Performative transactions: worlding compositional ecosystems

Lukawski, A.

Chapter 3. Artists as System-Builders

Encoding Process, Enforcing Trust

In Chapter 1, we explored the figure of the artist as operator: a practitioner engaging creatively with the outputs of computational tools such as machine learning models, generative systems, and agentic AI. These connectionist technologies introduced forms of automation, unpredictability, and emergence—qualities that artists learn to navigate through iterative exploration. Chapter 2 expanded this view by presenting the artist as a curator of both human and non-human agencies: someone who arranges and configures heterogeneous elements with their own autonomous agency into posthuman agential assemblages. This chapter focusses on artists who are not only users of pre-existing systems and arrangers of relations between various agencies, but also design autonomous systems from the ground up. These are artists as system-builders—those who construct processes that can run independently of their creator.

Part 1 of this chapter—titled *Composing Autonomous Systems*—first discusses the difference between aleatoric and stochastic composition, contrasting John Cage's use of chance operations—rooted in ethical and spiritual openness—with Iannis Xenakis's mathematically structured approach to probabilistic control. These positions introduce two divergent understandings of how composers relate to unpredictability and system behaviour. Part 1 then tracks the views of the musical work as a field of possibilities and affordances, and culminates in today's machine-learning milieu where composers become system-builders responsible for curating data, writing code, and balancing constraint with indeterminacy to shape emergent outcomes. The chapter then expands into the broader field of generative composition, where authorship and musical identity are displaced from singular works to the design of systems that produce variable outputs. Importantly, generative music is presented as a third alternative to live and recorded modes of music dissemination. The text further explores how compositional agency is increasingly understood as the crafting of structured spaces of potential—open systems governed by interaction affordances, codes, and dynamic rules. Crucially, it reintroduces an older but powerful concept—the machine.

To build a machine is to construct a system that functions reliably and autonomously within predefined parameters. The machine, in this sense, is a closed system designed to execute particular operations with minimal ambiguity. Many artists working in the tradition of algorithmic composition, cybernetics, or symbolic AI embrace this machinic logic. From stochastic music to generative scores, these practices aim to encode procedures so they can be executed consistently and precisely, even if they involve randomness as a variable within a tightly defined system. In such cases, the artist seeks reliability—designing a structure whose behaviour can be trusted precisely because it is calculable.

Yet contemporary artistic practice increasingly scales beyond any single machine or studio. Composers collaborate across geographic locations, delegate tasks to remote performers and AI services, or release executable works into online environments where they evolve without direct oversight. Once the locus of execution becomes distributed, predictable behaviour can no longer be guaranteed by deterministic design alone. A new set of concerns comes into focus: How can one verify that a remote instance of the process really is the process? Who owns or inherits its outcomes? How can contributors be sure the rules they agreed on will be honoured?

When deterministic control gives way to distributed or opaque behaviour, trust becomes the mechanism that supplements reliability. In the context of autonomous and distributed systems, we can distinguish:

- **Trust in process**: confidence that an opaque or adaptive system will act within an expected range, even if its internal workings remain inaccessible or stochastic. Niklas Luhmann—sociologist and founder of social systems theory—framed this as the reduction of environmental complexity to manageable horizons of expectation (Luhmann, 1979/1973, 25-27), emphasising that actors proceed as if the future was sufficiently knowable to risk continued interaction.

- **Trust in record**: confidence that the fact of what was executed—authorship, parameters, sequence—has been faithfully recorded and cannot be altered retroactively. As Diego Gambetta—sociologist noted for his work on trust, cooperation, and organised crime—notes "if evidence could solve the problem of trust, then trust would not be a problem at all". Cooperation depends on verifiable evidence that past actions occurred as claimed. Without such assurance, strategic partners cannot evaluate reputation or enforce commitments (Gambetta, 1988, 233).

Both forms are pertinent when composers release autonomous or generative works into open, networked environments. Process-trust relates to aesthetic behaviour (will the piece "perform" acceptably when re-instantiated?), whereas record-trust underpins attribution, provenance, and any contractual or economic obligations that may follow from subsequent iterations or sales.

For these reasons, Part 2 of this chapter—titled *Blockchain For The Arts*—shifts its focus from composing as building on the level of algorithmic processes that execute the artistic outcome of a system, to building on the level of algorithmic conditions: conditions that decide whether the underlying algorithmic processes can be performed at all, in which circumstances, and based on what evidence, the actant intending to start such a process is entitled to do so. It is here that the chapter discusses the possible role of blockchain technology in the arts—not only as a financial technology, but as a system for algorithmically enforcing trust.

Where traditional machines aim for reliability through deterministic design, blockchain-based systems enforce agreement through cryptographic verification, decentralised consensus, and smart contracts. In other words, they create machinic structures for managing trust, enabling a transformation of the artist's role from the artist as curator of agencies to artist as a builder of formal systems on the level of executable agential procedures. For artists working across networks, across institutions, or with autonomous agents, blockchain provides a programmable infrastructure for defining how participants interact, how permissions are enforced, and how provenance is secured and can be traced. It enables compositional work not only with media or operations, but with relationships, responsibilities, and constraints.

This chapter argues that the system-building turn in contemporary art—particularly in algorithmic, generative, and decentralised practices—demands new frameworks for distinguishing between reliability and trust. Artists build machines when they seek reliable automation, but they also look for ways of confirming reliability when engaging with systems and in settings where behaviour and the final outcome cannot be fully anticipated. Blockchain, as will be presented, uniquely allows for the construction of machinic environments in which trust itself can be encoded and executed on the compositional level.

Within the broader context of this doctoral dissertation, this chapter highlights a significant gap: at present, there exists no general compositional infrastructure that enables multiple actants to collaboratively compose on the blockchain in a seamless, accessible manner. While blockchain technology holds promise as a means of enabling such compositional interaction, its current implementations, which will be discussed, remain limited to a small number of specialised applications—each offering only partial and highly specific affordances for artistic participation. Accordingly, the purpose of this chapter's second part is twofold: first, to provide a critical overview of the current state-of-the-art of blockchain technologies for the arts and for music; and second, to lay the technical conceptual groundwork for the introduction of the concept of Performative Transactions—the primary contribution of this dissertation defined in Chapter 6—a blockchain-based compositional framework that enables diverse actants to transparently and modularly build on each other's contributions.

Part 1—Composing Autonomous Systems

From Chance to Probability

In 1951, composer John Cage famously turned to the *I Ching*, the ancient Chinese divination text known as the "Book of Changes", to compose *Music of Changes*. By employing the *I Ching*, Cage relinquished direct compositional control, instead allowing chance operations—determined by coin tosses corresponding to the text's hexagrams—to dictate the musical material. This approach introduced a radically new understanding of artistic agency: rather than meticulously crafting each detail, Cage allowed an external, non-human system to guide his creative decisions, effectively surrendering traditional notions of personal authorship. As Cage put it: "Most people who believe that I'm interested in chance don't realize that I use chance as a discipline. \[...] My choices consist in choosing what questions to ask" (Kostelanetz, 2003, 70). His embrace of chance thus foregrounded unpredictability, introducing indeterminacy as a fundamental aspect of the compositional process.

Traditionally, aleatoric practices in music composition, including Cage's, have been seen as challenges to Platonist ideals—attacks upon the notion of fixed, ideal structures underlying musical form. In *Formalized Music* (1963), composer Iannis Xenakis explicitly critiqued deterministic causality, rooted historically in Plato's assertion from the Timaeus that "it is impossible for anything to come into being without cause" (Xenakis, 2001/1963, 1). Yet Cage's turn toward the *I Ching* and Zen Buddhism imbued his work with other metaphysical and transcendental assumptions. Cage's indeterminacy was not a nihilistic retreat from structure, but an ethical and spiritual exercise in openness. "Instead of talking about [non-intentionality], I wanted to do it", Cage explained, "and that would be done by making the music nonintentional, and starting from an empty mind" (Kostelanetz, 2003, 35). This surrender to the process of chance—as something meaningful beyond the self—suggests a belief in latent forms waiting to emerge: "When I ask the *I Ching* a question, I say, 'What do you have to say about this?' and then I just listen to what it says" (70). Cage's intent behind his indeterminacy in music was one that sees chance not merely as the absence of authorial intent, but as a vehicle through which pre-existing potentialities—latent forms—manifest themselves.

This act of relinquishing compositional responsibility was precisely what Iannis Xenakis challenged in his formulation of stochastic music. Xenakis instead saw chance as something that could—and should—be systematically rationalised and controlled: "It is only recently that knowledge has been able to penetrate chance and has discovered how to separate its degrees—in other words to rationalise it progressively, without, however, succeeding in a definitive and total explanation of the problem of 'pure chance'" (Xenakis, 2001/1963, 4). Where Cage allowed aleatory processes to introduce genuine indeterminacy without pre-calculated probabilistic controls, Xenakis insisted that chance itself needed rigorous mathematical framing: "Chance needs to be calculated (…) In any case—to play with sounds like dice—what a truly simplistic activity! But once one has emerged from this primary field of chance worthless to a musician, the calculation of the aleatory, that is to say stochastics, guarantees first that in a region of precise definition slips will not be made, and then furnishes a powerful method of reasoning and enrichment of sonic processes" (Xenakis, 2001/1963, 38). Thus, Xenakis's stochastic music diverges fundamentally from Cage's indeterminacy. While Cage relinquished direct control as a spiritual-philosophical act, Xenakis's stochastic approach deliberately retained compositional authority by encoding probability distributions within carefully constructed mathematical frameworks.

Automatism and Affordances of a Musical System

Pierre Boulez, reflecting on the issues of "automatism and decision", recognised how "[a]lthough the computer, more than any other type of equipment, symbolises the urgency of this question, it did not create the problem. The question of the relationship between chance and determination in composers' intentions, and the means used to realise those intentions in actual works, is already familiar" (Boulez, 2018b). In a later lecture, Boulez develops this further, describing compositional practice not as a process of selection but composing entire fields of possibilities, within which the performer makes the final choices at the moment of execution. "I compose all of the possibilities", he writes, "and I leave it to the performer to choose the possibility at the very moment of performance" (Boulez, 2018a, 263). This mode of composition displaces authorship from the production of a single work toward the construction of a structured space of potential outcomes.

Such an understanding of the system as the site of artistic identity also reframes composition as an act of construction rather than inscription. For composer Helmut Lachenmann, to compose is not to express but to intervene—to "take apart" existing materials and construct new relations from their interconnections. "Composition is by no means a 'putting together'", he writes, "but rather a 'taking apart' and more: a confrontation with the interconnections and necessities of the musical substance" (Lachenmann, [1979] 1996, as cited in de Assis, 2018, 227). In this spirit, Lachenmann describes composing as the building of a new instrument—not necessarily a physical tool, but a virtual structure that captures forces and channels them into sound. For him, this "instrument" is a diagrammatic field. It is the complete collection of sounds and materials used in a given piece, crucially in the moment of performance. It is well-thought before the notation process starts, but it remains audible throughout. In the context of systems aesthetics and computational creation, such structures are not preparatory—they are the artwork. As Paulo de Assis puts it, the act of composition becomes one of mapping possibilities and constructing the conditions for their emergence: not to say something, but to do something (de Assis, 2018, 230). In this view, the artwork is not a singular object or fixed sequence of sounds, but a dynamic, metastable system—a machine capable of generating its own actualisations.

Generative systems, in this expanded sense, may take the form of a score—such as composer Witold Lutosławski's composition *Venetian Games*, which uses principles of limited aleatoricism—a live-coded algorithm, or a sound-producing machine. What unites these diverse media is their systemic logic: each defines a structured field of potentiality from which variable outputs may emerge. In all of these cases, the identity of the artwork does not reside in any single instantiation or performance. It resides in the system itself—in its capacity to generate, to transform, and to persist across iterations. This reconceptualisation of the artwork as a generative system opens directly onto the broader domain of generative aesthetics. Whether embedded in scores, or algorithms, such systems foreground process over product, modulation over inscription, and structural coherence over surface identity. They mark a shift from composing works to composing spaces of possible outputs that hold the identity of the artwork.

Similarly, composer Artemi-Maria Gioti argues that in the context of interactivity, machine listening, and generative processes, the contemporary musical work increasingly operates as a space of possibilities rather than a fixed structure (Gioti, 2021, 64). In this framework, musical identity is defined not by concrete sound structures but by the interaction affordances and action spaces embedded in the system. “The musical text”, she writes, “is understood as having an evocative, rather than a directive function. It delineates a space of action for the musicians to explore” (6). These spaces, governed by a system's affordances, can include composed, improvised, or algorithmically generated material—coexisting within a continuum of control and emergence. Such systems-oriented perspective leads to a redefinition of authorship and identity: the system becomes the artwork. Different instances of performance are not instantiations of a fixed original, but variations within an aesthetic topology defined by the system’s affordances. “Work identity”, Gioti writes, “is constituted of interaction affordances and spaces of sonic possibilities, rather than concrete structures of sounds” (9). It is precisely here that the notion of a machine re-enters. In Gilbert Simondon’s discussion of technical objects, machines evolve by reducing their internal margin of indeterminacy through successive codings of operation. Yet, paradoxically, Simondon clarifies that the progressive perfecting of a machine, its increasing technicity, does not correspond to greater automatism. Rather, it depends upon maintaining a certain margin of indeterminacy within its operations: “A purely automatic machine completely closed in on itself in a predetermined way of operating would only be capable of yielding perfunctory results. The machine endowed with a high degree of technicity is an open machine” (Simondon, 2017, 17). In Simondon’s view, such open machines inherently rely upon human participation, positioning the human as a permanent organiser and interpreter—guiding interactions between technical entities (17). Thus, the machine defines the operational space within which processes unfold with controlled variation.

If artists take on the roles of system-builders—designers of machinic infrastructures that also learn, adapt, and generate—then compositional responsibility extends once again: it is both about crafting outputs and encoding rules, but also about designing the conditions under which learning and emergence unfold. In this space, the challenge is to navigate both sides of the generative systems, to constrain and guide them in ways that preserve creative agency and aesthetic coherence—without foreclosing the very indeterminacy that makes them generative in the first place.

Artists Building Generative Systems

Generative music, as composer and producer Brian Eno articulates, is characterised by systems or "machines" designed to produce musical experiences by combining pre-defined materials and processes in ways the artist does not explicitly determine. As he writes: “One of my long-term interests has been the invention of ‘machines’ and ‘systems’ that could produce musical and visual experiences. Most often these ‘machines’ were more conceptual than physical; the point of them was to make music with materials and processes I specified, but in combinations and interactions that I did not” (Eno, 1996, 330). Central to Eno’s conception of generative music is the recognition of its inherent uniqueness and its novel position relative to traditional music forms. Reflecting on his experiences with generative systems such as the Koan software—generative system used by Eno for his first generative music pieces—he observes: “Until 100 years ago every musical event was unique: music was ephemeral and unrepeatable, and even classical scoring couldn’t guarantee precise duplication. Then came the gramophone record, which captured particular performances and made it possible to hear them identically over and over again. But Koan and other recent experiments like it are the beginning of something new” (Eno, 1996, 332).

Eno identifies generative music as distinct from—and complementary to—the two established forms of musical experience, live and recorded music. In his view, generative music bridges and extends the virtues of both: “From now on there are three alternatives: live music, recorded music and generative music. Generative music enjoys some of the benefits of both its ancestors. Like live music, it is always different. Like recorded music, it is free of time-and-place limitations—you can hear it when you want and where you want. And it confers one of the other great advantages of the recorded form: it can be composed empirically. By this I mean that you can hear it as you work it out—it doesn’t suffer from the long feedback loop characteristic of scored-and-performed music” (Eno, 1996, 332). This observation is further underscored by Eno’s reflection on a remark by Edgar Wind, who, in his 1963 Reith Lectures, compared listening to recorded music to experiencing a "superior kind of musical clock”. Eno extends this perspective to imagine generative music as a future standard, noting: “I too think it’s possible that our grandchildren will look at us in wonder and say, ‘You mean you used to listen to exactly the same thing over and over again?’” (Eno 1996, 332).

Despite the frequent association of generative art with contemporary digital practices, Philip Galanter clarifies that the field is not inherently tied to computer technology. He argues: "Because contemporary generative art is so very often computer based many assume it is a subset of computer art. (…) generative art preceded computer art, and in fact is as old as art itself." (Galanter, 2003, 15). A frequently presented example is the ancient Aeolian harp. First described by Athanasius Kircher in his book *Musurgia Universalis* (1650), it is a music instrument played by wind, rather than by humans. Indeed, Galanter positions generative art within a historical lineage far preceding any recent artistic movement, asserting explicitly that "generative art (…) precedes modernism, post-modernism, and just about any other 'ism' on record" (19) and explaining that "Equally important is the virtual certainty that new forms of generative art will come after the computer as well. Nanotechnology, genetic engineering, robotics, and other technologies will no doubt offer generative artists some wonderful opportunities"(Galanter, 2003, 15). He further broadens this claim, highlighting that "if we accept this paradigm, that generative art is defined by the use of systems, and that systems can be best understood in the context of complexity theory, we are led to an unusually broad and inclusive understanding of what generative art really is." (12).

Galanter offers a foundational definition of generative art that has become widely cited in the field: "Generative art refers to any art practice where the artist uses a system, such as a set of natural language rules, a computer program, a machine, or other procedural invention, which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art" (2003, 4). Central to this definition is the notion that the artist relinquishes partial or total control to the system: "The key element in generative art is then the system to which the artist cedes partial or total subsequent control" (4). He adds several clarifying observations. First, the term refers strictly to how the work is made, not why it is made, nor what it might express. Second, generative art is not tethered to any specific medium or level of technological sophistication. Finally, Galanter emphasises that such a system must be "well defined and self-contained enough to operate autonomously" (4). In short, generative art is not an aesthetic style but a procedural approach—"a way to create art rather than an art style" (15)—in which the outcome emerges from a system with built-in autonomy and variability. As a consequence, Galanter insists that generative art is best understood as a subset of art defined by its capacity to produce multiple outcomes: "The word 'generative' simply directs attention to a subset of art, a subset where potentially multiple results can be produced by using some kind of generating system" (4). The existence of variation across iterations is not merely a byproduct, but a defining characteristic. This attention to multiplicity also expands the role of the artist beyond conventional intentions. As Galanter poetically suggests, "The generative artist can remind us that the universe itself is a generative system. And through generative art we can regain our sense of place and participation in that universe" (19). Generative methods, in this broader context, align artistic practice with the principles of emergence, complexity, and systemic thinking.

Among the most illustrative historical examples of composers incorporating expert systems within their compositional workflows are Iannis Xenakis's *GENDY*, George Lewis's *Voyager*, and David Cope's *Experiments in Musical Intelligence (EMI)*. Xenakis's *GENDY3*, developed in the early 1990s, implemented his theory of stochastic synthesis—a process in which waveform generation is governed by statistical functions, producing continuously evolving sonic material without relying on traditional oscillators or sampled waveforms (Luque, 2009). While the stochastic synthesis engine acted as a sound generator rather than a standalone composition, it exemplified Xenakis's broader conceptual ambition of encoding compositional logic into autonomous algorithmic processes (Xenakis, 2001/1963). By contrast, George Lewis's *Voyager* (1987) was conceived not merely as a tool but as the composition itself: a real-time interactive musical system that engages with human performers as an autonomous improvising partner. Lewis explicitly frames *Voyager* as a nonhierarchical, multi-agent system that enacts its own musical logic, positioning the computer not as an instrument but as a collaborator capable of contributing to the unfolding musical form (Lewis, 2000). David Cope's *EMI* (begun in the 1980s) represents yet another model: a symbolic AI system that analyses stylistic patterns in existing compositions and uses them to generate new works in the same style. In contrast to Lewis's improvisational openness, *EMI* foregrounds stylistic replication, raising questions about authorship, creativity, and the boundaries of generative automation (Cope, 2001). Together, these three projects exemplify distinct approaches to expert systems in music—as generative engines, as improvisational agents, and as stylistic emulators—each contributing to an expanded view of composition as the design of autonomous, rule-based, and often performative systems.

Composers as Programmers

Trevor Wishart—electroacoustic composer and improvising vocal performer—writing in 1985, anticipated a future in which the computer would become a "truly generalised tool for sonic art" by enabling composers to define and articulate *operational fields*—rule-based groupings of musical parameters that can interact dynamically with higher-level behaviours or performative inputs (Wishart, 1996, 329). These operational fields mark a crucial conceptual shift: rather than composing fixed sequences of musical material, the composer designs systems that organise potentialities—configurable textures, articulation types, and timbral or morphological transformations—governed by abstract but pliable rules. In this model, composition becomes the design of generative grammars capable of continuous variation. Around the same time, in IRCAM, Miller Puckette—computer-music pioneer, software inventor, and academic—began his work on Max—a visual programming language for music and multimedia, enabling the development of various digital modular "instruments" for live-electronic performances and processual musical systems. Developed in 1996 by James McCartney, SuperCollider—an environment and audio programming language for real-time synthesis and algorithmic composition—offered composers another powerful framework for building their own instruments, systems, and compositional processes, enabling a shift from writing music to designing the conditions under which digital music processes can be shaped.

While today, in the age of ubiquitous Internet access, this might seem trivial, at the time these software environments, though enabling the creation of generative systems, were largely limited to the local machines on which they were installed. Given the constraints of computing power, bandwidth, network availability in the mid-1990s, and dependence on the software, a generative composition could not yet be easily distributed as a dynamic system across devices—in a sense it remained bound to the specific hardware that executed it until it was deliberately transferred, with all of its required software, to another computer. As Thor Magnusson—researcher, composer, and music technologist—points out, Brian Eno's *Generative Music 1*, released in 1996 on floppy disk and built on the Koan software was "one of the earliest examples of distributable generative music" (2019, 198), offering not a fixed audio file but a procedural system capable of generating unique musical outputs at every playback. This innovation helped shift the focus of composition from fixed media to executable logic—musical processes encoded in systems rather than scores or recordings. Eno's work inspired a wave of artists and developers, including Antoine Schmitt and Vincent Epplay's *The Infinite CD* (1999), John Eacott's *The Morpheus CD-ROM* (2001), and later collaborative works like *SameSameButDifferent* (2003–2008), each of which explored novel distribution formats for generative music systems (Magnusson 2019, 198–199).

However, the lack of standardisation in distribution formats has remained a persistent issue. As Magnusson observes, attempts to develop stable delivery mechanisms—from CD-ROMs to embedded hardware like Spyros Polychronopoulos's *Sound Object* (2016)—have been met with ongoing fragmentation, partly due to the diverse programming environments and devices involved (2019, 199). The shift to browser-based tools like the Web Audio API represents a promising new phase in this development, enabling generative music to be composed and performed within standard web environments using JavaScript (199–200). In parallel with these technological changes, a deeper transformation of the composer's role has taken place. The generative artist becomes a programmer of processes—an author of algorithmic environments capable of producing aesthetic output in their absence. As Magnusson writes, this transformation marks the rise of "composers as programmers", evident in the work of figures like Nick Collins, Bob Sturm, and François Pachet (2019, 201–202), where "experimental musical instruments and interfaces are often intended to be presented as artworks themselves" (Giannoutakis, 2024, 12).

As artists gain technical fluency, they are not only using tools but designing them—developing compositional environments, interfaces, and data pipelines that allow for bespoke creative practices. The Future Art Ecosystems 4 have tracked this shift closely, describing how artists now operate as system-builders: creators of integrated technological processes that blur the boundary between tool and artwork. "Some artists", they note, "look outward, building tools into more complex, operative systems for research or production. These artists craft tools, craft with tools, or craft tools into bigger systems" (Serpentine Arts Technologies, 2024, 102). Rather than merely producing objects, artists working with advanced technologies frequently develop systems with potential impact beyond the art world—sometimes innovating in fields such as interface design, artificial intelligence, or decentralised networks (Serpentine Arts Technologies, 2020, 26, 64–65). The metaphor proposed by Takashi Kudo (Tokyo-based Global Brand Director for art collective teamLab)—"It's like you are a painter, but you also have to invent paint" (Serpentine Arts Technologies, 2020, 8)—aptly captures the dual role of the artist as both creator and technologist. These roles converge in the figure of the composer-programmer, who composes not only musical structures but also the technological and computational infrastructures through which musical experience is produced. This shift is especially pronounced in contexts where artists fine-tune or reconfigure machine learning models, as a form of creative authorship—what the Future Art Ecosystems 4 report calls "model-making as meaning-making" (2024, 100–101). In this view, the generative field expands again: no longer ends with the generation of musical forms but includes the generation of the tools and systems by which those forms become possible.

Part 2—Blockchain for the Arts

What is Blockchain

If the first part of this chapter explored the artistic turn of some composers in the twentieth century towards a design of autonomously working artistic systems from the ground up, the second part turns to a more recent kind of machinic infrastructure that is gradually more interesting for artists—blockchain.

At its core, "blockchain is a way to trust a network of computers run by strangers—so you don't have to trust the individual people" (ECAD Labs, n.d.). This deceptively simple definition gets to the heart of blockchain's conceptual innovation. Instead of relying on a single source of truth—such as a website located on a single server—blockchain protocols enable confirming the identity of digital objects and agents through code. All the computers in the network—called nodes—run the same software, store the same history, and apply the same rules. If one computer lies or makes a mistake, the algorithms automatically ignore it. Through cryptographic verification, each participant can independently confirm that data hasn't been tampered with. No central authority is required. The result is a distributed system in which consensus between the nodes (that some piece of information is considered true by the majority) is enforced by code, not by institutional oversight or any central authority. Importantly for this thesis, blockchain introduces, among various new solutions, a way of distributing digital art in which trust is not assumed or negotiated—but formally encoded, a feature crucial when questions of ownership and scarcity are explored in digital contexts, and when various actants collaborate remotely.

This architecture of trust was invented to solve a very specific technical problem in digital commerce: the double-spending problem. When something exists only as digital information, it can in principle be copied endlessly. Unlike a physical coin or banknote, which can only be handed from one person to another, a digital token can be duplicated and spent more than once. This undermines the value and integrity of digital money. Until recently, the only way to prevent this was to use centralised intermediaries—like banks or payment processors—to track ownership and verify each transaction. But this reintroduces a structure of dependence: users must trust these intermediaries not to lie, fail, or extract value.

In 2008, a white paper by the anonymous figure Satoshi Nakamoto proposed a radically different solution: to use a peer-to-peer network to collectively verify and record every transaction—publicly, permanently, and without any central control. "The problem", as Nakamoto writes in the paper, "is that the payee can't verify that one of the owners didn't double-spend the coin" (Nakamoto, 2008, 2). The solution is to make every transaction part of a shared public history—a timestamped chain of verified data blocks, each sealed with cryptographic proof. As Nakamoto explains: "We propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work" (Nakamoto, 2008, 1).

To understand how this system works, we need to begin with one of its core building blocks: a hash. The hashing algorithm takes any piece of information—whether a sentence, a file, or a full transaction record—and transforms it into a short, fixed-length string of characters[4]. The specific kind of hashing used in Bitcoin and many other blockchains is called SHA-256, a secure cryptographic algorithm designed to be fast to compute, but impossible to reverse. It is part of the SHA-2 family, developed by the U.S. National Security Agency (NSA) and published by the National Institute of Standards and Technology (NIST) in 2001. SHA-256 takes input data of any length and produces a fixed-size, 256-bit (32-byte) output known as a hash or digest. What makes it powerful is that even a tiny change in the input creates a completely different hash, making it highly sensitive to input changes. For example, changing a single letter in a sentence will produce a radically different string of characters. This makes it very easy to detect if something has been altered. A hash acts as a digital fingerprint for the input data: even a tiny change in the input will result in a completely different hash. The process is one-way as it is computationally infeasible to reverse the hash to retrieve the original input, which is crucial for security applications like blockchain, but also for everyday online password storage, and digital signatures in various applications besides the blockchain. This means you can easily turn data into a hash, but you can't go backwards from the hash to guess the original data. This one-way nature is essential: it allows the system to verify that something hasn't changed without needing to see or reveal the original content. In this way, to check that some

[4] For example: 8c872aa8ab6955cb91897efba8630850dd2b2ce918c8f3c21d2ce86881bc9f87.

data is the same as the original takes only to compare two hashes, without the need to inspect the hashed data itself.

Now that we understand hashing, we can look at how it's used to validate blocks. A block is a collection of data—such as a text or file—that we want to encode with a special kind of hash. But not just any hash will do. In systems like Bitcoin, a block is only considered valid if its hash meets a particular requirement: for instance, the hash must begin with four zeros. Since hash outputs are unpredictable, there's no way to simply choose the right input to get a valid hash. Instead, what the system does is repeatedly modify the block's contents by changing a small, meaningless number called a nonce. Each time the nonce is changed and added to the data that we want to hash with a valid hash (again: starting with four zeros), the hash is recalculated. Most of the time, the resulting hash is invalid (does not start with the four zeros). But eventually, by trying enough different nonces, the system will stumble upon one that produces a hash with the required number of leading zeros[5]. When such a hash is found, we say that the block has been *signed*. This process—called mining—is like solving a puzzle: the solution proves that someone has spent real computational effort to find a valid hash. Once found, anyone can verify the result by checking the hash themselves. What's important is that while verifying a valid hash is quick, finding it is deliberately hard. This built-in difficulty is what gave early blockchain systems their security and trustworthiness.

Once we understand that a block contains some data—along with a nonce chosen to produce a valid hash—we can begin to see how these blocks are connected into a larger system. This is where the idea of the chain comes in. Each time a new block is created, it doesn't just contain its own data and a valid nonce. It also includes the hash of the previous block as part of its own data. This means that the hash for any given block depends not only on its own contents, but also on the hash of the block that came before it. That block, in turn, contains the hash of the block before it, and so on, forming a continuous chain that links all blocks together. The result is a chain of blocks—a blockchain—in which each block secures the one before it, and is itself secured by the one after. What makes this structure powerful is how it protects against tampering. If someone tries to change the contents of a block from earlier in the chain—even by a single character—the hash of that block will change. But because

[5] For example: 00001fcbc5a02912336fac9a3875049673f43127d1d804cde8078409795a1f8e.

that hash is now embedded in the next block, the next block's data will no longer be valid. Its hash will be incorrect, and so will the hash of every block that follows. To cover their tracks, the attacker would have to recalculate the valid nonce and hash for every block from the point of the change to the present, solving a new computational puzzle each time. In a large blockchain, this becomes computationally infeasible, especially since new blocks are constantly being added by others.

What makes this structure especially powerful is how it behaves in a distributed network. In a blockchain system, many different computers store the same copy of the entire blockchain. Each node has its own copy, but all are expected to agree on the same version of history. When we compare that modified version of the blockchain to the versions stored on other nodes, the inconsistency becomes immediately visible. All it takes is to compare the hashes at each block to see whether a particular copy is in agreement—or in consensus—with the rest of the network. If it's not, the altered version is rejected. The chain with the most valid work behind it—meaning the one that required the most computation and remains unbroken—is considered authoritative. In this way, blockchain doesn't rely on any single trusted source. Instead, it uses redundancy, verification, and shared computation to maintain a collectively agreed history.

So far, we've been referring to each block as containing some unspecified "data". In the original blockchain—Bitcoin—this data is a list of financial transactions. Each transaction recorded the transfer of digital and fungible (meaning: exchangeable 1:1) tokens from one address to another. The blockchain's role is to act as a public ledger—a chronological and verifiable record of who sent tokens to whom, and when. Each block contains multiple transactions, and once a block is mined and added to the chain, those transactions become part of the permanent history. This makes it nearly impossible to alter records retroactively without detection. What Bitcoin invented, then, is not just a new form of money, but a new infrastructure for digital accounting: a distributed system for collectively agreeing on the state of a shared ledger, without relying on any central institution to maintain it.

To interact with a blockchain, users need a digital wallet—a software tool that generates a public address (used like an account number) and a private key (used like a password) which are also cryptographic hashes themselves. These keys define ownership: whoever holds the private key controls the assets associated with its public address. Users on the blockchain network can take different roles. Full nodes store the entire blockchain history and validate new blocks; mining nodes (in proof-of-work systems) compete to solve cryptographic puzzles (like finding the hashes with four zeros at the beginning) and add blocks; and lightweight nodes (or light clients) use simplified methods to verify transactions without storing the full chain. To agree on which block gets added next, blockchain networks use consensus mechanisms. The original system, proof-of-work, requires nodes to perform energy-intensive calculations—a process that secures the network but consumes vast resources. A more recent alternative, proof-of-stake, selects validators based on how many tokens they hold and are willing to "stake", offering similar security with dramatically lower energy use. While Bitcoin still uses proof-of-work, many newer blockchain systems—including Ethereum since its 2022 upgrade—have adopted proof-of-stake, reflecting a broader shift toward more energy-efficient consensus models.

Machines that Keep the Score

Blockchain architectures took a major leap forward in 2014, when Canadian computer programmer Vitalik Buterin proposed a new kind of blockchain platform: Ethereum. In his white paper, Buterin argued that blockchain could support more than just financial transactions. The crucial change he introduced was that whereas Bitcoin stored only a list of financial transactions and data required for its technical operations, Ethereum allowed that data to include any kind of executable program. By embedding a fully Turing-complete scripting language into the blockchain, Ethereum made it possible to write, deploy, and run arbitrary code directly on the blockchain. These programs became known as smart contracts (Buterin, 2014). What makes smart contracts radically different from traditional code is the environment in which they operate. They don't run on a single server or app, but inside a global virtual computer: the Ethereum Virtual Machine (EVM). While each computer in the Ethereum network runs the same software independently, together they simulate a shared computational environment. Every time someone interacts with a smart contract, all full nodes verify the code and reach agreement on its outcome. This means that the contract's execution (a piece of code running) is not only decentralised, but publicly verifiable and persistently stored on the blockchain.

As Akira Summers—software developer and lecturer in digital media technologies—puts it, smart contracts are to blockchain what vending machines are to shops: they automate transactions by enforcing rules with perfect consistency (2022, 119). In simple terms: a smart contract is a computer program “installed” on the virtual machine consisting of many computers, which always has a condition encoded that specifies when it should execute its code (de Assis & Łukawski, 2024, Glossary). Such condition could be financial—in this case the code executes when some public address receives a payment, or it could be non-financial—such as a threshold number of users participating, external data like weather reports, or even artistic criteria—such as a particular input or signal from a sensor, performer, or AI system. Crucially, these contracts do not require human intermediaries to function. Anyone with the Internet access and a digital wallet can run any smart contract on the virtual machine, if they can meet the conditions of that smart contract. When the conditions are met, the code runs. Like a vending machine, the contract does not discriminate between users or interpret intent—it simply executes. This introduces a new kind of infrastructural trust: not based on institutions or reputation, but on code that is publicly auditable and resistant to tampering. In this sense, smart contracts function as small, composable units of logic that can be assembled into broader systems—enabling, among other things Non-Fungible Tokens (NFTs), Decentralised Applications (dApps) (such as art marketplaces), and Decentralised Autonomous Organisations (DAOs).

Non Fungible Tokens (NFTs)—a Non-Fungible Token (NFT) is a blockchain-based mechanism for automatically tracking the identity of a digital object. Normally, digital files—images, sounds, texts—can be copied indefinitely, making it difficult to establish scarcity of such objects and attribute to them such properties as uniqueness, ownership, and provenance (de Assis & Łukawski, 2024, Glossary). This is a variant of the double-spending problem explained earlier: if a digital asset can be duplicated without constraint, its value and authorship become unstable. NFTs offer a solution by recording, on-chain, a verifiable reference to a single version of a digital asset, or in some cases even that asset itself as part of the block’s metadata on chain. Each NFT is defined by a unique identifier embedded in a smart contract, ensuring that even if the associated file is widely distributed, the token representing it remains singular. In this way, NFTs enable digital artworks to behave more like singular artefacts than infinitely replicable files, allowing a simulation of so called digital scarcity.

Decentralised Applications—a Decentralised Application, or dApp, is a software application that interacts with smart contracts on a blockchain. While smart contracts themselves operate on the backend—executing logic and storing data on-chain—a dApp typically includes a user interface that allows people to interact with these contracts more intuitively, often through a web browser or a mobile app (Summers, 2022, 124). Unlike traditional apps, which depend on central servers, dApps rely on decentralised infrastructure: their core logic runs not on a private server but on the blockchain itself. In this way, dApps inherit the properties of the blockchain they run on—transparency, immutability, and censorship resistance—while enabling users to engage with smart contracts through familiar and accessible user interfaces. A common example of a decentralised application is a digital art marketplace, such as *objkt.com* or *teia.art* on the Tezos blockchain, or *zora.co* on Ethereum. These platforms allow artists to mint, list, and sell blockchain-based artworks directly, using smart contracts to govern ownership, provenance, and exchange.

An Example: Hic Et Nunc

A particularly significant example of a decentralised application in the context of the arts was Hic et Nunc, an art marketplace for trading NFTs launched in March 2021 by Brazilian developer Rafael Lima. As Diane Drubay—museum & digital innovation strategist, Web3 consultant, and environmental visual artist—recounts (2024, 98), the platform quickly distinguished itself from the dominant NFT marketplaces of the time—many of which were characterised by high transaction fees, selective gatekeeping, and an emphasis on speculative value. By contrast, Hic et Nunc positioned itself as an open, low-cost, and experimental alternative. Built on the Tezos blockchain—chosen for its energy-efficient proof-of-stake model—Hic et Nunc enabled artists to mint and sell NFTs at a fraction of the cost of Ethereum-based platforms, significantly lowering the barrier to entry for creators worldwide. What made the platform particularly notable was not only its affordability, but its minimal design, lack of curation, and open-source ethos. Artists could mint digital artworks within seconds using a simple interface, and without needing approval from gatekeepers. This structural openness allowed a wide range of creators—from Brazil to Malaysia—to participate on equal footing. According to Drubay, Hic et Nunc's emergence catalysed a moment of cultural reorientation: artists began to experiment with blockchain as a space of shared activity, where new forms of community, value, and visibility could emerge.

When the original platform was abruptly discontinued by its founder in November 2021 (as its frontend depended on a self-hosted website), its community responded by reconstituting the project. Thanks to the decentralised architecture of the blockchain—where the NFTs created through the Hic Et Nunc platform were stored on the Tezos blockchain and the files of the artworks on IPFS (InterPlanetary File System, a peer-to-peer file system)—the community was able to preserve access to the works and launch mirror sites (alternative frontends referencing the same smart contracts and assets). One of these mirrors, *teia.art*, evolved into a Decentralised Autonomous Organisation (DAO): a collectively governed structure in which decisions about the platform's future are made directly by its contributors. As Paulo de Assis and I noted elsewhere (de Assis & Łukawski, 2024, 104), this transition marked an instrumental moment in the history of blockchain-based artistic platforms, demonstrating that ownership and continuity could be sustained even in the absence of a central figure or company.

Ruth Catlow (curator, writer and blockchain researcher), and Penny Rafferty (art-tech researcher) (2022, 28), as well as Catherine Mulligan (blockchain researcher and tech-policy advisor) (2024, 108-110) observe that Decentralised Autonomous Organisations represent a novel organisational form that challenges centralised models of coordination and governance. Because their rules are encoded in smart contracts, DAOs enable decentralised communities to manage membership, voting, and collective funds without requiring external intermediaries. Participation is often token-based, meaning that holding a specific token can grant decision-making rights or access to shared resources. This opens up possibilities for experimenting with alternative institutional models that are more transparent, participatory, and adaptable than conventional structures. Mulligan (2024, 109) suggests that DAOs may offer a basis for new kinds of cultural institutions—less hierarchical, more inclusive, and potentially more responsive to the needs of decentralised communities. For this reason, blockchain can be understood not just as a technical tool, but as a space for organisational and societal experimentation. As Primavera De Filippi—legal scholar, blockchain governance expert, and creative technologist—puts it, "We don't otherwise have a chance to experiment with governance anywhere, so it's a very valuable space for exploring and testing new governance structures"

(Serpentine Arts Technologies, 2022, 55).

Yet, as Claudio Tessone—complexity scientist and blockchain researcher—argues, current blockchain applications often replicate the patterns of inequality found in traditional socioeconomic systems (2024, 88). Markets built on token economies tend to reward those who already hold value, following the familiar logic that "the growth of an entity within a system is proportional to its existing size" (88). Without careful intervention, decentralised systems risk reproducing the same concentration of power and wealth that they claim to resist. In this sense, the future of blockchain depends less on its technical architecture than on the imaginative and political choices made by its communities. As Tessone writes, "Blockchain technology, in its nascent stages, has yet to significantly deviate from [existing] paths. […] To achieve this, a concerted effort in education, research, and practical experimentation is essential" (94). Some proposals—such as Mark Alizart's notion of cryptocommunism (2020)—have begun to explore how blockchain could support alternative value systems entirely. Whether or not such visions are realised, what matters is that blockchain infrastructures remain open to experimentation—not only with money and media, but with the very forms of collective life they make possible.

Inte Gloerich—critical media & technology researcher and cultural programmer—argues that it is more productive to view blockchain not in terms of static technical features, but as a heterogeneous assemblage encompassing everything from smart contracts and digital wallets to institutional actors, cultural imaginaries, and regulatory frameworks (2025, 27). These assemblages condition what kinds of relations and behaviours can emerge and trust is encoded into the system's architecture. Blockchain networks can then be understood as machinic assemblages: dynamic configurations of code, protocols, interfaces, infrastructures, narratives, and agents—both human and non-human. Future Art Ecosystems 3 suggests that artists working with decentralised technologies do not merely create artworks, but design systems in which art can be produced, circulated, and interacted with differently (2021, 24). For artists working with multi-agent systems, decentralised authorship, or distributed networks of value, blockchain offers more than a tool for tokenising media—it is a medium for designing the very rules by which participation, provenance, and power are constituted.

Generative Music NFTs

One of the most innovative intersections of blockchain technology and artistic practice to date can be found in generative art NFTs. Generative art NFTs can be defined as 'unique digital art pieces created through coded algorithms and stored as NFTs on the blockchain' (Stoykov, 2023, as emphasised in Łukawski, 2024a, 37-38). The same algorithm can generate many unique NFTs by using random functions or dynamic inputs. For example, the colours, shapes, or musical structure might change depending on the time of day, the wallet address of the collector, or external data like weather conditions.

The degree to which the generative process itself is integrated into the blockchain can vary widely. In the most autonomous configurations, not only is the NFT minted on-chain, but the algorithm that generates the artwork is also encoded directly into the smart contract. This means that the entire lifecycle of the artwork—from generation to tokenisation—is managed without relying on external systems. Rhapsody Labs (RLXYZ 2022, and as discussed in Łukawski, 2024a, 37-38) has proposed a useful framework for understanding this, introducing the notion of on-chain purity. This framework evaluates how permanent, transparent, and self-contained a generative artwork is when considered over long timescales—centuries, even. The fewer external dependencies an artwork has (such as off-chain storage or mutable metadata), the "purer" it is considered. At the highest level of purity, all components—the artwork, smart contract, generation function, metadata, and unique identifiers—are encoded entirely on the blockchain, ensuring that the artwork remains independently accessible and verifiable far into the future.

To understand the practical applications of generative art NFTs, it's essential to examine the platforms that facilitate their creation and distribution. Two prominent platforms in this space are *Art Blocks* (n.d.) and *fxhash* (n.d.), each offering unique approaches to generative art on the blockchain. Art Blocks is an Ethereum-based platform that enables artists to publish algorithmic systems, which collectors can then mint into unique, on-demand generative artworks. When a collector initiates a mint, the platform uses a combination of the artist's code and a random seed to generate a one-of-a-kind piece, ensuring that each generated artwork is distinct. Fxhash is an open generative art platform built on the Tezos blockchain and currently also integrated with Ethereum and Ethereum Base. It allows artists to upload generative scripts, which are then used to automatically generate unique new artworks as NFTs upon collection by the users of the platform (typically, when the user buys a new iteration of the artwork with cryptocurrencies). The platform emphasises accessibility and openness, enabling any artist to publish projects without a formal curation process (although many strict rules on what is allowed within the code of the artwork are listed in the platform’s documentation). Fxhash also allows artists to define parameters that collectors can adjust during the minting process (through their *fx(params)* functionality), adding an interactive dimension to the artwork. Both Art Blocks and fxhash exemplify how blockchain technology can be leveraged to create, distribute, and collect generative art in innovative ways. By providing platforms that integrate the generative process with the minting of NFTs, they offer artists new mediums for creativity and collectors new forms of engagement with digital art.

Long-Form Generative Art

In his essay “The Rise of Long-Form Generative Art”, Tyler Hobbs (2021) delineates a significant shift in generative art practices facilitated by blockchain technology. Traditionally, generative artists would produce numerous outputs from an algorithm and selectively curate only a few of them for public presentation. This "short-form" approach allowed artists to showcase only the most compelling iterations, effectively concealing less successful outputs (Hobbs, 2021). The advent of platforms like Art Blocks introduced a "long-form" paradigm, wherein artists deploy a generative script to the blockchain, specifying a fixed number of iterations—often ranging from 500 to 1,000—that collectors can mint directly. Each minted piece is generated in real-time, with neither the artist nor the collector knowing the outcome beforehand. This process eliminates the possibility of post-generation curation, compelling artists to ensure that every potential output meets a high standard of quality. Hobbs emphasises that this model requires artists to design algorithms capable of producing consistently compelling results across a broad output space. The responsibility shifts from selecting the best outputs to crafting a generative system where every piece holds artistic merit. This transparency fosters a deeper engagement between collectors and the generative process, as they gain insight into the algorithm's full range and variability. In contrast, short-form generative art relies heavily on the artist's discretion to present a curated selection, potentially obscuring the algorithm's limitations. Long-form generative art, by exposing the entirety of the algorithm's output, democratises the creative process and challenges artists to develop more robust and versatile generative systems (Hobbs, 2021).

While the field of generative visual art NFTs has seen rapid innovation, particularly through platforms like Art Blocks and fxhash, similar developments in music and sound-based practices have been comparatively limited. As Kosmas Giannoutakis—composer, media artist, computer-musician researcher—notes, visual artists were among the first to explore the creative potential of blockchain technologies, with musicians only more recently beginning to experiment with these tools in meaningful ways (2024, 61). Martin Zeilinger (2024, 57) similarly observes that although blockchain and AI technologies have already led to a range of sophisticated experiments, "creative experimentation and theoretical reflection appears to have taken place [...] with a primary focus on visual, non-time-based artforms". The affordances of blockchain for encoding logic, temporality, and interactivity remain underexplored in musical contexts, despite their potential to support new models of compositional authorship, distribution, and performance.

Some notable projects have explored how blockchain-based systems can be used for generative audio. Among the most known is the work of digital artist Deafbeef, whose fully on-chain series—*Synth Poems*, *Angular*, *Transmission*, and others—combine generative sound synthesis with minimal code to produce audiovisual artworks directly from smart contracts (0xDEAFBEEF, n.d.). These works are executable algorithms that generate sound deterministically on-chain, exemplifying what a high degree of blockchain-native authorship might look like in the sonic domain. Other artists, such as Memo Akten and James Paterson (a.k.a. Presstube), have explored the hybrid terrain of generative audiovisual NFTs with projects like *A Strange Loop*, which investigates recurrence and evolution in algorithmic patterns. The artist known as poperbu has also created generative sound-based NFT collections including *BITXO* and *Amnèsia*, each experimenting with sound as a time-based, procedurally generated medium embedded in blockchain infrastructures.

Three Experimental Prototypes: *Canons, Tone Row, Subtraction* (2024)

As part of my artistic research into the compositional affordances of blockchain infrastructures, I developed three experimental generative NFT projects on the fxhash platform in 2024: *Canons* (2024b), *Tone Row* (2024c), and *Subtraction* (2024d). Each of these projects investigates a different aspect of generativity, authorship, and decentralised infrastructure, treating the NFT not only as a distribution mechanism, but as a medium for algorithmic music-making. These works were intended as exploratory prototypes, artistic experiments designed to test the creative and technical limits of existing blockchain platforms. By treating generative algorithms, minting processes, and smart contracts as compositional materials, each project models a distinct approach to encoding musical logic, interaction, and variability. The first of these, *Canons* (2024), is a limited-edition generative NFT that algorithmically composes short two-part canons in the style of 16th-century counterpoint. Developed in consultation with historical flute player and artistic researcher Jonty Coy, the algorithm draws on the rules laid out by Zarlino, Bathe, and Morley for improvising canons and ornamenting lines through diminution. Each mint generates a new, unique canon, rendered both as audible playback and as a visual score, with MIDI export functionality. The project was conceived as an academic response to Coy's presentation at the Academy of Creative and Performing Arts (Leiden University), and functions as both a scholarly commentary and a compositional system.

The next two projects, *Tone Row* and *Subtraction*, were commissioned by composer Zach Dawson as interpretations of his experimental text scores from the book *ImageAudio: A Post-Digital Event Score Anthology* (2024). Tone Row is an audiovisual generative music NFT that produces a new 1–4 minute piece on each mint, combining a large collection of pre-recorded synthesizer samples with algorithmic transpositions through a tone.js sampler device, form-generation, and audio-reactive visuals. The algorithm randomly selects between 12-tone equal temperament and microtonal tuning systems, resulting in a broad sonic spectrum across editions and is capable of generating stylistically coherent though widely varied outputs. Since its release on May 14, 2024, over 140 diverse iterations have been minted. This work illustrates how a generative NFT can function not as a fixed output but as an expert system capable of composing with variability—one that requires multiple versions to be fully grasped.

The *Subtraction* generative music NFT, by contrast, reimagines the minting process itself as an act of distributed composition. Rather than generating full works at once, each NFT in the *Subtraction* series represents a single chord, whose parameters—duration, pitch range, density—are selected by the minter (the person initiating the generation process becoming the owner of the generated NFT). These chords are then aggregated into a continuous composition rendered on an external website (2024e), which queries the fxhash API to assemble the complete composition timeline in real time and allows its playback in various configurations (from start to end, shuffled, or allowing the user to play each chord separately to examine who is its current owner and the original minter). In this model, the artwork exists not as a single object but as an emergent assemblage of contributions—an ongoing, collaborative system that evolves with each mint. The artwork's minting phase was algorithmically set to one month and resulted in a short composition consisting of 37 chords. As I wrote in a related article for fxhash (2024f), *Subtraction* enacts a kind of "posthumanist musical assemblage", where the roles of composer, performer, and listener are redistributed across a network of human and non-human actors.

Together, these three projects explore how generative music can be operationalised through blockchain not only as a form of aesthetic output, but as a medium for rethinking authorship, temporality, and collaborative agency. They gesture toward new forms of compositional infrastructure in which code, protocol, and procedurality become central to the musical experience.

More Long-Form Generative Music?

A useful overview of early generative audio-based NFT projects can be found in an article by Erik West (2022), published on the fxhash platform. West traces a range of experiments at the intersection of sound, code, and blockchain, acknowledging the pioneering nature of these efforts (such as as a collaboration between kaigani and Dadabots, ByteBeats, or Rami Awad's generative piece *Basil*) while noting that "generative music is still a newborn compared to the relative infancy of the entire NFT space". He frames the problem not as a lack of artistic potential, but of infrastructural maturity and cultural adoption. As West puts it, "I dream of hearing Brian Eno's next release in a truly generative way, instead of by streaming or buying a single version, with a set beginning and ending".

One important reason for this limited engagement might lie in the technical divergence between tools traditionally used by algorithmic sound artists and those required for generative art development on blockchain platforms. Many sound artists, especially those trained in conservatory or experimental music contexts, work with specialised audio programming environments such as Max/MSP, Pure Data (Pd), or SuperCollider. These platforms are optimised for real-time audio synthesis, granular control over signal flow, and complex musical structures, but they are not natively compatible with the web-based programming environments required for generative art NFT platforms. By contrast, platforms like fxhash and Art Blocks rely on technologies rooted in web development—particularly HTML, CSS, JavaScript, and libraries such as p5.js and tone.js—which are more commonly used in visual or browser-based creative coding communities. For many algorithmic musicians, transitioning to these environments can present a significant learning curve, particularly when it comes to sound synthesis in the browser, which lacks the low-level control and audio fidelity of desktop-native tools. That said, bridging solutions are beginning to emerge. Cycling '74, the developer of Max/MSP, recently released the RNBO framework (Cycling '74, n.d.), which allows Max patchers to be exported as web-compatible audio applications. Similarly, developer Sébastien Piquemal has created WebPd (Piquemal, S., n.d.), a library that enables Pure Data patches to run in web browsers using JavaScript. Notably, the generative artist poperbu has already alluded to these developments in his work: in the NFT piece *bang.pd* (poperbu, 2023), he incorporates the visual motif of a Pure Data patcher as a conceptual gesture. The artwork seems to imagine a future in which Pd-native audio systems can be deployed and collected through blockchain platforms like fxhash. This gesture suggests perhaps not only a technical desire, but a broader cultural statement that tools familiar to experimental musicians could become interoperable with decentralised, generative infrastructures for distribution and collection.

One reason why generative art platforms have so far been more widely adopted by visual artists than by musicians or sound artists may lie in how each field relates to the ontology of the artwork. As Domenico Quaranta—art critic, curator and researcher—notes, visual art has traditionally operated within an autographic framework—where each work is a unique object—whereas every notated music is in principle allographic: what matters is the underlying score, code, or system, not any single performance or output (2021, 85-87). A famous example is that traditionally a copy of the Mona Lisa painting is not considered the same as the original Mona Lisa, but various distinct performances of the Beethoven's symphony are all considered instantiations of the original piece. Blockchain infrastructures, particularly NFTs, introduce mechanisms of artificial authenticity that align more easily with autographic thinking—supporting scarcity, uniqueness, and the object-hood of digital works. For many sound artists, however, whose focus remains on reproducible processes or live performative systems, this shift toward collectible artefacts may feel less intuitive. As a result, the affordances of blockchain for experimental music remain underexplored—not because they are lacking, but because they require a different conceptual framing of what constitutes a "work". Nevertheless, the relatively new phenomenon of generative are NFTs has a strong potential to attract time-based experimental artists, exactly because it enables a medium in which it is the unfolding of a system in time that perpetuates the experience of the artwork.

Whereas traditionally visual artists might more naturally consider a single generative output to constitute a unique object, composers and sound artists are traditionally used to be thinking in terms of time-based performative re-execution and iterative variation. Blockchain-based systems, by contrast, offer a technical mechanism to individuate and register each execution as a distinct artefact. This opens up new possibilities, where composition's multiple performances or parameter-based instantiations can be each minted as unique NFTs, establishing a direct provenance for what might otherwise remain ephemeral or repeatable. A parallel can be drawn to recent developments in dance. During the 2022 symposium "Dance and the Blockchain: Commodification and Ownership of Embodied Creativity in the Crypto Space", held at Ghent University and convened by Jorge Poveda Yáñez (Poveda Yáñez, 2022), researchers explored how motion-tracking systems could capture the gestures of dancers and encode them as NFTs (Poveda Yánez & Davies). The goal was not simply to document performance, but to inscribe bodily movement as reusable digital assets—capable of being invoked, recombined, or reinterpreted in other choreographic or virtual contexts.

This gap between visual and musical experimentation on the blockchain may also reflect a lack of a distinct conceptual anchor within music's Web3[6] culture. As Marcus

[6] "Web3 represents the next phase of the internet, focusing on decentralised networks and blockchain technology. It aims to create a user-owned internet, where individuals have control over their data, identities, and transactions. Web3 leverages cryptocurrencies, decentralized finance (DeFi), and non-fungible tokens (NFTs) to facilitate peer-to-peer interactions without central authorities." (de Assis & Łukawski, 2024, Glossary).

O'Dair—researcher of innovation and musician—observes, while the aesthetic of Web3 visual art—particularly the iconic profile picture collections like *CryptoPunks* and *Bored Ape Yacht Club*—has become instantly recognisable, "Web3 music doesn't sound any different to Web1 or Web2 music" (2024, 78). These visual projects introduced not only a distinctive generative format, but also a collective identity and recognisable function: the profile picture (PFP). They effectively became symbols of participation and status in the Web3 ecosystem. Music, by contrast, has yet to produce a comparable conceptual format—one that would embody its generative, distributed, or performative qualities while also being legible as a new form of cultural value. Without a widely accepted conceptual form that links blockchain-native affordances with musical or sonic expression, Web3 music risks remaining tethered to distribution models inherited from earlier platforms. In visual art, NFTs facilitated the creation of recognisable categories—such as generative long-form series or fully on-chain artworks—that became integral to the experience and valuation of the medium. In music, however, NFT-based projects often reproduce familiar formats: static audio files, streaming rights, or token-gated access to pre-existing compositions. As a result, the experiential and compositional affordances of blockchain—such as real-time algorithmic generation, interactive participation, or verifiable versioning—are underutilised.

Critics argue that the decentralization ideal of blockchain can mask new forms of corporate or monopolistic control. As observed by the media artist Hito Steyerl, rather than empowering creators, blockchain platforms are seen as *"onboarding tools"* that ultimately serve Big Tech conglomerates' interests (Brown, 2023). Furthermore, the libertarian and techno-utopian rhetoric of blockchain clashes with many artists' community values. The promise of *democratizing* art through decentralization is met with skepticism: while crypto advocates speak of radical inclusivity, critics note that the NFT boom often reproduced the same exclusivity and inequities of the traditional art market (Casemajor, 2022).

Some art theorists argue that tokenization hasn't truly empowered artists but instead introduced new intermediaries (e.g. crypto-exchanges, NFT platforms) and speculative investors into the art ecosystem, often diluting artistic value in favor of financial value. As one commentator put it, the NFT market "disengages art from objects" and treats art purely as a tradeable asset (Taylor & Sloane, 2021, as cited in Calvo, 2024). Additionally, legal and ethical issues around ownership and authenticity complicate the economics. Cases of art theft and copyright infringement—where works were minted as NFTs without artists' consent—have undermined trust (Gibson, 2021). Intellectual property scholars note that in a digital realm of infinite reproducibility, the concept of owning a "unique" copy of an artwork is problematic (Frye, 2021; Joselit, 2021, both as cited in Calvo, 2024). This critique holds that NFTs create artificial scarcity that may be more valuable to speculators than to artists' creative vision or remuneration.

A widely cited barrier to adoption has been the ecological impact of blockchain, especially evident during the NFT boom of 2020–2021. Early on, most NFTs were built on energy-intensive proof-of-work blockchains (notably Ethereum pre-2022), which consume enormous electricity for mining and transactions. This gave rise to heated controversy: numerous artists and activists decried NFTs for their carbon footprint and contribution to climate change. Scholarly critiques likewise include environmental unsustainability as a key ethical issue with cryptoart (Calvo, 2024, 2). Although technological developments (such as Ethereum's later shift to a less energy-intensive model) significantly reduced the carbon costs (notably by 99% in the case of Ethereum), the damage to blockchain's reputation among many artists has been significant. Many remain unconvinced, noting that efficiency improvements do not fully erase the fundamental issue: the push for artificial scarcity and incessant transactions inherently encourages waste. Some theorists argue that even in a future with green blockchain energy, "scarcity is not something any free society should be aiming to produce" as a matter of principle (Whyman, 2021, as cited in Notaro, 2022).

This points to a broader tension in how blockchain is being operationalised within artistic contexts. As Future Art Ecosystems 3 notes, blockchain platforms are often framed not as spaces of artistic experimentation, but as instruments for building community or enabling transactions. "Art", in this context, becomes a narrative device—a means to attract users or bootstrap platforms—rather than a domain of experimentation in its own right (Serpentine Arts Technologies, 2022, 79–80). This framing risks reducing artists to content providers within a predefined economic infrastructure, rather than recognising them as infrastructural creators in their own right—agents capable of shaping protocols, interfaces, and systems through which artistic processes unfold.

The report warns that if blockchain systems continue to treat art primarily as a speculative asset or marketing tool, the medium could quickly lose relevance within the broader field of artistic innovation. The infrastructure itself must remain open to experimentation—across disciplines, across technologies, and beyond the narrow confines of the Web3-native community. As the report puts it, “There is a tendency to ignore historical institutional precedents, and where ‘art’ is more commonly a narrative vehicle for catalysing community and infrastructural development, rather than the site of experimentation itself” (80). To avoid this trap, artists must be empowered not only to make work with blockchain, but to reshape the very conditions under which blockchain art is made—by designing new protocols (for example Kosmas Giannoutakis (2024) imagined an alternative artistic consensus mechanism called a proof-of-creative-contribution), creating alternative economies, and proposing speculative futures for compositional and curatorial practice (such as the projects of Holly Herndon and Mat Dryhurst discussed in Zeilinger 2024). In this sense, the artist becomes not only a user of platforms, but a world-builder—someone who actively prototypes the infrastructures. This expanded role—as infrastructural designer and, ultimately, world-builder—sets the stage for the next chapter, which turns to the question of how artists can actively shape the ontologies, imaginaries, and ecologies of future art ecosystems.

Artists as System-Builders

Artists today are no longer defined solely by the works they produce or the tools they use. Increasingly, they operate as designers of systems—constructing dynamic processes, rule sets, and machinic environments through which aesthetic experiences and creative relations unfold. In this expanded role, the act of composition shifts from crafting outcomes to engineering conditions of their generation. From Cage's use of chance to Xenakis's formalisation of stochastic structures, from Brian Eno's generative software to the machinic grammars of web-based generative music, we witness a shift from composition as inscription to composition as system design. This shift intensifies in blockchain-based practices, where smart contracts, NFTs, and various decentralised applications allow artists to embed rules of interaction, reproduction, and authorship directly. What unites these otherwise disparate practices is a shared commitment to building systems that do not merely produce outputs, but redistribute agency and remain open to recomposition. In this context, the system becomes the artwork.

The system-builder operates through assemblages composed of both human and non-human actors, but also code, interfaces, sensors, protocols, and feedback loops. Agency is distributed—located not in isolated decisions, but in the patterned interactions that unfold within and across these systems. The artist's role becomes not only one of determining how autonomous processes are activated, how transformations are possible, and how responsibility is shared across networks, but also how are they made, what do they consist of, and what operations enable them. Whether through the stochastic logic of musical scores, or the programmable autonomy of smart contracts, the artist as system-builder choreographs a field of relations in which action is structured, negotiated, and refracted through the new systems built. The system-builder composes infrastructures for action—frameworks that structure how meaning, authorship, and participation unfold. Composition becomes the construction of systems that define procedures, automate transformations, and embed logic into the artwork itself. These systems are not neutral tools but expressive media—technical forms through which artists encode compositional intent, distribute control, and structure variability. In building such systems, artists redefine authorship at the level of process, creating works that remain active, procedural, and open to future reconfiguration.