.
- The f_c and Q of equation (3) and coefficients of equation (4) were fitted afterwards on Red Rocks. They were then kept fixed for predictions on the other layouts, supporting their treatment as layout-independent constants within the tested scope. The a(d) and b(d) in equation (5) are tabulated. Interpolation errors between d-grid steps of 0.05 were not directly measured; 120 output values with random d and a maximum error of 0.023 dB provide indirect evidence.
- These are static transfer-function measurements through a digital electrical path, not listening tests or room-acoustic measurements. The 1 kHz gain is highly sensitive to Width's comb structure, so Width must also be interpreted using the full one-third-octave responses in `data/measurements.csv`.
- To reduce ambiguity where a gain distribution alone does not uniquely identify an internal position, direction inferred from ratios (within 1e-4 degrees) and frequency response (shelf shape) were used as independent evidence. Even so, the Width D kernel was not uniquely determined by three speakers in Red Rocks: grid spacings of 8-12 degrees all fit exactly.

9 Unresolved modules

Module	Established in the tests	Unresolved	Numerical bound
Width, w > 0	Two taps: 0 and 24 samples; +/-30-degree*w support; E sign from azimuth parity and magnitude rule; E = 0 at pair center	D kernel for general layouts; shelf normalization at w > 0; general-layout validity of E rule	Ignoring Width: max 42.9 dB. W* before measurement: D error 0.096; E/D error 0.50
Coplanar-face triangulation, 3D	UUID dependence (c18); no effect from hwOut, XML/group order, names, position, or aiming in the controls; quadrilateral prediction success for (7), c20	Faces with five or more vertices; one mirrored mixed-quadrilateral exception	Exception: max 0.49 dB. Wrong diagonal: up to 54.4 dB
Static/dynamic input delay	OSC enable ineffective; no audio delay even with UI ON	Hardware Processor behavior; dynamic mapping	Applied delay: 0 +/-0.14 samples
HF coefficients a(d), b(d)	First-order pole-zero filter; Nyquist -20 dB at d = 1	Closed form	Table interpolation; fit <=0.014 dB
VirtualDistance / MaxGainGradient	Difference on tested 2D direct-sound path: 0.0000 dB	Reverb, Spatial Fill, and 3D paths	-

10 Conclusion

The tested direct-sound rendering of L-ISA Studio's Processor Desktop v3.2.2 can be described by origin-referenced 2D/3D VBAP, a frequency-dependent normalization shelf, a cubic distance gain, and a first-order HF filter. Predictions fixed before the respective measurements were within 0.03 dB on five classes of new layouts and random combinations at w = 0.

Termination criterion 1 (50 random combinations including Width) was not met because Width remains unresolved. Criterion 2 was met for Island A and the asymmetric layout, but not overall: Island B Elevation failed by 54 dB because of coplanar-face tie-breaking. Subsequent tests identified the cause in the tested quadrilateral cases (descending-UUID incremental insertion; successful c20 prediction). Width was identified only structurally: two-tap decorrelation, angular support, and an E rule, with its unresolved parts explicitly bounded.

The experiment ended at the 20-cycle limit. Next steps are to separate Width's D kernel and normalization at w > 0 using several nonuniform dense layouts, and to test equation (7) on coplanar faces with five or more vertices, such as hexagonal Tops.

Materials and reproducibility

The following paths identify materials in the original research workspace. Listing a path does not mean the file is included in this PDF or publicly distributed with it.

Path	Contents
`model.py`	Reference implementation: load_layout, gains_1k, response, delays; raises an exception for w > 0
`validate.py`	Recomputes Table A (post-hoc) and Table B (before measurement) from all raw data, writing `data/validation_tables.md`
`data/measurements.csv`	21,520 output-level rows: session, template, 2D/3D, pan range, origin, group Usage/VD/MGG, p w d e s and physical values, switches, actual UI delay state, version, speaker coordinates, gain, delay, and one-third-octave response
`data/raw/*.jsonl`, `data/irs/*.npy`	Per-capture raw records and test-source impulse responses
`predictions/*.json`	Predictions fixed before measurement, with creation time and model.py SHA-1
`sessions/`	Measurement and derived session copies, plus output-mapping JSON files
`tools/rig.py`, `tools/measure_session.py`	Measurement tools; `tools/predict_*.py` fixes predictions; `tools/a0*_*.py` performs post-hoc analysis
`NOTES.md`	Cycle logs: hypotheses, prediction files, measurements, residuals, and conclusions

To reproduce: load the session in the UI, fix the prediction with `tools/predict_*.py`, measure with `tools/measure_session.py <tag>`, then run `python validate.py`. Before measurement, Source 1 and Source 2 must grant OSC control of PWDES, with reverb off and, if possible, the reverb master muted. `rig.Rig` applies the source switches.