The Matter Stack — Living Instrument

Art & Expressive-Medium Addendum: an 8×8 sound-reactive matter display

Companion to The Matter Stack: A Layered Physical–Acoustic–Photonic Display · Design & build addendum — July 2026

1. Purpose and Scope

This addendum adapts the base Matter Stack build into an expressive art instrument rather than a data display or a science exhibit. It specifies three changes to the core guide: an upgrade of the magnetic layer from a 4×4 to an **8×8 electromagnet grid** for smoother, higher-count terrain; a **sound-reactive engine** that turns live room audio into synchronized motion of matter, sound, and light; and a **timeline/choreography system** that supports scripted “micro-theater” pieces. It closes with a staged roadmap that ends at a hardened, collaborative, museum-grade installation — the direction to build toward, not to build first.

Two creative modes share one machine. In **data-as-art** mode the installation listens to the room and renders the ambient soundscape as living terrain, a darting light, and shifting color — the crowd itself becomes the composer. In **micro-theater** mode a scripted score drives the three layers through a short narrative: terrain as landscape, the levitated bead as a character or star, light as sky and emotion. The two modes are not exclusive; the strongest pieces use a scripted scaffold that the room's sound modulates in real time.

> **The governing idea: temporal composition, not spatial resolution.** Even at 8×8 the terrain holds only a handful of independent features. Treat the system not as a low-resolution screen but as an instrument with a few slow “keys,” one fast moving voice, and a full palette of light — expressive through *sequence and timing*, the way most music is expressive with few notes. This reframing is what keeps a low feature count from reading as a limitation.

2. The Instrument: A Two-Speed Voice

The three layers divide naturally by timescale, and that division is the instrument's design. The **ferrofluid terrain is the slow voice** — features change over roughly 0.2–1 second, so it behaves like a pad or drone: evolving harmonies, swelling landscapes, textures that breathe. The **levitated bead is the fast voice** — it moves at metres per second, so it is the lead: quick melodic runs, accents, a point of light that darts and hovers. The **photonic layer is the palette** — projected color, intensity, and label, painting mood across both physical layers and carrying any detail the matter cannot.

Composing across these three at once, locked to a shared clock and a shared coordinate frame, is the actual art form here. No single-medium instrument can do it, and the registration problem it forces — keeping light aligned to matter that is moving — is the same problem the base guide identifies as the holodeck's real research question. Building this piece is therefore both an artwork and a genuine step down that road.

Why the bead is also the anchor

The bead's trajectory is known in advance because the drive code commands it, so it doubles as a bright, moving optical fiducial that anchors the whole coordinate frame at zero latency. The terrain's shape, by contrast, is an emergent physical result and must be *sensed* — which is why the depth camera (Section 4) matters for any piece where light must register precisely to the fluid rather than to commanded values.

3. Hardware Delta — From 4×4 to 8×8

The magnetic layer is the only subsystem that changes substantially. Moving from 16 coils at 30 mm pitch to 64 coils at roughly 15 mm pitch roughly doubles spatial resolution per axis, letting terrain features be smaller, closer, and morph more smoothly. Three engineering realities come with that jump.

3.1 Addressing 64 channels

An Arduino Mega has only about fifteen hardware-PWM pins — nowhere near 64. Bit-banging that many soft-PWM channels is unreliable, so the drive path changes to dedicated PWM driver ICs. The recommended route is four **PCA9685** boards (16 channels each, 12-bit PWM, driven over a single I²C bus): beginner-friendly, cheap, and more than fast enough, since coil switching needs only hundreds of hertz. Each of the 64 channels still drives its own logic-level MOSFET and — non-negotiably, exactly as in the base guide — its own flyback diode across the coil. Omitting even one diode still eventually destroys the drive electronics.

3.2 Power and heat become the real ceiling

Sixty-four holding coils drawing 0.3–0.5 A each cannot all run at once. Budget a **12 V, 20–30 A** supply with the coil rail fused and 16 AWG or heavier distribution, and treat the thermal budget as firmware, not a suggestion: cap duty cycle around 60–70 percent, cap how many coils may exceed 50 percent simultaneously, and add active cooling — heatsinks on the coil bank and forced-air fans underneath. As the base guide states, heat is this subsystem's main reliability enemy; at 4× the coil count it is the dominant design constraint. The practical consequence for composition: you have a fixed “power budget” to spend each frame — a few tall features, or many gentle ones, never everything at full height.

3.3 The fluid, not the grid, sets the new limit

Below a certain spacing, neighboring peaks merge no matter how fine the coil pitch, because the ferrofluid's own spike spacing (the Rosensweig instability, set by fluid density and surface tension) takes over as the binding constraint. In practice an 8×8 grid yields roughly **8–12 cleanly separable features** rather than a literal 64. Shaped pole pieces on each coil sharpen the field and buy a little more separation; going denser than 8×8 mostly wastes coils. Keep everything else from the base build unchanged: the sealed saline-suspended ferrofluid cell, the horizontally-mounted TinyLev gantry 10–15 cm above the fluid, and a ≥15 cm standoff between driver boards and the strongest magnets.

4. Software Architecture — The Choreography Engine

This is where most of the new work lives. A host computer runs a real-time engine that keeps all three layers locked to one clock and one coordinate frame, and drives the two Arduinos over USB serial. The engine is organized as five cooperating modules feeding a shared timeline.

Module	Job	Output
Master clock & scene graph	Single timeline all layers subscribe to; holds the current scene, cues, and blend state.	Frame tick + scene parameters
Terrain engine	Turns a target “heightfield” into per-coil PWM with easing/cross-fade so bumps travel smoothly; enforces the thermal budget.	64-channel duty values → Mega
Bead choreographer	Turns a target path into the levitator's phase-shift commands; keeps the bead as a known fiducial.	Trajectory + phase → Nano
Photonic renderer	Projection-maps color, intensity, and labels onto terrain and bead; warped to the cell.	Projector / AR frame
Room-sound analyzer	Captures a microphone feed and extracts musical features in real time (Section 5).	Feature vector per frame

A sixth, optional **registration** service closes the loop: a depth camera above the cell measures the fluid's actual surface and corrects the photonic renderer so light lands on real, not commanded, geometry. It is not required for a first gallery piece but is the pivot toward the museum build.

The data flow is simple to state: microphone → analyzer → (blended with the scripted scene) → the three layer engines → hardware. The scene graph decides, per piece, how much control comes from the live sound versus the script.

5. Data-as-Art: Mapping Room Sound to the Stack

The chosen data source is the sound of the room itself — voices, footsteps, music, the hush of a quiet gallery. A single microphone feeds a real-time analyzer (an FFT plus a few feature extractors) that produces a compact feature vector every frame. Each feature is routed to the layer whose timescale matches it. The core insight is to **match signal speed to actuator speed**: slow envelopes drive the slow terrain, fast transients drive the fast bead.

Audio feature	Speed	Drives	Effect
Low-band energy (bass, 20–250 Hz)	Slow envelope	Central terrain coils	A broad swell rises and falls with the low end

Mid-band energy (250 Hz–4 kHz)	Medium	Edge / ring terrain coils	Ripples and secondary peaks around the rim
Onsets / transients (beats, claps)	Fast	Bead choreographer	The bead darts or jumps on each hit — the lead voice
Overall loudness (RMS)	Slow	Photonic intensity	Brighter, more saturated light as the room gets louder
Spectral centroid (“brightness”)	Medium	Photonic hue	Warm hues for dark/bassy sound, cool hues for bright sound
Silence / low activity	—	All layers	Falls back to a slow generative idle so the piece never dies

Smoothing, latency, and safety of the mapping

Because the terrain lags by up to a second, its features are driven through envelope followers (attack/release smoothing) so the fluid never chases detail it cannot render — it reads the *mood* of the sound, not its waveform. The bead, being fast, is driven by discrete onset events rather than continuous level, so it feels responsive and rhythmic. Note that the base guide's “speed of sound is a latency floor” caution concerns projecting acoustic *force* across a room; here sound is only an *input* captured at one microphone, so that limit does not apply. Keep the fluid within its stable regime — hold duty within the thermal budget and avoid commanding spikes faster than the measured tear-speed from Build Part B, so features shear smoothly instead of splattering.

6. Micro-Theater: Scripted Choreography

Micro-theater treats the stack as a tiny stage. A scripted score sequences the three layers through a short narrative arc; the same scene graph that carries live-sound parameters also carries authored keyframes, so a piece can be fully scripted, fully reactive, or — best — a scripted skeleton whose details the room's sound fills in.

The scene model

A piece is a list of **scenes**, each a set of keyframes across three tracks — a *terrain* track (target heightfields over time), a *bead* track (waypoints and dwell), and a *light* track (color, intensity, labels). Scenes blend with cross-fades; a cue system advances them on a timer, an operator key, or an audio trigger (for example, “hold until the room falls quiet, then begin”). Authoring is the creative instrument, and building this sequencer is the bulk of the software effort.

An example three-beat piece: “Landfall”

- **Beat 1 — Calm sea.** Terrain flat and glassy; the bead drifts low and slow like a distant boat light; cool blue wash. Room sound barely stirs the surface.
- **Beat 2 — The storm.** Low-band energy from the room (or a scripted swell) raises a central peak; mid-band ripples spread to the rim; the bead climbs and circles faster; light shifts to warm amber, intensity tracking loudness. This is where live sound and script interlock.

- **Beat 3 — Landfall and stillness.** The peak subsides into a low ridge that stays put; the bead settles at its crest; light cools and dims to a single held point. A slow generative idle takes over until the next cycle.

7. Parts and Cost Delta

These items are *in addition to or replacing* the base build's magnetic-layer parts; the acoustic layer, sealed fluid cell, and integration items from the core guide are unchanged. Figures are indicative.

Item	Qty	Est. cost	Notes
12 V electromagnets, 25 mm barrel	64	\$140–200	Replaces the 16-coil grid; 8×8 at ~15 mm pitch.
Logic-level MOSFET modules (IRLZ44N)	64	\$60–100	One channel per coil.
PCA9685 16-ch PWM driver boards	4	\$12–20	I²C-chained; solves the 64-channel addressing problem.
Flyback diodes (1N5819)	70	\$8	One across every coil. Still non-negotiable.
12 V 20–30 A power supply	1	\$35–55	Headroom for many simultaneous coils; fuse the rail.
Heatsinks + 12 V fans (active cooling)	—	\$20–35	Coil-bank cooling for sustained gallery runtime.
USB microphone / audio interface	1	\$15–40	Room-sound capture for the analyzer.
Mini LED projector	1	\$50–90	The photonic layer; glossy ferrofluid takes projection well.
USB depth camera (for closed loop)	1	\$60–150	Optional in v1; required for the museum build.
Perfboard, screw terminals, 16 AWG wire	—	\$20–30	Power distribution for 64 channels.
Delta subtotal (art rig, no depth cam)		**≈ \$360–580**	**over base build**

The 8×8 magnetic-layer parts alone add roughly \$180–320 versus the base 4×4; the projector, microphone, and cooling account for the rest. A depth camera adds \$60–150 when you move toward the closed-loop and museum stages.

8. Roadmap: From Gallery Piece to Collaborative Exhibit

Build in stages. Each stage is a working, showable artifact, and each de-risks the next — so the museum-grade collaborative instrument is the destination, not the starting point.

Stage	What you build	What it proves / adds
V1 — Living instrument (this doc)	8×8 rig, room-sound engine, micro-theater sequencer, projector.	The two-speed instrument works; sound → matter/light is expressive and stable. A gallery-ready reactive piece.

V2 — Closed loop	Add the depth camera; register light to the fluid's measured shape.	Light lands on real geometry; the registration problem is solved on the bench. Visual polish jumps.
V3 — Hardened	Ruggedized sealed cell, full active cooling, conservative duty caps, attract/idle loop, graceful bead-loss recovery, tamper-proof enclosure, acoustic guarding.	Survives all-day unattended runtime and a public that touches everything. Turns an artwork into an exhibit.
V4 — Collaborative	Multiple public inputs (knobs/pads or several microphones), one input per terrain feature; ensemble mixing in the engine.	Two or three visitors play together — the low feature count becomes a feature: one person per channel. The queue-forming museum piece.

9. Safety Notes (Deltas)

Heat, scaled up. Sixty-four coils multiply the base build's thermal load. Mount on a non-flammable base, enforce duty and simultaneity caps in firmware, run the fans whenever coils are energized, and power down when unattended until the enclosure and interlocks of V3 exist.

Power. A 20–30 A supply can melt thin wire and start fires. Fuse the coil rail, use heavy gauge for distribution, and verify polarity before first power-up.

Ultrasound. The 40 kHz field is inaudible but intense near the arrays; keep ears and pets away during long runs. Public proximity is a V3 concern — the levitation chamber must be guarded before any exhibit deployment.

Ferrofluid and magnets. Keep the cell sealed; ferrofluid stains permanently and must never touch bare neodymium. Keep strong fields away from pacemakers, cards, and mechanical watches.

10. References

Hirayama, R., Martinez Plasencia, D., Masuda, N., & Subramanian, S. (2019). A volumetric display for visual, tactile and audio presentation using acoustic trapping. Nature, 575, 320–323.

Marzo, A., Barnes, A., & Drinkwater, B. W. (2017). TinyLev: A multi-emitter single-axis acoustic levitator. Review of Scientific Instruments, 88, 085105. Build guide: instructables.com/Acoustic-Levitator (UpnaLab).

Marzo, A., et al. Ultraino: an open phased-array platform for airborne ultrasound; successor software platform OpenMPD. (Per-transducer phase control — the route to a steerable bead and to the array emitting audible sound via amplitude modulation.)

Jordà, S., Geiger, G., Alonso, M., & Kaltenbrunner, M. (2007). The reacTable: exploring the synergy between live music performance and tabletop tangible interfaces. Proc. TEI '07. (Reference for a tangible, collaborative, real-time musical instrument.)

Kodama, S. Protrude, Flow (2001) and Morpho Tower (2006–2007). Ferrofluid sculpture — the artistic precedent for a magnetically animated matter display.

Shi, S., Hübl, M. C., Grosjean, G., Goodrich, C. P., & Waitukaitis, S. (2025). Electrostatics overcome acoustic collapse to assemble, adapt, and activate levitated matter. PNAS, 122(50). (Path to higher bead counts.)

StableLev: Data-Driven Stability Enhancement for Multi-Particle Acoustic Levitation. CHI 2024.