

An Album in 1,000 Variations: Notes on the composition and distribution of a parametric musical work

Preprint. To appear in *Leonardo Journal*, MIT Press. Accepted April 2013.

Oliver Bown

Design Lab, University of Sydney, NSW, 2006, Australia

ollie@icarus.nu

and

Sam Britton

School of Arts, Brunel University, Middlesex, UB8 3PH, UK

sam@icarus.nu

Abstract

In this paper we discuss the making and distribution of an album that was created using parametric techniques and released in 1,000 distinct variations, as a kind of limited edition for the age of digital distribution. After describing the project we discuss how the project has affected our thinking about the production of electronic music, the process of musical distribution and the notion of tracks, musical works and uniqueness.

The Object

In 2011, under our collective moniker Icarus, we created an album that was available to download in 1,000 subtly varying versions, released with the title “Fake Fish Distribution” (FFD), on 6th February 2012 [NOTE1]. The album’s download site was set up such that each purchaser would get a unique version, and once all the versions were sold the album would no longer be available for sale. Through this process we engaged with issues of rights and ownership, collectability and uniqueness, authorship and methods of music distribution. The album was composed using a combination of software development and digital music production using the new possibilities afforded by the integration of Cycling 74’s *MaxMSP* [NOTE2] --- a visual programming environment for computer music --- and Ableton’s *Live* [NOTE3] --- a popular modern Digital Audio Workstation (DAW) with a strong emphasis on live performance. In building the software and composing the album we engaged with issues of systems music, parametric composition, designing variation, remix, issues of creative intent and determinacy, and the search for a creatively satisfying workflow when using the above methods. In this article we describe the project and discuss the pertinent issues that arose during its

development, with the hope of identifying and clarifying some key themes in current technologically mediated creative music practice.

Musical Background

Through Icarus, we have explored various forms of electronic music production, creating music with elements of electronica, breakbeats, free improvisation and electroacoustic composition. Beginning in the late 1990s we produced a series of studio-based albums in which the musical material was created from scratch through detailed editing on a time-line, either a MIDI piano-roll controlling multiple synthesizers and samplers, or a multi-track DAW, the traditional 'project studio' or 'bedroom studio' [NOTE4].

From around 2002, real-time algorithmic processes entered our repertoire through patches made in MaxMSP, and we became involved in improvising with computers, leading to a number of recorded works based on live material. During this time our creative process evolved towards the production of electronic music through a combination of algorithmic generative processes and collectively improvised long-term structure. In this respect we were part of a wider movement, initially inspired by the futurism of electronic dance music, which --- somewhat intuitively --- explored the methodologies pioneered in 20th century avant-garde music through creative programming.

Mechanical Variation

In 2011, we set out to return to studio production and make a full-length album from scratch, meaning a rendered rather than performed work, but still drawing upon our experience working with live improvised electronic music performance. At the time of our last full studio album, six years previously, Apple's iPod was four years old and the "mp3 revolution" was only a looming threat to CDs. In light of the subsequent revolution, returning to studio production in 2011 involved a rethink of our motivations and an appraisal of the context in which the music would be presented. With the age of mechanical reproduction fast approaching a point, in the digital realm, where all objects may be instantly available to all users all of the time, paywalls notwithstanding, the burgeoning presence of mechanical variation seemed to us to fit both the challenges of the age and an appropriate ambition given current technological-creative possibilities. This is not a new idea; parametric design is a mature subject in many fields, particularly architecture. However, the idea remains relatively unexplored in commercial musical culture.

The idea of producing an album in 1,000 variations was motivated by the desire to seek a balance between various forces: to produce music in the standard distribution medium --- the digitally-distributed album --- as opposed to producing

custom generative software, but with mechanical variation as a key element; to work with code and algorithms as a key production element, but still to creatively compose in timelines with audio and MIDI and virtual studio elements; to act as creative authors, whilst handing over part of the production process to an automated system; to provide a listening experience that is familiar, in that it is repeatable, and that can be experienced as any other album, but can also be perceived as a vast multifaceted entity, with musical qualities that are noticeably the product of a process of automated variation. Choosing 1,000 variations led to a quantity of music (about 1 month) with which the work's creator could conceivably engage, but not in any practical way. This quantity of music could be considered out of reach, but only just.

Music that adopts systems and processes, which might otherwise be described as “extra musical”, as part of a compositional procedure has a rich history; from early examples such as Mozart's dice games to the introduction of chance operations in the work of John Cage and the mapping of stochastic processes in the work of Iannis Xenakis. Amongst the 20th Century avant-garde, algorithmic and mathematical strategies are not uncommon. Examples include the player piano compositions of Conlon Nancarrow or works such as ‘Plus Minus’ by Karlheinz Stockhausen and György Ligeti's piano ‘Etudes’. By the later half of the 20th Century, as a result of the increasing proliferation of computers within music production and composition, such strategies could be said to have become commonplace (the chapter “Representations and Strategies for Algorithmic Composition” in Curtis Roads' (1996) “The Computer Music Tutorial” gives a good overview).

Within this ecosystem of systematized music, our focus was to approach parametric composition by fitting into a traditional electronic dance music, DAW-based approach to music creation. This approach is not based on generative process but on human creativity in a timeline-based composition environment. In order to make it parametric, a number of elements in a work that would typically be fixed are left as controllable variables, the adjustment of which results in a different output. Creating parametric music is therefore in the first instance a simple step back from normal electronic music composition, in which the final positions of various dials and sliders are left unknown. In reality, a simple leaving-open of possibilities does not in itself solve the problem of how to achieve mechanical variation, as the capacity for variation needs to be *designed* through the appropriate mapping of parametric variables onto musical outcomes, along with a suitable high-level strategy for then representing variation within the work.

Making and Distributing an Album in 1,000 Variations

FFD was made using Ableton Live and a series of additional tools created by the authors using Max for Live. At the core of this suite of Max for Live tools was a single parametric control, which could be used to determine which one of 1,000 different

versions, was currently being played back. This control, rendered as a large dial in the main Ableton Live interface (Figure 1), was used to interpolate between a series of keyframe timelines across multiple control channels, defining parameters that were then used to control various aspects of the Ableton Live composition, either directly or via more or less complex transformations in MaxMSP.

Through the version dial, the one-dimensional timeline of the traditional DAW becomes joined by a second one-dimensional control, which can be freely set by the composer at any time during the compositional process, just as the playback transport position can be freely set along the timeline. Together, these allow a composer to move seamlessly around the two-dimensional time-version space that defines each piece. A composition can then be thought of as a two-dimensional structure, slices of which form actual musical outputs.

The various time-dependent parameters that go on to define the musical content form surfaces over this 2D domain. The simplest way to build such surfaces without manually specifying the entire set of resulting points (this being the purpose of parameterization), is to create keyframe timelines and then use the version dial to interpolate between them (Figure 2). A system for creating, managing and describing the interpolation between such timelines was therefore built, using Emmanuel Jourdan's advanced breakpoint editor for MaxMSP [NOTE5]. The result of these steps is a system for parametric music composition in which any aspect of a musical piece can be specified in the resulting two-dimensional time-version space. Parameterized tracks, once defined, can be rendered into their 1,000 constituent audio file outputs through an automated process, across multiple machines if necessary. After each version is rendered, the version dial is incremented by a single step and the next rendering is initiated. Since the version parameter is used to interpolate between different parameter timelines, the result is a smooth variation from one version to the next.

What was controlled and how was it parameterized?

Broadly, three strategies were pursued. In the first approach, a piece of music was conceived of as a set of states, arranged in a space, through which any given trajectory would result in a specific instantiation of the music. The parametric variation consisted of interpolating between different trajectories through the musical state space. The arrangement of the state space and the choice of which paths to interpolate between defined the characteristics of the ultimate composition, both in terms of individual outputs and in terms of the entire set of outputs. Actual states consisted of groups of parameters, such as specific clips or tracks being activated. This approach allowed a fairly preconceived idea of the final piece to be developed whilst working on the musical material. For example, in one track we decided upon a fixed ending, and conceived of a number of musical trajectories that would arrive at that ending [NOTE6].

The second approach was more intuitive and bricolage-like, and involved iteration between parameters and compositional material, with all elements being mutually adapted until satisfying results were found. For example, a number of musical components were created, with parameterized elements identified, and a number of abstract parameter trajectories were defined. Whilst auditioning the track in a number of different states, connections between abstract parameters and parameterized elements were made, while both musical material and keyframe parameter timelines were altered. In this case the ability to rapidly audition the musical output over a range of variation indices and at different points in time was invaluable in gaining a broad enough understanding of the musical dynamics [NOTE7].

In the third approach, the mapping from parameters to musical outputs took a more opaque form, through the use of parameterized number sequences and other generative processes. In the track *Colour Field*, the version number was fed directly into a quadratic residue equation to produce a number sequence, which determined the durations between sequential updates. The process was applied to a number of distinct tracks, each stepping through a preconceived sequence but at different rates. In this case, the notion of interpolation did not apply.

In reality, a mixture of each of these processes, and variations of them, was applied freely in the composition of all of the works.

Reflections: Singletons, Formats and the Standard Listening Model

Owing to our mix of motivations, the final form of the record and the experience of composing it were elements about which we had few expectations; the project provided an opportunity to reflect on both the inherent structure and entrenched habits in our technologically-mediated compositional process. It has shone some light on our intuitive understanding of aspects of musical production, suggesting notions and creative technological possibilities that could be relevant in the near future of music production and distribution.

Production

A characteristic that distinguishes contemporary studio-based music production for commercial release from other forms of music production is that the musical work may never be conceptualized by its makers or experienced by its audience as a separate, or separable entity to the musical *recording*. This view applies most clearly to the practice of studio composition in electronic music, in which the studio acts as a tool in the creative process. That said, much studio activity involves recording songs, in which case the recording is unlikely to be the only instance of the song and is conceptually distinct. Yet, through the mechanical reproduction and distribution

of music --- be it album, EP or single --- the main way people experience those musical works is by listening to instances of a single recording [NOTE8].

Borrowing from programming terminology, such a product can be thought of as a 'singleton', an object of which there is only once real instance, no matter how many times it is referenced. Whilst some pieces of music exist in many manifestations, a large number effectively have a definitive form, the *album version*. When music is performed live, by contrast, no matter how constrained the demands are for interpreting that music, there is always variation; each performance is different, but all performances that derive from a specific composition are also distinct instances of the same musical work. In the case of singleton studio music, the instantiation occurs simultaneously with the development of the musical concepts, and may never be experienced as a separate entity. At the other extreme, a completely improvised piece of music, albeit the antithesis of a studio composition, may go full-circle and share this quality with a studio-composed product: a specific recording of a specific performance, once cast as a mass-produced 'record', may also become a singleton in exactly the same way that a studio composition does [NOTE9], and a number of recorded works fit into this category [NOTE10].

Our previous approach to singleton studio work has generally been centered on the affordances of the DAW as a linear structuring tool, which, like the studio in general, allows the composition of musical work to occur simultaneously with the creation of the work's one and only sonic manifestation, the "cut" or "bounce to disk". Live electronic music performance has matured with the advent of 'next generation' tools derived from DAWs, including Ableton Live. Live combines the DAW production process with functionality that enables realtime remixing of predefined structures in live performance. Its popularity demonstrates the musical relevance of this capability. Creative practice emerging around tools such as Live involves the entangling of musical concepts and actual sound, and remains in a state of flux, cut to a singleton output for a commercial release, but then also remixed in myriad ways by its creator in a performance context, or by other remix artists (either working from the original project file or from stems).

Distribution

The promotion and broadcast of music through digital networks of communication has radically affected the reach of artists, particularly independent producers, and naturally this is now the prevailing medium for the distribution of music.

Despite this transition, the traditional commercial music forms of albums (LPs), EPs and singles still predominate in musical output. This is understandable, as most commercial releases still appear on CD and sometimes vinyl, and the requirements of mainstream radio show playlists continue to influence decisions about song duration. Finally, music reviewers from both traditional and new media remain attached to these forms as legitimate targets for their attention and criticism. For this reason, the album, for most music producers, may continue to be a dominant

form of output, although its duration may no longer have an upper limit imposed by its storage medium. Meanwhile, vinyl continues to be a specialist, in-demand format.

The predominant tools of music consumption today, such as iTunes, also reinforce these formats through conventions such as the central use of the 'album' in any organizing structure, although increasingly the dominant structure is the playlist (descendent of the mixtape). Naturally, in that these tools are first and foremost 'playback devices', they do not change the fundamental relationship between music consumers and their music, largely remaining loyal to pre-internet music-listening experiences, aside from increases in speed and quantity of access. Whilst digital networks of communication have dramatically changed the relationship between producers and consumers, including the structure of music access rights (e.g., through services such as Spotify), this is less clearly reflected in the typical process of listening to a body of music: the consumer owns or has access to a collection of music, a static entity, which he or she listens to passively (i.e., without interaction, except in the selection of tracks) [NOTE11]. Although a consumer's music collection may be a unique set, the fingerprint of their musical interests, the elements of that set are not unique, but are shared amongst other individual's collections.

Systems music, algorithmic composition and generativity as tools for the everyday producer offer a fascinating new sense of the 'large-scale' in creative works [NOTE12], complementing interactivity as another digitally-facilitated enhancement to musical consumption [NOTE13]. As the avant-garde of the 20th century have shown, these processes have the potential to extend and augment the relationship between artists, their audiences and the context within which music is perceived to operate.

Although commentators such as Jaron Lanier (2010) despair that the modern internet is ruining creativity by undermining artists' revenue and promoting homogeneity, in our opinion these processes have equally been responsible for a renewed search for novelty amongst small dedicated niches of enthusiasts. These digital approaches to varied music consumption experience are good examples.

In the digital domain, the fact that generative and interactive works cannot be experienced directly through standard digital tools such as iTunes and portable playback devices (although they can be experienced on computers, mobile phones, tablets and other devices capable of running custom software [NOTE14]) reinforces the sense that they are outside normal music experience. The requirement for non-standard playback tools means that such works cannot currently be integrated into users' collections in the same way as other music, i.e., in playlists or through standard services such as SoundCloud or Spotify, and they would contravene the existing user experience. This does not of course mean that generative and interactive processes can't be used offline to produce musical outputs that do conform to the standard playback paradigm, and in certain cases standard albums may be produced to distribute or promote generative music systems, just as they may be to distribute or promote multichannel compositions (i.e., it should be

understood that these are not the definitive or preferred outputs). Nevertheless, although listeners could put add FFD to their playlists, the project required its own server for the custom form of distribution it required, and could not have been delivered through any of the standard services.

It is likely that generatively produced, but ultimately static, outputs will play an important role in forthcoming changes to music production and distribution. Static outputs allow producers to produce work that is in keeping with the dominant modes of musical distribution, whilst experimenting with new methods of mechanical variation. They invite listeners to gain familiarity with a musical work, which we think is a desirable feature of our current mode of music consumption. If Brian Eno's re-envisioning of his ambient music from the 1970s through the release of dedicated applications for smartphones is a success story for generative music, this may be in part due to the process of acclimatizing audiences to novel music production techniques via traditional static recordings.

Philosophy

Relevant philosophies of contemporary digital creative practice include Jacques Attali's (1977) vision of socialized music creation and Luciano Floridi's observations concerning a 'Philosophy of Information'. These both privilege the idea of interaction, and share the view that some critical identification with difference or variance in the broadest sense is a key component for the future of music.

In attempting to engage with the attributes of both static and interactive music, FFD can also be defined as an assemblage [NOTE15] that not only harnesses an array of generative techniques that inform its sound, but also desires subjective interpretation of this sound world by the listener, as an album should. Unlike many experiments in generative music, the mechanics of how the work is created needn't necessarily inform the work as an aesthetic experience. In practice, however, FFD opens up the larger inter-subjective territory regarding how each unique version differs musically from another [NOTE 16].

Critical Appraisal

The commercial release of FFD was supported by a publicity campaign. We received a number of reactions, and FFD was featured in numerous contexts, from the traditional music press to music technology websites and blogs. Overall it was positively received.

Reviews included the following:

"Ultimately, this album may rely a lot on computer input, and may sound, at times, like an extremely complex machine, yet [Icarus] have perhaps created here their warmest and most human record to date." The Milk Factory.

“The process breaks two commonly-understood notions about recordings: one, that digital files can’t be released as a “limited edition” in the way a tangible object can, and two, that recordings are identical copies of a fixed, pre-composed structure ... The conceptual experiment is all-encompassing. Just to prove the file is “yours,” you can even use it to earn royalties --- in theory.” Peter Kirn, Create Digital Music

A predictable outcome was that reviews focused mostly on the concept and not on the musical content. But, in general, reviews, both positive and negative, indicated that musically FFD was viewed similarly to our previous work.

The project concept also provoked negative reactions. In these cases, FFD’s distribution was interpreted either partly as a gimmick and not seen as integral to the genesis of the album itself, or elsewhere as an affront to musical industry and human creative authorship. According to a mainstream music critic (private communication):

“... if an artist doesn't put out a definitive version of their work but allows it to be repeatedly changed without their autonomy in the name of commerce it no longer can be said to be art.”

Several blog comments referred to the work as a technological affront to human artistic creativity, and described the music as cold or inhuman.

FFD was conceived as a labyrinth whose geography remains only partially known even to its designers. The diversity of outputs is made relatively inaccessible to listeners, who are welcome to seek alternative variations if they want to, but only by communicating with other version owners. In this sense, there is substance to criticism that sees this production of variation as more conceptual than tangible, but our decision to focus each listener’s attention only on their version should be seen as part of the conscious design of the work. Despite the fact that we initiated a discourse with the techniques we have described, we make no claim to be originators in this field. Our intent has been to develop it as a legitimate territory for artistic and creative exploration, and we hope that FFD will be accepted as a useful experiment.

References

Jacques Attali. *Noise: The Political Economy of Music*. University of Minnesota Press, 1977.

Manuel De Landa. *Philosophy and Simulation: The Emergence of Synthetic Reason*. Continuum, 2011.

Luciano Floridi. *Information: A Very Short Introduction*. Oxford University Press, 2010.

Jaron Lanier. *You Are Not a Gadget: A Manifesto*. Knopf, 2010.

Clay Shirky. *Cognitive Surplus: Creativity and Generosity in a Connected Age*. Penguin, 2010.

Barry Truax. *The Handbook for Acoustic Ecology*. World Soundscape Project, Simon Fraser University, 1978.

Curtis Roads, *The Computer Music Tutorial*, MIT Press, 1996.

Notes

[NOTE1] Available from <http://www.icarus.nu/FFD>.

[NOTE2] <http://www.cycling74.com>.

[NOTE3] <http://www.ableton.com>.

[NOTE4] The latter being a derivation of the former, which has its origins in the 1950s with the many experimental radio studios that were dedicated specifically to electronic music creation. Bedroom studio production emerged as a low-budget, grass-roots approach to 'commercial' electronic music (as opposed to academic music), the traditional output being the album, EP or single, particularly the 12" single used by DJs.

[NOTE5] From Emmanuel Jourdan's Ejies collection, available from <http://www.e-j.com>.

[NOTE6] Examples of this are the tracks *Old D* and *Two Mbiras*.

[NOTE7] Examples of this are the tracks *Spineez of Breakout* and *Three Stupidities*.

[NOTE8] Clay Shirkey makes the point that although we talk of 'copies' of a piece of music, in the most abstract sense digital reproductions are not even copies, but instead occurrences of the same thing. "Nobody asks you for a *copy* of your phone number" (Shirky, 2010).

[NOTE9] We can take the programming metaphor one step further and say that an unrecorded musical improvisation is akin to an anonymous class, an entity that does not have a life beyond the given event in which it occurs.

[NOTE10] An example might be Keith Jarrett's *Köln Concert*, celebrated by many as a definitive performance.

[NOTE11] Although according to Truax (1978), listening is never passive. Music is that which is constructed internally as a result of actively transforming sound

though perception. We do not disagree with this notion, but it is simply more practical in this context to think of the recorded sound as the music.

[NOTE12] Although not itself generative, an excellent, recent semi-musical example of a 'large scale' media work is Christian Marclay's "The Clock". Here 'large-scale' refers to the quantity of content created or brought together.

[NOTE13] RJDJ, is an example of an interactive music application made for the iPhone that is both able to 'personalize' songs and 'generate' music by analyzing contextual data collected by the host device. Many bands and groups have also released dedicated applications available for mobile devices that allow the user to remix and interact with songs otherwise available as fixed format recordings

[NOTE14] We exclude the use of Internet radio as a means to stream software-generated music. This exclusion is made on the basis that radio is a one-to-many service and we are talking about personal on-demand listening. We realize of course that a system can perfectly well spawn personalized streams.

[NOTE15] Here we refer to the theory of assemblage as outlined by Manuel De Landa (2011) in "Philosophy & Simulation: The Emergence of Synthetic Reason".

[NOTE16] Our audience was invited to participate by sharing their copy on a private mailing list reserved for other version owners.

Glossary

MIDI: Musical Instrument Digital Interface

DAW: Digital Audio Workstation

Biographies

Oliver Bown works across the borders of music composition and performance, creative software development and theoretical research. His recent work explores electro-acoustic improvisation, and the design of autonomous improvising software systems, which includes collaborations with trumpeter Tom Arthurs, along with Lothar Ohlmeier, Maurizio Ravalico and Isambard Khroustaliov in the Not Applicable Artists, and in Australia with clarinetist Brigid Burke and trombonist Adrian Sherriff. In 2010 he curated a series of events showcasing autonomous improvising software systems at Melbourne's Guildford Lane Gallery. His performance and installation software, including solo work and collaborations with the multimedia artists Squidsoup, has appeared at the Bath International Music Festival, the Huddersfield Contemporary Music Festival, The Science Museum's Dana Centre, Club Transmediale, The Sonic Arts Expo, The Kinetica Art Fair, Futuresonic and The UK's Future of Sound Tour. He recently completed an Australia Research Council funded project with the Centre for Electronic Media Art in Melbourne, exploring the creative aesthetics of ecosystem computer models. He is current a postdoctoral fellow at the University of Sydney's Design Lab.

Sam Britton studied architecture at the Architectural Association from 1996 to 1999 and began writing electronic music in 1996. In 2006, he attended a masters course in electronic music and composition at IRCAM in Paris (one of the only students in its history to have been accepted without having completed an undergraduate degree in composition), and began to write composed electro-acoustic and concert music for which he was the recipient of a SACEM composers bursary. He is currently working as part of the London Sinfonietta's Writing the Future commissioning scheme and researching a PhD in composition under the supervision of Richard Barrett, Peter Wiegold and Christopher Fox at Brunel University. In addition to realizing music with Ollie Bown as Icarus and under the alias Isambard Khroustaliov, he has collaborated with numerous artists, directors and software developers on work for film, media and performing arts. Productions featuring his work have been exhibited at MOMA New York, The Tate Gallery, The Colorado Contemporary Art Museum, ARS Electronica, broadcast by Channel 4, BBC 2, BBC 3, MTV, HBO and ITV, and have been the recipients of numerous international accolades including a Creative Review Creative Futures award and a Best of Nemo and ARS festival awards.

Figures

Figure 1: A master Max for Live device with version dial. The device took care of recording and all aspects of variation management, and could be used to dynamically link in track-specific Max patches and read time-line data from a track-specific project folder. Besides pop-up timeline editors, the only essential (non-debugging) interface element is the version dial, which immediately adapts the contents of the track to a different version state.



Figure 2: A live project showing the process of interpolation between breakpoint functions (lower two track lanes) to produce a resulting interpolated pattern (upper right breakpoint function). The upper left breakpoint function describes a mapping from the version dial (x-axis) to the actual value used to interpolate between the input breakpoint functions. This could be repeated multiple times for each piece. In the top right the outputted renderings are shown being recorded into a directory.

ICARUS Colour Field 04 [The Parametric Album]

Generative Music
1 item, 79.36 GB available

Name	Date Modified
0.wav	Today, 1:43 PM

Compressor

Envelope Threshold

Peak

Ratio

Release

Lookahead

Gain

Limiting

ICARUS - Colour Field

Track Index 0

beat 31

stop 508

Drop Audio Effects Here

[filibuster_fx]

[multi_interp_player]

function0

fb_fx1

grid

range

fb_fx2

grid

range