ORGANIC REFLECTIONS: A STOCHASTIC FRAMEWORK FOR REAL-TIME AUDIOVISUAL COMPOSITION ACROSS AR, VR, AND PROJECTION

Deniz CAGLARCAN[1]

[1]**University of California, Santa Barbara**, CA, USA

ABSTRACT

This paper presents a real-time audiovisual composition framework in which sound and image are shaped by the same stochastic decisions, producing structurally coupled output without post-hoc synchronisation. Built entirely within Unreal Engine 5, the system uses a weighted-probability scheduler to simultaneously select from five heterogeneous sound engines, implemented in MetaSounds using analogue-modelled synthesis, granular processing, and sample playback, and route the resulting amplitude envelopes to Niagara GPU particle systems, shaping mesh speed, birth rate, force, and scale in real time. The framework externalises compositional decision-making into a generative system: the composer authors a possibility space of gestural and textural audiovisual behaviours, and the scheduler navigates it in real time, producing structurally guaranteed synchresis through shared kineto-morphological evolution. The system deploys from a single project to three presentation modes, mobile augmented reality, virtual reality on Meta Quest, and projection, preserving the parametric audiovisual bond while reshaping the audience's spatial relationship to the work. The framework is demonstrated through the generative audiovisual work *Organic Reflections*.

Keywords: audiovisual composition, real-time synthesis, parameter mapping, particle systems, Unreal Engine 5, MetaSounds, Niagara, visual music

A video demonstration of Organic Reflections is available at: `https://youtu.be/AJKcg9nryEQ`
An archived audiovisual documentation is available at: `doi.org/10.5281/zenodo.20710728`

1. INTRODUCTION

Audiovisual composition has pursued increasingly tight bonds between sound and image [1]. At one end of the spectrum, composers work iteratively across domains, shaping a sound in response to an image, then reworking the image in light of the revised sound, building synchresis [2] through craft and repetition. Parameter mapping [3, 4] formalises this relationship, establishing explicit correspondences between sonic and visual attributes so that the two media share structural behaviour rather than merely coexisting on a timeline. Co-generation [5] takes the principle further still, deriving both modalities from a single mathematical process so that their coupling is not designed between two outputs but is inherent in a shared source.

This paper presents a real-time audiovisual framework built entirely within Unreal Engine 5 [6] that sits at the intersection of algorithmic composition [7] and parameter mapping. A weighted-probability scheduler, a form of controlled randomness, makes stochastic decisions that simultaneously select from heterogeneous sound engines (analogue-synth models in MetaSounds and sample-based players), assign amplitude envelopes, and route the resulting energy profiles to a Niagara particle system, shaping speed, birth rate, applied force, and size in real time. Because the same amplitude envelope shapes both the sonic output and the particle behaviour, the two domains are structurally coupled rather than independently authored. The framework is demonstrated through the audiovisual work *Organic Reflections*. The contribution is not any single technique, since stochastic scheduling, parameter mapping, and particle visualisation are each established, but their integration into one pipeline in which a single stationary stochastic process drives synthesis, particle behaviour, and multi-platform presentation from one project, so that the coupling is structural rather than retrofitted.

A distinctive feature of the system is its deployment across three presentation modes from a single Unreal Engine project. In augmented reality, viewed through a mobile phone, the particle field is composited into the viewer's physical surroundings via camera passthrough. In virtual reality, experienced through a Meta Quest 2 or 3 headset, the viewer is immersed within the particle environment with spatialised audio. In projection, the particles are cast onto a wall via a projector, transforming the surface into an audiovisual field. Each mode preserves the same generative coupling between sound and image while fundamentally reshaping the audience's spatial and perceptual relationship to the work.

The framework responds to the conference theme of INTERPLAY at multiple levels: the reciprocal interplay between sound engines and visual particle systems, mediated by stochastic decision-making; the interplay between the generative work and the viewer's physical environment, which differs across AR, VR, and projection; and the interplay between a single computational system and the radically different audience experiences it produces across platforms.

The remainder of the paper is organised as follows. Section 2 surveys related work on parameter mapping, au-

Copyright: © 2026. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 Unported License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

diovisual co-generation, and real-time environments. Section 3 details the system architecture, including the sound engine design, envelope and amplitude mapping, particle system configuration, and multi-platform deployment. Section 4 presents the audiovisual output and analyses the perceptual results. Section 5 discusses limitations, future directions, and conclusions.

2. RELATED WORK

Garro [4] provides a systematic framework for connecting sound and moving image by enumerating the phenomenological parameters relevant to each domain and mapping them as desired. On the visual side, he identifies colour (saturation, hue, brightness), shape (size and geometry), surface texture, granularity (clusters, groups, clouds, single objects), position and movement (speed, trajectory, acceleration in 2-D or 3-D space), and surrogacy (recognisability and representation). On the sonic side, he highlights spectrum (frequency, pitch, harmony), amplitude envelope (energy profile), granularity (sound grains, sequences, aggregates), spatial behaviour (trajectory, position, speed), and surrogacy (recognisability and connection to reality). While Garro's taxonomy establishes a rich conceptual vocabulary for audiovisual cross-modal correspondence, it also builds on earlier discussions of mapping as a constitutive layer of expressive system design [3], extending those concerns from electronic instrument design into the audiovisual domain. Callear [8] draws a related parallel between audio grains and visual particles, a correspondence directly relevant to systems coupling audio synthesis to GPU particle generation. However, the question of how to realise these mappings within a single real-time generative system remains open.

Putnam [5] addresses this question at the signal level through the orbit generator, a rational function of Fourier series implemented as a ratio of two inverse discrete Fourier transforms. The output is a sequence of complex numbers: plotting these values in the plane yields a visual trajectory, while extracting the real or imaginary part yields an audio waveform. Because both modalities are projections of one underlying mathematical process, their structural coupling is inherent rather than imposed. However, the orbit generator operates at the level of single complex sequences and does not readily extend to audiovisual scenes involving multiple independent agents, heterogeneous sound sources, or particle-based visual systems.

A number of real-time environments are widely used for audiovisual creation. TouchDesigner [9] offers a node-based visual programming interface with strong GPU compute capabilities and is commonly used for projection and generative installations. Max/MSP with Jitter [10] provides deep audio signal processing alongside real-time video and OpenGL rendering and has a long history in audiovisual performance. Processing and openFrameworks offer code-based approaches with broad community support. Each of these environments, however, typically requires bridging separate subsystems for audio and visual generation, often through OSC, MIDI, or shared-memory protocols. Unreal Engine 5 [6] offers a different proposition: a game-engine environment in which a node-based audio synthesis engine (MetaSounds [11]), a GPU-accelerated particle system (Niagara), a Blueprint visual scripting layer, and a native XR deployment pipeline coexist within a single project. This integration eliminates inter-application communication overhead and allows the same project to be built for screen, AR, VR, and projection without maintaining separate codebases. Prior work has embedded procedural audio within game engines for artistic purposes: Atherton and Wang [12] integrated ChucK in Unity; Hughes and Johnston [13] coupled Unity's GPU particles with modular synthesis; and Hughes et al. [14] argued game engines are optimal for audiovisual art. Lucas et al. [15] recently compared audio engines for XR on Meta Quest 3 in Unity. The present work extends this direction by using Unreal Engine 5 not only for audio synthesis but for tightly coupled audiovisual generation, leveraging MetaSounds and Niagara within a single project.

Audiovisual art has been explored across immersive platforms, VR, AR, and projection mapping, surveyed comprehensively by Turchet et al. [16]. Recent SMC contributions include the restoration of an audiovisual installation in VR [17] and a motion-controlled immersive installation [18]. However, most such works rely on separate toolchains for audio generation, visual rendering, and platform deployment, which can introduce latency between modalities. The framework presented here generates both sound and image within a single engine and deploys to all three platforms from one project, preserving the parametric audiovisual bond at every stage of the pipeline.

The perceptual framework for evaluating the audiovisual output draws on two established concepts. Chion [2] defines synchresis as the spontaneous perceptual bond formed when a sound and a visual event occur simultaneously, even when their causal relationship is implausible. Boucher and Piché [19] extend this with a taxonomy of synchresis descriptors, including direct, gestural, behavioral, and metric synchresis, and introduce the audiovisual object as the smallest unit of audiovisual discourse. In Section 4, these concepts are used as analytic tools for examining the perceptual outcomes of the system, not as the contribution of this paper.

3. SYSTEM ARCHITECTURE

3.1 Overview

The system is built entirely within Unreal Engine 5 and consists of three interconnected layers: a Blueprint-based stochastic scheduler, a set of heterogeneous sound engines implemented in MetaSounds, and multiple Niagara GPU particle systems. At each decision point the scheduler selects sound sources, assigns an amplitude envelope curve, and routes the resulting energy profile to the particle systems, where it controls mesh behaviour in real time. Because the scheduler's decisions simultaneously determine what is heard and what is seen, the audiovisual coupling is intrinsic to the generative process rather than imposed through post-hoc alignment. The complete signal flow is shown in Figure 1.

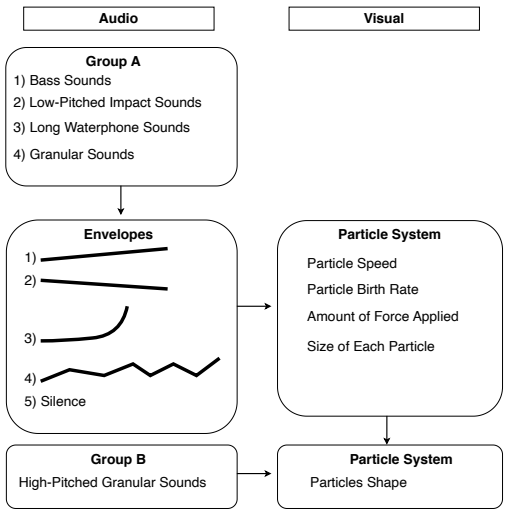


Figure 1. System architecture: the stochastic scheduler selects from five sound engines across two groups and assigns amplitude envelope curves. Amplitude signals drive Niagara particle parameters. The system deploys to AR, VR, and projection from a single UE5 project.

No element of the system is pre-rendered. Each run produces a unique audiovisual sequence: the stochastic scheduler ensures that no two performances follow the same path, while the parametric coupling ensures that whatever path is taken, sound and image remain structurally bonded.

3.2 Sound engine design

The sound layer comprises five distinct engines, four of which are selected by the stochastic scheduler and a fifth that intervenes independently.

Engine 1 — Sub bass. A synthesizer built in MetaSounds using a wavetable derived from analogue Moog synthesizer recordings, shaped through ladder filters and LFOs, to produce low-frequency sustained tones. Map Range nodes scale the oscillator output into the desired frequency and amplitude regions. A tremolo variant of this engine adds iterative amplitude modulation, producing a rhythmically pulsing low-frequency layer. The tremolo rate and depth are parameterised within the MetaSounds patch. When active, the tremolo sub may sound simultaneously with the sustained sub bass, creating interference patterns and slow beating that enrich the low-frequency texture.

Engine 2 — Impact sounds. Twelve pre-recorded impact samples are loaded into a MetaSounds patch. At each trigger event, one sample is selected at random, producing short, punch-like percussive transients that serve musically to initiate or close a phrase, creating gestural counterpoint against the surrounding textural layers.

Engine 3 — Waterphone chords. Pre-recorded waterphone samples are loaded into a MetaSounds patch as a `WaveAsset` array. At each trigger event, a `Random Get` node selects from the array with weighted probabilities and a no-repeat window of 13, preventing immediate repetition. Up to five samples may be triggered simultaneously through parallel delay chains (Figure 2), so that

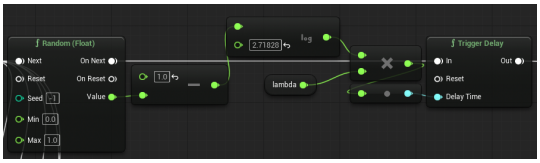


Figure 2. MetaSounds patch for Engine 3 (waterphone chords). Each parallel chain generates exponentially distributed delay times via inverse transform sampling: a uniform random float is log-transformed and scaled by $-\lambda$ to produce the delay interval. Multiple chains run concurrently, enabling up to five simultaneous layers.

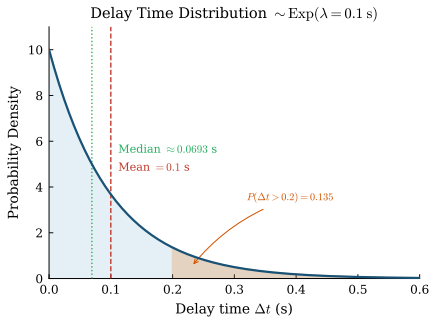


Figure 3. Probability density of the exponential delay time distribution used in Engine 3, with rate parameter λ yielding a mean of 0.1 s. Short delays predominate (median $\approx$ 0.069 s), producing temporal clustering, while longer gaps occur with decreasing probability ($P(\Delta t > 0.2) = 0.135$).

between one and five layers sound at any given moment; silence is also a possible outcome. The layering produces chord-like textures whose density varies stochastically.

The temporal spacing between successive triggers is governed by an exponential distribution, generated within MetaSounds via inverse transform sampling:

$$t_{\text{delay}} = -\frac{\ln(U)}{\lambda} \tag{1}$$

where $U \sim \text{Uniform}(0, 1)$ is drawn by a `Random Float` node and λ is a rate parameter controlling the mean inter-arrival time. In the current configuration, λ is set to produce a mean delay of 0.1 s (Figure 3). The exponential distribution was chosen for its memoryless property: the probability that the next sound event occurs within a given interval is independent of how much time has already elapsed. For the listener, this means that anticipation never consolidates into expectation, the temporal structure remains perpetually open, resisting the formation of periodic patterns that would make the system predictable.

An Envelope Follower (attack time 0.01 s, release time 0.2 s, peak mode) tracks the instantaneous amplitude of the mixed output.

Engine 4 — Granular. Drawing on microsound techniques [20], a hybrid MetaSounds patch combines synthesised grains with processed audio-file fragments. Multiple oscillator chains feed through bandpass filters and LFOs, producing dense micro-textural clusters. The patch includes Map Range nodes that constrain grain density and spectral bandwidth.

Engine 5 — Jitter (interventional). The most complex synthesis patch, built from multiple parallel wavetable oscillator chains with extensive filtering, modulation, and spectral shaping. Unlike the first four engines, the jitter sound is not part of the regular rotation; it has a low probability of being triggered and, when active, plays on top of whatever is currently sounding. Its role is to disrupt: both sonically, by injecting an unstable, spectrally erratic layer, and visually, by deforming the particle field (described in Section 3.4).

The stochastic scheduler selects among Engines 1–4 using weighted random draws. At each scheduling event, one or more of these engines may be activated; silence is also a possible outcome. The scheduling process is stationary and memoryless: the weights are fixed for a given performance rather than evolving over time, and each event is an independent draw not conditioned on previous selections. With the memoryless inter-arrival times of Engine 3 (Equation 1), the system carries no internal state across a run; its behaviour is fully specified by the fixed weights and the per-event draws. The weights were tuned iteratively by ear across multiple development sessions, balancing variety against coherence: too-frequent waterphone chords, for instance, produced a static harmonic field, while over-weighting the granular engine overwhelmed the textural balance. More broadly, the weight configuration is the principal control over the pacing and density of the work: weighting silence and the sub-bass slows the field into sparse, drifting clusters, whereas weighting the impact and granular engines raises event density and gestural turnover. Because the impact and jitter engines map to high-amplitude onsets, their weights most strongly govern gestural continuity against the textural ground. Engine 5 (jitter) operates on an independent low-probability trigger, so its appearances are sparse and unpredictable.

Each engine outputs its signal through a dedicated Envelope Follower, which extracts a continuous amplitude value. These amplitude streams become the control signals for the particle systems.

3.3 Envelope selection and amplitude mapping

Each activated engine follows one of five amplitude envelope curves, selected by a second layer of controlled randomness. The envelopes are pre-authored as curve assets:

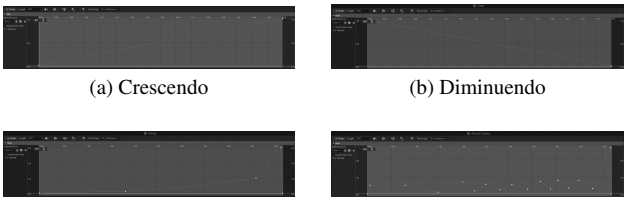
1. *Crescendo* — a gradual increase in amplitude over approximately 5 seconds, with a slow exponential rise that accelerates toward the end.

2. *Diminuendo* — the inverse: a sustained level that decays toward silence.

3. *Delta crescendo* — an envelope that remains at near-zero amplitude for most of its duration, then surges dramatically at the very end, producing a long accumulation of tension followed by a sudden burst of sonic and visual energy.

4. *Crescendo with zigzags* — an irregular, oscillating increase that introduces unpredictable dynamic fluctuation within an overall rising trajectory (5.2 seconds in duration, with amplitude values oscillating between approximately 0.0 and 1.0 throughout the curve).

5. *Silence* — zero amplitude output, producing a visual pause in the paired particle field while other engines may continue sounding.



(a) Crescendo (b) Diminuendo

(c) Delta crescendo (d) Crescendo with zigzags

Figure 4. Four of the five amplitude envelope curves used by the scheduler (the fifth is silence). Each curve shapes both the sonic energy profile and the corresponding particle behaviour.

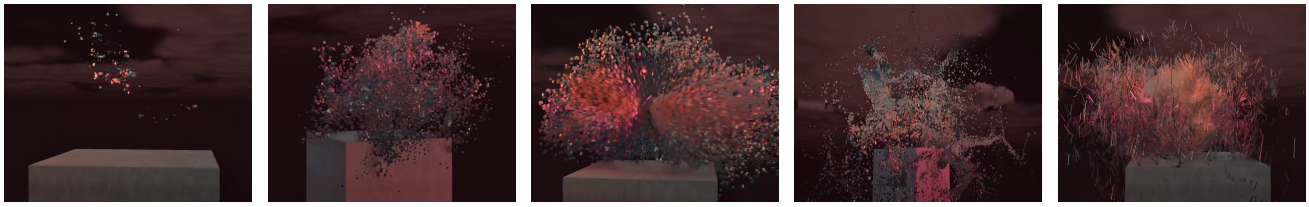
Envelope selection is stochastic: at each scheduling event, a weighted random process assigns one of the five shapes to the active engine. Because the envelope can change independently of the sound source, stochastic shifts in envelope selection directly reshape the particle field even when the same sounds continue to play. A sustained waterphone chord, for instance, may drive a crescendo, a silence, or a zigzag pattern in the visual domain, yielding varied audiovisual phrasing from a limited sonic palette.

The amplitude signal is routed from each engine's Envelope Follower to the Niagara particle systems through Blueprint logic. The raw amplitude is log-scaled (base 10) to compress the dynamic range, ensuring that quiet sounds still produce visible particle activity while loud sounds do not saturate the visual field. The log-scaled value is then passed through `Map Range Clamped` nodes that remap it into the target range for each Niagara user variable (e.g. `User.Envelope1`, `User.curl`), with per-parameter ranges configured to produce the desired visual response.

3.4 Particle system design

The visual layer comprises a Niagara particle system containing multiple emitters running concurrently. Each sound engine has a designated emitter whose behaviour is driven by that engine's amplitude output. A continuous emitter provides the textural ground, responding to the combined energy of all active engines, while a dedicated emitter for the jitter engine handles mesh deformation independently. All emitters use mesh-based rendering rather than sprites, giving the particles volumetric, three-dimensional presence, and spawn from sphere-shaped locations.

The continuous emitter serves as the aggregate textural ground: its spawn rate is the sum of remapped amplitude values from multiple sound engines, so that the overall particle density reflects the combined energy of the active soundscape. Each sound engine's designated emitter shares the same module structure. Spawn rates are remapped from amplitude to a range of 0–300 particles,



(a) Waterphone: sparse floating cubes, slow textural drift.

(b) Sub-bass: drag pulls particles into a dense central cluster.

(c) Impact: high initial amplitude scatters particles outward.

(d) Granular: curl noise drives spiralling, web-like expansion.

(e) Jitter: three-axis mesh deformation produces needle-like forms.

Figure 5. Visual signatures of the five sound engines in *Organic Reflections*. Each engine's amplitude envelope drives different dominant particle parameters, producing perceptually distinct behaviours from the same underlying Niagara system.

meaning emitters produce densely when their associated sound is loud and fall silent when it is not.

Each emitter is configured with a stack of amplitude-driven force and behavioural modules that translate sonic energy into visible motion. Turbulence intensifies with louder sounds, producing swirling, non-repetitive particle trajectories. Directional drift increases with amplitude, adding a sense of wind-like movement. Resistance is mapped inversely to amplitude: louder sounds release particles into freer motion while quieter passages constrain them, a deliberate compositional choice that gives gestural events a sense of physical liberation. Sudden repulsive forces, driven by amplitude peaks, scatter particles explosively outward from the centre. Positional jitter perturbs particle trajectories in proportion to amplitude, and particle colour evolves over each particle's lifetime via a gradient curve, fading from warm tones to transparency as particles age.

When the jitter engine is triggered, it drives a dedicated intervention across all three spatial axes of the particle mesh. The jitter amplitude simultaneously compresses the X and Y axes (scaling from 1.0 down to 0.2) while stretching the Z axis (scaling from 1.0 up to a randomised maximum between 4.0 and 10.0). This produces a dramatic elongation effect: particles that were roughly cubic become needle-like, with the degree of stretch varying unpredictably due to the randomised Z maximum. The sonic instability of the jitter engine, erratic spectral fluctuation and unpredictable amplitude, is mirrored by this visible morphological distortion, reinforcing the audiovisual bond through shared behavioural dynamics.

Although all five engines share the same underlying amplitude-to-particle mechanism, each produces a distinct visual signature through its dominant parameter mappings. The waterphone chords, with their gentle amplitude profiles, create small textural movements, sparse cubic meshes floating slowly through the field (Figure 5a). The sub-bass engine activates Drag most strongly, pulling all particles toward the centre into a dense, compressed cluster (Figure 5b). The low-pitched impact sounds (Engine 2), whose amplitude begins at a high peak, scatter particles explosively outward, rapidly expanding the volume of the field (Figure 5c). The granular engine drives Curl Noise Strength, causing the particle system to grow and curl outward in spiralling tendrils, expanding like a web from the centre (Figure 5d). The jitter engine transforms cubic meshes into needle-like forms, as described above (Figure 5e). The result is that each engine is visually legible: a viewer familiar with the system can identify which sound source is active from the particle behaviour alone.

The overall visual result is a continuously evolving particle field in which the continuous emitter provides a textural ground reflecting aggregate sonic energy, the amplitude-driven emitters articulate individual sound events, and the jitter emitter periodically disrupts the field's equilibrium through three-axis mesh deformation synchronised with its sonic counterpart.

3.5 Multi-platform deployment

The generative core of the system, the scheduler, sound engines, amplitude analysis, and particle to parameter coupling, is identical across all deployment targets. What changes between platforms is the rendering output, the spatial audio configuration, and the viewer's physical relationship to the work.

In augmented reality, the system runs on a mobile phone using Unreal Engine's AR pipeline. The Niagara particle field is composited into the camera passthrough feed, placing the generative visuals within the viewer's physical surroundings. The viewer can move around the particle field, observing it from different angles as the phone's position and orientation update the camera perspective in real time. Audio is delivered through the phone's speakers or headphones in stereo.

In virtual reality, the system runs natively on Meta Quest 2 or 3. The viewer is immersed within the particle environment, with the field surrounding them in three-dimensional space. Head tracking allows free exploration of the visual field, and audio is spatialised through the headset so that the directional position of sound sources corresponds to the perceived location of particle clusters. The immersive enclosure fundamentally changes the perceptual experience: rather than observing the particle field as an object in the room (as in AR), the viewer inhabits it.

In projection, the Niagara output is rendered and projected onto physical walls, transforming the architectural surface into an audiovisual field (Figure 6). Audio is delivered through external speakers positioned in the room.

In all three modes, the parametric bond between amplitude and particle behaviour is preserved exactly. The audience experiences the same generative audiovisual organism; what differs is whether that organism inhabits their physical room (AR), surrounds them (VR), or transforms

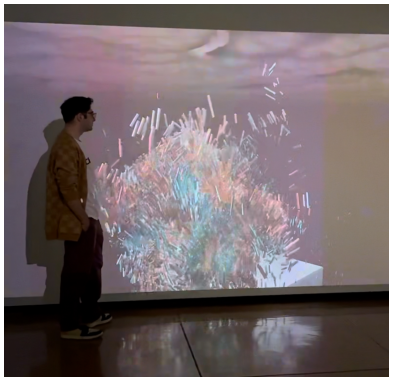


Figure 6. *Organic Reflections* in projection mode: the particle field is projected onto a physical wall, transforming the architectural surface into an audiovisual field.

the wall into a responsive audiovisual surface (projection).

4. RESULTS AND AUDIOVISUAL ANALYSIS

4.1 Compositional design rationale

The stochastic architecture described in Section 3 was not designed to produce arbitrary audiovisual output. Rather, it was composed in the same way one would compose a fixed electroacoustic work [21]: by defining the sonic materials, their possible energy profiles, their roles in the musical texture, and the conditions under which they appear. The difference is that instead of fixing these decisions to a timeline, they are encoded as weighted probabilities within a possibility space. The scheduler navigates this space in real time, making decisions that the composer would otherwise make manually during the compositional process.

This approach draws on the distinction between gesture and texture as articulated by Smalley [22]. In fixed-media composition, the composer places gestural events, characterised by steep energy change, clear onset, and a sense of directed motion, against a textural ground defined by internal repetition, stability, and self-containment. In *Organic Reflections*, these roles are distributed across the system's components. The burst emitters, driven by impulsive sound events from the granular and sub-bass engines, produce gestural audiovisual objects: particles that appear suddenly, accelerate, and dissipate, articulating energy onsets that the viewer perceives as directional and caused. The continuous Chaos emitters, sustained by the waterphone chords and the sub-bass drone, provide a textural audiovisual ground: a slow, internally active particle field that maintains itself and recycles energy without demanding focal attention. The jitter engine, when triggered, introduces a gestural disruption into the textural field, an intervention whose rarity and morphological intensity make it perceptually salient.

Because both the sonic and visual domains are driven by the same amplitude signal, they necessarily share what Boucher and Piché [19] term kineto-morphological behaviour: the same speed, the same directional dynamics, the same energy trajectory. When a crescendo envelope drives both a rising waterphone texture and an accelerating particle birth rate, the two modalities progress morphologically at the same perceptual speed with coherent shared directional dynamics. This is precisely the condition under which Boucher and Piché identify gestural synchresis, a perceptual bond formed not through causal plausibility but through shared energy-motion trajectory. The synchresis in *Organic Reflections* is therefore not designed between two finished outputs; it is a structural consequence of the shared amplitude source.

4.2 Exhibition and audience reception

Organic Reflections was presented at *Shadows*, a contemporary audiovisual immersive art festival at the University of Nevada, Reno (Sheppard Contemporary), on 22–26 January 2024, where the author was composer and artistic director. It was shown in two of its three modes at once: a viewer in an Oculus Quest 2 headset explored the piece in VR while the same generative output was projected onto the gallery wall, so that the headset user and the surrounding audience shared one live performance through different media.

The following observations are informal rather than a controlled study. Viewers in VR walked into the particle field to stand inside it and stepped back to view it from any angle, as one might navigate a game environment. Audiences repeatedly described the work as organic or biological, reading the sound-directed visual gestures as the behaviour of a living thing, consistent with the title. Despite the heterogeneity of the five engines, listeners reported a strong perceptual bond between sound and image across all of them, consistent with the shared kineto-morphological trajectory that couples the two domains. A formal perceptual evaluation across the AR, VR, and projection modes remains future work.

5. CONCLUSION

This paper presented a real-time audiovisual framework in Unreal Engine 5 coupling MetaSounds engines to Niagara particle systems through a stochastic scheduler, producing emergent gesture–texture interplay across AR, VR, and projection from a single project. The system is not a general-purpose audiovisual tool but a compositional instrument: the possibility space is authored for a specific aesthetic, and the stochastic scheduler is a performer of that space. Although the runtime mapping flows one way, from audio amplitude to visual parameters, the work is an audiovisual composition in the visual-music tradition, not a sonification [23]: both domains are authored, and the composer fixes which sonic events drive which visual parameters. This single direction is a compositional decision, not a constraint; beginning from the visual domain, or coupling the two bidirectionally, is an equally valid alternative rather than a missing capability. The work is conceived as a continuous generative installation rather than a fixed-duration concert piece: designed to be entered and left, it runs indefinitely without exact repetition and so does not impose a global narrative arc. The weighted scheduler instead encodes the composer's gestural and textural language so that local phrasing approximates hand composi-

tion, while the memoryless turnover of selections sustains variety and resists stagnation. An optional higher-level layer modulating the weights over time, shaping macro-level tension and sectionality, remains future work.

The title *Organic Reflections* derives from this property: sound and image share a common amplitude source, so the particle field responds to sonic energy with the immediacy of a living system, each modality reflecting the other in real time. The framework responds to the notion of interplay at multiple levels: interplay between sound engines and particle systems, mediated by controlled randomness; interplay between compositional craft and stochastic emergence, where the artist composes a possibility space rather than a fixed sequence; and interplay between a single generative system and the radically different perceptual experiences it produces across platforms.

Acknowledgments

This work draws on research conducted as part of the author's doctoral studies at the University of California, Santa Barbara.

6. REFERENCES

[1] K. Brougher, J. Strick, A. Wiseman, and J. Zilczer, Eds., *Visual Music: Synaesthesia in Art and Music Since 1900*. London and New York: Thames & Hudson, 2005.

[2] M. Chion, *Audio-Vision: Sound on Screen*. Columbia University Press, 1994, original work published 1990.

[3] A. Hunt, M. M. Wanderley, and M. Paradis, "The importance of parameter mapping in electronic instrument design," *Journal of New Music Research*, vol. 32, no. 4, pp. 429–440, 2003.

[4] D. Garro, "Connected media, connected idioms: The relationship between video and electroacoustic music from a composer's perspective," in *Sound and Image: Aesthetics and Practices*, A. Knight-Hill, Ed. Routledge, 2020, pp. 1–12.

[5] L. J. Putnam, "A mathematical model of planar curves for audiovisual synthesis," Doctoral dissertation, University of California, Santa Barbara, 2011.

[6] Epic Games, "Unreal Engine 5 [computer software]," https://www.unrealengine.com/, 2024.

[7] I. Xenakis, *Formalized Music: Thought and Mathematics in Composition*, revised ed. Stuyvesant, NY: Pendragon Press, 1992.

[8] S. Callear, "Audiovisual particles: Parameter mapping as a framework for audiovisual composition," Ph.D. dissertation, Bath Spa University, 2012.

[9] Derivative, "TouchDesigner [computer software]," https://derivative.ca/, 2024.

[10] Cycling '74, "Max/MSP [computer software]," https://cycling74.com/, 2024.

[11] Epic Games, "MetaSounds documentation," https://dev.epicgames.com/documentation/en-us/unreal-engine/metasounds-in-unreal-engine, 2024.

[12] J. Atherton and G. Wang, "Chunity: Integrated audiovisual programming in Unity," in *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME 2018)*, Blacksburg, Virginia, 2018, pp. 102–107.

[13] M. Hughes and A. Johnston, "URack: Audio-visual composition and performance using Unity and VCV Rack," in *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME 2020)*, Birmingham, UK, 2020.

[14] M. Hughes, J. A. Garcia, F. Wilcox, R. Sazdov, A. Johnston, and A. Bluff, "Immerse: Game engines for audio-visual art in the future of ubiquitous mixed reality," in *Proceedings of the International Conference on Live Interfaces (ICLI 2020)*, Trondheim, Norway, 2020.

[15] P. P. Lucas, S. Fasciani, and K. Glette, "Csound vs. ChucK: Sound generation for XR multi-agent audio systems in the Meta Quest 3 using the Unity game engine," in *Proceedings of the Sound and Music Computing Conference (SMC 2024)*, Porto, Portugal, 2024.

[16] L. Turchet, R. Hamilton, and A. Çamcı, "Music in extended realities," *IEEE Access*, vol. 9, pp. 15 810–15 832, 2021.

[17] M. Russo *et al.*, "Revitalizing multimedia art: Restoration and re-activation of an audiovisual installation in virtual reality," in *Proceedings of the Sound and Music Computing Conference (SMC 2024)*, Porto, Portugal, 2024.

[18] A. E. Memis and T. Kelkar, "IEMI: An immersive, evolving motion-controlled installation," in *Proceedings of the Sound and Music Computing Conference (SMC 2024)*, Porto, Portugal, 2024.

[19] M. Boucher and J. Piché, "Sound/image relations in videomusic: A typological proposition," in *Sound and Image: Aesthetics and Practices*, A. Knight-Hill, Ed. Routledge, 2020, pp. 13–29.

[20] C. Roads, *Microsound*. Cambridge, MA: MIT Press, 2001.

[21] M. McDonnell, "Constructing visual music images with electroacoustic music concepts," in *Sound and Image: Aesthetics and Practices*, A. Knight-Hill, Ed. Routledge, 2020.

[22] D. Smalley, "Spectromorphology: Explaining soundshapes," *Organised Sound*, vol. 2, no. 2, pp. 107–126, 1997.

[23] D. Çağlarcan, "From sound to sight: Translating electroacoustic music principles into visual music composition," PhD dissertation, University of California, Santa Barbara, 2025, https://escholarship.org/uc/item/0sd6p43k.