

# Alternator: A General-Purpose Generative Music Player

Ian Clester

Georgia Institute of Technology  
Atlanta, Georgia, United States  
ijc@gatech.edu

Jason Freeman

Georgia Institute of Technology  
Atlanta, Georgia, United States  
jason.freeman@gatech.edu

## ABSTRACT

Computers are a powerful technology for music playback: as general-purpose computing machines with capabilities beyond the fixed-recording playback devices of the past, they can play generative music with multiple outcomes, or computational compositions that are not fully determined until they are played. However, there is no convenient platform for distributing generative music in a way that captures the space of all possible outputs. This absence hinders composers’ and listeners’ access to the possibilities of computational playback.

In this paper, we address the problem of distributing generative music. We present a) a format for bundling computational compositions with static assets in self-contained packages and b) a music player for finding, fetching, and playing/executing these compositions. These tools are built for generality to support a variety of approaches to making music with code and remain language-agnostic. We take advantage of WebAssembly and related tools to enable the use of general-purpose languages such as C, Rust, JavaScript, and Python and audio languages such as Pure Data, RTcmix, Csound, and ChuckK. We use Audio Worklets and Web Workers to enable scalable distribution via client-side playback. And we present the user with a music player interface that aims to be familiar while also presenting the possibilities of generative music.

## 1. INTRODUCTION

*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.*

— Brian Eno, *A Year with Swollen Appendices*

In the essay “Generative Music” [5], Eno describes three kinds of music. The first is live music, in which the music



Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** owner/author(s).

Web Audio Conference WAC-2022, July 6–8, 2022, Cannes, France.

© 2022 Copyright held by the owner/author(s).

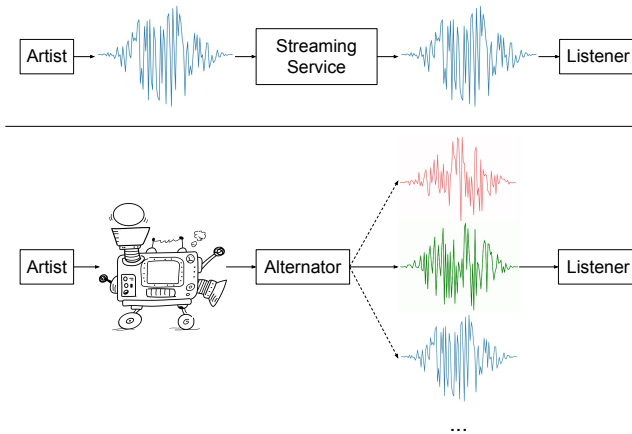


Figure 1: Conventional streaming model (top) vs. Alternator

comes out differently every time. Even the same musicians playing the same piece will sound a little different between performances. In musical traditions that emphasize improvisation on a form, each performance of a piece may differ significantly, but still be recognizably the same piece. Freer practices go further still, taking each performance in different directions bound together only by the improvisational spirit of the performers — or by a score that embraces randomness, with variability built-in by the composer.

The second kind is recorded music, in which music is frozen in time. The recording consists of whatever particular sounds occurred that time, e.g. samples output by an ADC or grooves etched into a record, and the recording does not change each time it is played (other than, perhaps, the medium degrading). Nonetheless, recorded music is wildly successful because it is convenient for both artist and listener. Once an artist records their music, their fans can listen to it across time and space, whenever they want and however they want. A recording, as information, is inherently less scarce than a performance—a performer can only play one thing in one place at a time. Digital recordings in particular are non-rivalrous: my listening to a recording does not prevent your listening to an identical copy of the same recording, and there is no venue to limit the size of the audience.

So on the one hand we have live performance, which is dynamic but constrained and scarce, and on the other record-

ings, which are static but easy to distribute and reproduce. Then there’s the third alternative: generative music. By representing music as *systems that generate audio* rather than audio itself, we can have the best of both worlds: dynamism and availability; possibilities and scalability.

Eno described that third alternative in 1996, but where is it today? Live music and recorded music are ubiquitous, but generative music remains niche. Modern music production invariably involves a computer, but playback uses almost none of the machine’s potential. Languages and environments abound for making computer music, but distribution remains challenging, and the solutions that exist are typically custom-made for a particular musical work, artist, or language. From the perspective of the listener, these forms of distribution present a barrier to listening and separate generative music not only from other music, but even from other generative music.

In this paper, we address this state of affairs and investigate the problem of distributing generative music. We describe Alternator: a system for distributing and playing computational compositions, making the possibilities of such compositions more accessible to listeners and more useful to composers.<sup>1</sup>

Fig. 1 summarizes the conceptual model of Alternator. In conventional music distribution platforms (top), an artist produces an audio recording of their music and uploads it to a streaming service, which then streams it to listeners. A fixed, unchanging waveform represents the musical work, and listeners hear the same thing every time. In the Alternator model (bottom), the artist creates and uploads a *computational* recording of their music instead—rather than audio, a program that generates audio (represented by a contraction in the figure). At the time of playback, the program runs, and there may be many possible outcomes; listeners hear a different one each time.

## 2. BACKGROUND & RELATED WORK

### 2.1 Musical Inspirations

Our work is inspired by the long tradition of thinking about music in terms of systems and processes. Aleatoric or indeterminate music may include instructions in the score that allow for many possible outcomes at the time of performance, as in works by Cage, Brown, and Stockhausen [3].<sup>2</sup> Steve Reich [14] made explicit the notion of a piece of music as a process, and emphasized gradual, perceptible processes as opposed to the chance of Cage (or the “seldom audible” processes of serial composers). Brian Eno [5], inspired in turn by Reich, articulated the idea of generative music as a kind of hybrid between live and recorded music, and linked its future to advances in the capabilities of ubiquity of computing technology.

Alternator also builds on work in algorithmic composition; for a summary, see [6]. Algorithmic thinking long predates

<sup>1</sup>A live demo is available at <https://ijc8.me/alternator>, and the source code can be found at <https://github.com/ijc8/alternator>.

<sup>2</sup>Some aleatoric works are “indeterminate with respect to their composition,” as Cage puts it (or, as we might say, they vary only at ‘compile-time’). In this paper, we are more interested in those that are “indeterminate with respect to their performance” (which vary at ‘run-time’), as these compositions encompass multiple outcomes.

automatic computing machines, but as Essl notes, “[by] using generative composition algorithms on computers, music can be created in realtime by an autonomous and infinite automatic process”. Essl also observes that, due to the nature of generative music, “distribution on a reproductive medium such as a compact disc seems highly inappropriate.” While CDs may have since been eclipsed by streaming services, the fit is no better, as the medium (static audio recordings) remains reproductive rather than generative. The core aim of Alternator is to provide a more appropriate medium, suitable for the variable output and variable (potentially infinite) length of generative compositions.

Though it draws on these artistic movements and fields, Alternator does not enforce any particular aesthetic constraints. The sole requirement is that a musical work have a computable translation into sound:<sup>3</sup> thus, Alternator supports generative and aleatoric music as well as deterministic or conventionally-recorded music — and everything in between. Despite this aesthetic agnosticism, Alternator’s existence is predicated on the idea that there is value in generativity, and that the idea of the musical work as a computable process — ever-changing, describing a field of possibilities — is an idea worth sharing. Alternator embraces the vision of composers of aleatoric, generative, and algorithmic music, and deals with the problem *after* composition: getting music to a listener.

### 2.2 Technical Precursors

#### 2.2.1 Generation

Many projects have addressed the problem of allowing for many possible outcomes from a computer-playable piece of music. SSEYO’s Koan software (used by Brian Eno as described in [5]) and its successors explicitly address this under the umbrella of *generative music*, usually with an ambient emphasis. Programming languages designed for music offer a general solution, from Max Mathews’s original MUSIC onwards [17]. Today, a composer can use audio programming languages (Max/MSP, Pure Data, SuperCollider, ChuckK, RTcmix, Csound, etc.) to describe a computational program which may depend on random values or external input and thus encompass a broad range of possibilities. A composer may also use general-purpose languages toward the same end, possibly in conjunction with composition/synthesis libraries such as Aleatora [4] (for Python) or JSyn [2] (for Java).

Alternator does not offer a new language for composition, nor does it enforce the use of any particular existing language. Instead, it leverages the development of WebAssembly and open-source projects such as Emscripten and Pyodide to allow the composer to choose among the large set of existing options.

#### 2.2.2 Distribution

Generation is only part of the puzzle. A composer may create a piece as a Max patch, but they still need to distribute it. They can upload the patch somewhere and send out a link, but this requires listeners to install a new application to play it. Max/MSP allows for creating stan-

<sup>3</sup>More specifically, Alternator requires bundles that produce up to two channels of floating point samples at consumer-audio sample rates, but none of these practical restrictions are essential to the idea.

dalone patches (patches bundled with the Max runtime), but this requires listeners to run an untrusted binary executable. Even if listeners trust the artist completely, the playback experience is isolated: instead of a familiar music player, they see the interface of the patch (which may lack basic controls such as pause, rewind, and seek), and there is no way to put the patch in a playlist with the rest of their music.

Another option is to render the piece as an audio file and distribute that instead, but then the composer must choose one fixed rendering of the piece, eliminating all other possible outputs. To address this, they may opt to run a few renderings — generate and include multiple versions of the piece (as in jazz albums that include multiple takes for each tune). But doing this for every piece (or even a single piece, if there are thousands of possible outcomes) is not feasible.

Some artists have solved the distribution problem individually by leveraging programming expertise that is independent of their musical skills. For example, Jason Freeman’s *Grow Old EP* consists of pieces which vary their output as the days pass.<sup>4</sup> New audio files are automatically generated each day and replace the old ones. Thus, the pieces are always available in the same place, can be readily played in a web browser without requiring the user to execute code, and realize new possibilities day by day.

Alex Bainter’s *Generative.fm* [1] solves the same problem another way. On *Generative.fm*, listeners are presented with a web-based music player offering 50+ ambient music generators in a familiar interface. The interface is simple and consistent across pieces. Unlike the *Grow Old EP*, the pieces are generated at the time of playback. The generation occurs client-side, as the user’s browser executes the code for each piece on demand.

Though these both effectively distribute generative music, they are artist- and technology-specific. The *Grow Old EP* is an album with all pieces written in RTcmix by Jason Freeman and rendered server-side. *Generative.fm* is a collection of ambient works written by Alex Bainter with the Web Audio API. *Alternator* takes inspiration from these efforts and aims to generalize, serving as a platform where many composers, using many different languages, tools, and workflows, can share their music.

The nearest precursors to *Alternator* are automated Internet radio stations that stream real-time generative music, such as `rand()`<sup>5</sup> and *Streaaam* [9]. Unlike the previous examples, these projects serve as independent platforms where artists (at least within a particular community) can submit their own work using a variety of languages. As radio stations, all listeners hear the same thing per the station schedule. We present *Alternator* similarly, as a platform where artists can submit their own generative music, with the key distinction that it takes the form of a personal music player (such that listeners can independently listen to whatever they want, whenever they want) rather than a radio station. This would be difficult if *Alternator* rendered on the server-side like `rand()` or *Streaaam*, as each simultaneous listener would require their own sandboxed composition

<sup>4</sup><http://distributedmusic.gatech.edu/GrowOld/>

<sup>5</sup>`rand()` has been down since 2007, but some information is available on the Internet Archive ([https://web.archive.org/web/20070629095427/http://www.r4nd.org/rand\\_home.html](https://web.archive.org/web/20070629095427/http://www.r4nd.org/rand_home.html)) and at <https://www.bbc.co.uk/radio3/cutandsplice/rand.shtml>.

process running on the server. Instead, *Alternator* executes compositions in the browser, which allows listeners to choose what they hear while avoiding the scalability issues inherent to server-side execution.

### 2.2.3 Platforms & Archives

Though it has a different focus, *Alternator* bears some relation to the mobile apps *RjDj*, *MobMuPlat* [10], and *PdDroidParty*<sup>6</sup>, as well as *JSyn* [2] and *WebPd*.<sup>7</sup> The three apps play Pure Data patches on mobile devices, while *JSyn* and *WebPd* enable interactive musical applications to run in the browser. Like these, *Alternator* runs computational compositions in the browser and on mobile, but it aims to support many different languages and provide a consistent music player interface.

Looking further back, MPEG-4 Structured Audio [16] deserves special mention. This forward-thinking standard likewise dealt with the distribution and playback of computational audio. However, it focused on the *efficiency* (in terms of compression) of describing audio computationally—an ill-fated tack given increasing bandwidth and storage—rather than the new possibilities afforded by such a general representation. Additionally, it inherited a synthesis framework from *Csound* (and older MUSIC-N languages) and embedded a high-level language (SAOL) in the standard itself, making it difficult to use alternative languages or approaches to computational composition (and more difficult to implement the standard). In contrast, *Alternator* aims to be a “common carrier,” providing an executable format sufficiently low-level to accommodate all kinds of approaches to computational composition and sound synthesis.

*Alternator* differs from digital archival projects such as *Rhizome ArtBase*<sup>8</sup> and Miller Puckette’s *Pd Repertory Project* [13]; it aims to enable any composer to share their work, rather than attempting to preserve works of historical significance. It also differs from platforms such as *Scratch* [12], *EarSketch* [11], and *TunePad* [7]; like *Alternator*, these are focused on computational art & music and enable sharing, but they are geared towards pedagogy rather than distribution and lack a dedicated interface for the listener.

In short, *Alternator* aims to fill a void: a generative music player for any composer, in any language, on any device with a browser.

## 3. GOALS

*Alternator*’s goal is to enable composers to easily distribute and share their compositions—including generative music, without compromises to fit it into static media. From a complementary perspective, *Alternator*’s goal is to enable listeners to easily discover and listen to such compositions, experiencing the field of possibilities inherent in each piece without abandoning the features, interfaces, and portability they are used to finding in a music player.

This goal implies two others. The first is generality: to enable composers to easily distribute and share their compositions, *Alternator* must *play* their compositions. Therefore, it must avoid mandating a single chosen way to create computational compositions. If composers have to drop their existing tools, expertise, and workflows to fit into a mold,

<sup>6</sup><https://droidparty.net/>

<sup>7</sup><https://github.com/sebpiq/WebPd>

<sup>8</sup><https://artbase.rhizome.org/>

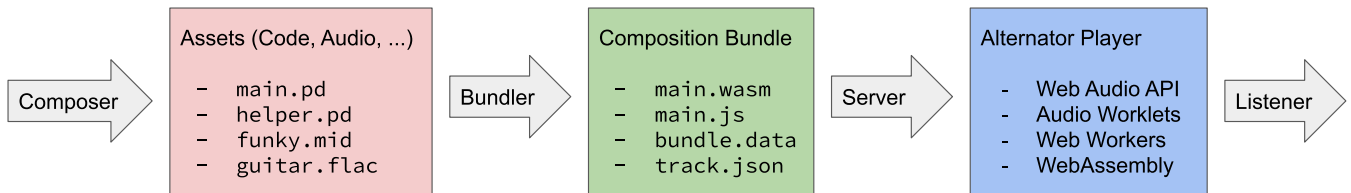


Figure 2: How music flows from the composer to the listener in Alternator.

the battle is already lost. Alternator must strive to provide a general platform, suitable for the tools and languages we have today and for those yet to be invented.

The second implied goal is stability. A platform is not useful to composers if it requires them to continually maintain and update their compositions every time the platform changes. This is unheard of in music streaming services for the simple reason that PCM data is forever—after a composer uploads the audio recording for a given track, no maintenance is required. Unfortunately, this situation is common in software, especially on the ever-changing web: an API for some service changes, or disappears entirely, and breaks anything that depends on it. This requires regular upkeep by software engineers to keep their software running as the foundations shift beneath it. As this is unacceptable for most composers, Alternator must endeavor to make its compositions more like audio files than living software in terms of maintenance requirements.

In the following section, we describe how Alternator meets this and the other goals in its design and implementation.

## 4. DESIGN AND IMPLEMENTATION

### 4.1 Architecture

The core of Alternator is a music player that executes computational compositions, generating samples on demand for playback. These computational compositions are represented as *bundles* containing all of the resources needed to run the piece: code (potentially in multiple languages), audio samples, MIDI data, models, etc.

These bundles could be executed server-side (streaming generated audio to each client) or client-side. Server-side generation, however, would require a backend capable of running a process for each concurrent listener (or else trade-off some of this computation cost for storage cost or compromised variability).<sup>9</sup> Given that most digital music listeners today are listening on relatively powerful devices (smartphones, tablets, laptops), client-side execution is feasible and inherently more scalable.

Thus, Alternator executes musical bundles client-side, in the browser. However, the browser does not understand languages like Csound or Pd. Until a few years ago, it only understood JavaScript; this limitation led to `asm.js` and then to development of WebAssembly (henceforth “Wasm”), a binary instruction format for a portable virtual machine. Al-

<sup>9</sup>Note that the Grow Old EP, which does use server-side generation, avoids this problem by only generating each piece once per day. This optimization is only possible because it fits the artistic intent of the album; the pieces evolve slowly, day-by-day, so there is no need to generate them at the exact time of playback.

ternator takes advantage of this development, using Wasm builds of `libpd`,<sup>10</sup> `libsound` [18], `ChucK`,<sup>11</sup> `RTcmix`,<sup>12</sup> and the like to enable the execution of patches and scores in the browser.

The complete flow of a piece of music in Alternator is depicted in Fig. 2. First, the composer creates a piece using their preferred tools. In this example, their piece consists of a Pure Data patch (`main.pd`), an abstraction (`helper.pd`) used by the main patch, a MIDI file (`funky.mid`) containing recorded performance data, and an audio sample (`guitar.flac`). The patches and MIDI/audio resources are then bundled together by the *bundler*, using the Pure Data template. The template consists of a Wasm blob and some glue JS. These are combined with the composer’s assets and some metadata about the track (title, artist, etc.) to produce a bundle. Finally, this bundle is statically served to a listener’s client, which executes the JS and Wasm (which may access the static assets) to generate audio in real-time when the listener plays the track.

### 4.2 Execution

In the last subsection, we used Pure Data as an example, but the Alternator player does not have special support for any particular audio language built-in. Rather, it supports a general executable format; the only requirement is that the executable can fill buffers with samples, and (optionally) indicate when it is finished.

An Alternator executable is called a *bundle*. A bundle includes some JavaScript that implements two functions: `setup()`, which takes the listener’s sample rate as an argument and does any necessary preparation, and `process()`, which takes in a buffer and fills it with samples. Unless the piece is written in JavaScript, the bundle also includes a Wasm blob. Typically this blob corresponds to a language runtime, but it may also correspond to the piece itself if it is written in e.g. C/C++ or Rust.

Thus far, we have implemented templates for Pure Data, `ChucK`, `Csound`, `RTcmix`, `Python` (with our composition framework `Aleatora` [4]), and static audio files (WAV and Ogg Vorbis). Alternator’s design is intended to be future-proof: regardless of what music programming environments exist in the future, they can work in Alternator as long as they 1) can spit out samples and 2) fit in a Wasm blob. #1 is a prerequisite for any audio programming environment (it must be audible), and #2 is in a good state today thanks to efforts such as `Emscripten` and the `Rust` toolchain. It is of course impossible to predict the future, but existing support

<sup>10</sup>Claude Heiland-Allen’s `empd`: <https://mathr.co.uk/empd/>

<sup>11</sup>[https://github.com/ccrma/chuck/tree/main/src/host\\_web](https://github.com/ccrma/chuck/tree/main/src/host_web)

<sup>12</sup>We created a Wasm build in a fork: <https://github.com/ije8/RTcmix>

and the simplicity of the WebAssembly specification bodes well for its continued viability as a compiler target.

One important execution detail is endings: how does Alternator know when a piece is done? Pieces can end at any time by returning fewer than the requested number of samples in `process()`. This signals to the player that the piece is finished, and it will not call `process()` again. Some languages lack built-in notions of endings. In these cases, the bundle template can implement this notion in whatever way is most convenient. For example, the Pure Data template allows a patch to signal that the piece is over by banging a `send` object named “finish”.

When the listener plays a piece, Alternator fetches the corresponding bundle. The JavaScript is executed in a Web Worker, and typically fetches and instantiates a Wasm module. An Audio Worklet communicates with the Web Worker, transferring buffers to fill with freshly-generated samples. The Audio Worklet uses double-buffering with a moderate buffer size (1024) to avoid hitches, connects to a `GainNode` for volume control, and finally gets audio out to the listener.

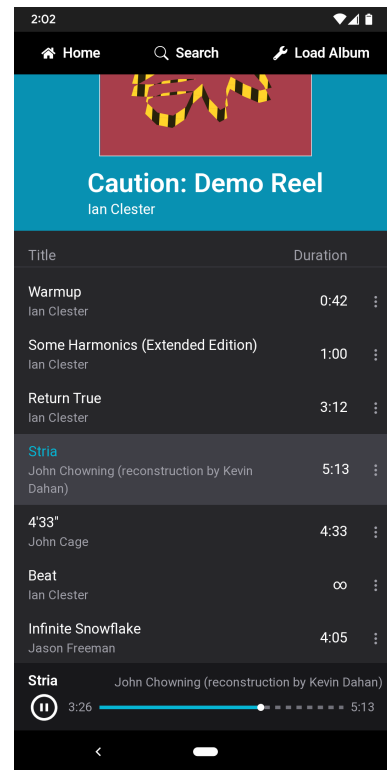
### 4.3 Listener Interface

In pursuit of its goal for listeners, Alternator adopts a familiar music player interface. The basic features of a music player are playing, pausing, resuming, seeking (skipping forward/backward within a track), and switching tracks. It is clear what these should do in the case of a static recording, which is pre-computed time-series data to be replayed. Alternator translates these core operations into a generative music context. From Alternator’s perspective, music is something that can generate sound. To play, then, means to start generating sound; pausing pauses generation; and resuming picks up from the same point. However, some aspects of the interface require special consideration in a computational music context.

Seeking is less straightforward, but nonetheless has clear analogs. Seeking backward should replay exactly what was heard the first time, as the listener typically uses seeking backward to hear something again. So, Alternator maintains a growing history buffer for the playing piece. Seeking backward 10 seconds, for example, will replay the previous 10 seconds of generated samples, and then return to generating fresh samples from where it left off.

Seeking forward, however, requires playing samples that have not been generated yet. The only way to generate those samples is to reach that point in the piece. Thus, when the user seeks forward, beyond what has already been generated, Alternator generates the intervening samples as fast as possible (faster than real-time, only limited by the performance of the bundle) and then resumes real-time audio generation and playback from the target position.

Alternator expands the visual language of the seek bar to convey these differences. Electric blue indicates the segment in the past, which have necessarily already been generated. Solid gray indicates samples in the future which have already been generated. Such samples only exist after a seek backwards; the user can seek forward to these samples instantly, because they were already generated. Dashed gray indicates samples in the future which have yet to be generated; this region is still “potential sound,” as yet unrealized in this playthrough and possibly indeterminate. When the current position is also the end of the generated samples (in other words, when there is no solid gray), the dashed line moves



**Figure 3:** It is feasible to generate and play compositions (such as this Csound reconstruction of “Stria”) smoothly, in real-time, in a browser, on a smartphone.

as it shrinks to convey activity: the system is in motion.

Another consideration for a computational music interface is duration. In an ordinary music player, duration is straightforward: because each recording is static, it is known in advance, and therefore it must be finite, with a known duration. In Alternator, none of this is necessarily true. A piece may be infinite, continuing until the user intervenes. Or a piece may be finite, but with many possible durations (say, anywhere from 2:30-3:00). Due to the Halting Problem, it is impossible for Alternator to determine the duration of a piece in advance (or whether it ever ends). From a UX perspective, however, it is valuable to the listener to know what to expect. Thus, the composer declares the duration of a piece in its metadata. If the piece is finite, Alternator will display the piece’s duration (or duration range) and use it in scaling the seekbar. If the piece is infinite, Alternator will display “∞” as the duration instead, and will use the time of the last generated sample + 10 seconds as the duration in the seekbar,<sup>13</sup> so that the user can still skip ahead of the already-generated audio.

Finally, there is one element of Alternator’s interface which has no analog in a conventional music player: the “view source” button. Pressing it opens a window showing the contents of the bundle: the code and assets that make up the piece. It also includes a link to the repository (elaborated in §4.4) containing the piece, which may have additional information about how the piece was made (such

<sup>13</sup>In the future, this seek-ahead window may be user-configurable.

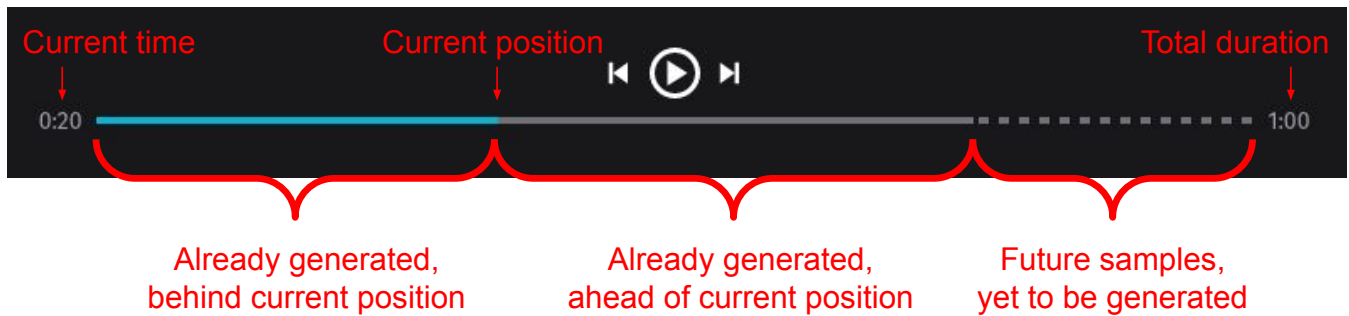


Figure 4: Alternator extends the visual language of conventional music players for computational music. Seeking to the dashed region, which represents the unknown future, will trigger faster-than-realtime rendering.

as documentation or more source code, if the piece is written in a compiled language). This implications of this feature are discussed further in §5.

#### 4.4 Backend

Our discussion so far has focused on the Alternator client and bundler. By comparison, the backend is simple: the only essential thing is static file hosting for bundles, and some mechanism for composers to upload their work. Other features such as search, recommendations, and playlist management require more from the backend, but these have been solved in existing music players and do not require changes in a computational music context.

Because the backend is not our focus, we leverage the existing open-source ecosystem. Drawing inspiration from *utterances*,<sup>14</sup> a comment widget backed by GitHub Issues, Alternator uses GitHub as its backend for this initial release. Albums in Alternator — which are both software and music — take the form of GitHub repositories. An album-repo includes all the track bundles, some JSON metadata, and cover art.<sup>15</sup> Alternator can discover and search through these albums using the GitHub API because the repositories are marked with the `#alternator-album` tag.

#### 4.5 Composer Interface

Up to this point, we have described how Alternator works and what it looks like for the listener. In this section, we demonstrate how it works for the composer with a concrete example.

To get started, we need a generative piece to bundle. For this example, we'll use the simple Pure Data patch shown in Fig. 5, which plays random harmonics with vibrato for 10 seconds. This is an ordinary patch which is playable on its own if Pure Data is installed, but there are two details of how this connects with Alternator that are worth mentioning.

First, as in normal Pd usage, `loadbang` sends out a bang when the patch is loaded. In Alternator, this happens when the user plays (or resets) the track. Second, the patch sends a delayed bang to a special destination called `finish`. As mentioned in §4.2, the Pure Data template for Alternator listens for this to indicate that the track has finished play-

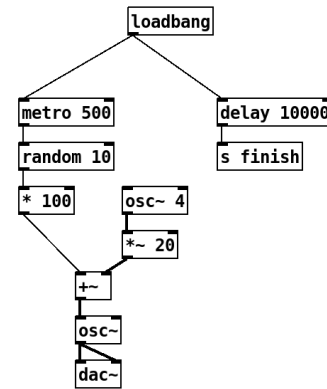


Figure 5: Example Pure Data Patch

ing. This is necessary because Pure Data, like other signal-processing-oriented languages, does not have a built-in notion of “endings,” and it demonstrates the adaptability of Alternator’s model to different languages and environments.

Next, we need to set up the right structure and provide some metadata. First, we create a directory for our album. Inside, we create `album.json` with the contents:

```
{
  "title": "My Cool Album",
  "artist": "Jen Rétive",
  "cover": "my-album-cover.svg",
  "tracks": ["bundles/my-first-track"]
}
```

and save our album art as `my-album-cover.svg`.

Next, we save the patch as `main.pd` in a subdirectory called `my-first-track`. Inside, we create a file called `track.json` with the contents:

```
{
  "title": "My First Track (Some Harmonics)",
  "artist": "Jen Rétive",
  "duration": 10,
  "channels": 1
}
```

Finally, we run the bundler:

```
$ ../alternator/bundle.py pd my-first-track
Creating output directory: bundles/my-first-track
```

<sup>14</sup><https://utteranc.es/>; repository at <https://github.com/utterance/utterances>

<sup>15</sup>Beyond the bundles and metadata, album-repos can contain anything else, such as the source code, a README, LICENSE, documentation, etc.

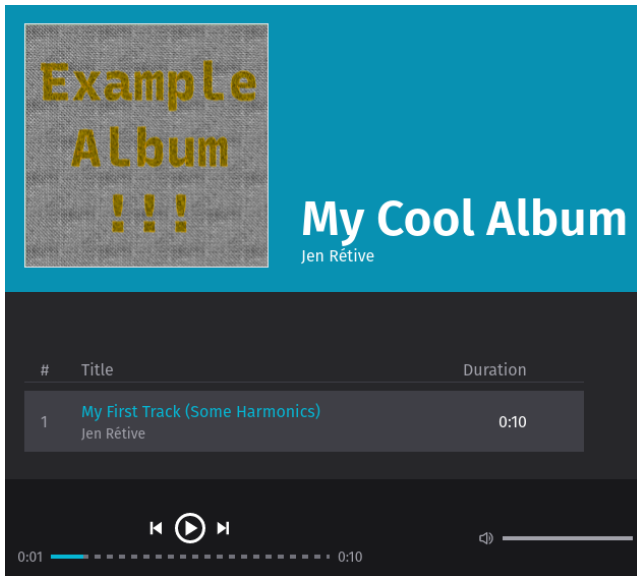


Figure 6: Our example patch is now playable in Alternator!

```
Copying main.js from pd template.
Copying main.wasm from pd template.
Bundling /main.pd
Saving finalized track.json.
Finished bundle: bundles/my-first-track
```

The bundler takes our patch and bundles it together with everything we need to run it (the Pure Data runtime compiled to WebAssembly and some glue code). Everything that Alternator needs to play the piece is stored in `bundles/my-first-track`. If our patch had any additional assets it needed (e.g. a `.wav` file or an abstraction in another patch), these would likewise be bundled in the track.

At this point, we have a playable bundle in Alternator (Fig. 6). If we host it somewhere on the Internet, we can then distribute a working link like <https://ijc8.me/alternator/?u=https://example.com/my-album>. If we want our album to be discoverable via Alternator’s search feature, we can put it on GitHub as discussed in §4.4.

## 5. DISCUSSION & LIMITATIONS

In this paper, we have presented Alternator, a platform for distributing and playing computational music compositions: music that is generated on-demand, such that each track may encompass a field of possibilities rather than a single fixed recording. In this section, we reflect on the implications of Alternator’s design and its limitations.

### 5.1 Music as Code

Alternator’s premise is distributing music as code (instructions that generate audio) rather than pre-rendered audio. Depending on how composers use this, it opens up exciting possibilities: an enthralled listener who thinks “this is amazing, how did they do it?” can look at the source and start to get a glimpse. Other musicians (typically avid listeners) can go even further, forking the repository and experimenting with it to make it their own. These ideals are shared by other contemporary projects, as in Fluid Music’s

[8] vision for music production and sound design shared as code.

Alternator does not guarantee any of these possibilities. Most music released today is opaque, with alternatives (releasing stems or a DAW project) being the exception rather than the norm. A composer is free to follow this broader trend and distribute their music as an opaque binary blob, as with ordinary closed-source software. Indeed, this may be necessary, at least in part, for proprietary systems such as Max/MSP to be usable in Alternator.<sup>16</sup> Alternator can suggest another way through its design (the “view source” button, the repository link, hosting on GitHub) and practices (releasing the player, templates, and demo albums as open-source), but ultimately the culture will be determined by the artists.

### 5.2 Interactive Music

One intentional limitation is that Alternator only deals with purely autonomous music. It does not handle interactive music, where the output depends on external user input. The reason for this is simple: Alternator is a music player, not an instrument. All autonomous music fits essentially the same interface (samples out), which translates well to the conventional music player of static recordings. In contrast, interactive music comes in many shapes and sizes, requiring different input devices that the user may or may not have, or that might make sense for a listener on a smartphone but not on a laptop. It also expects a different mode of engagement from the listener. Finally, from a technical standpoint, any potential interface for interactive bundles will be more complex and likely to change than the simple `setup()` + `process()` required for generative pieces. Such a platform would thus be more prone to breaking existing compositions, making it less reliable and thus less useful for composers who expect their pieces to remain playable without maintenance.

That said, interactive music is a vibrant area that offers unique experiences impossible with purely autonomous works. It may be worth exploring a separate UI/API in Alternator for interactive compositions, for composers willing to accept increased uncertainty (and maintenance requirements) in exchange for interactivity. The possibility of compositions (interactive or autonomous) generating video as well as audio is likewise intriguing, especially given the large body of work and literature on generative visual art.

### 5.3 Technical Limitations & Future Work

For the goal of stability, composition code would ideally be totally self-contained WebAssembly, with as little API surface area shared with the browser as possible. Currently, there is some JavaScript “glue code” in each bundle that runs in a Web Worker. A well-meaning composer could exploit this to fetch external data (e.g. the current weather) over the network for their compositions—rendering the composition broken if the external resource should move or disappear. A less-scrupulous composer could exploit this to e.g. run a browser-based cryptominer using spare cycles.

However, switching to a self-contained (no JS) Wasm model has some drawbacks. In the current version of Alternator, support for compositions written in JS is “free;”

<sup>16</sup>Hybrids are possible: the runtime engine may be closed, while the composer’s patch is open.

no Wasm bundle or added VM overhead is required, because browsers already come with JavaScript engines. The Python template also depends on JS glue code to load Pyodide [15]. These issues are surmountable, but it may be a mistake to prioritize these changes that (at least in the short-term) constrain Alternator’s features over other work.

Another issue is that Alternator currently uses a “statically-linked” model for bundles, in which they are expected to come with everything they need to run. This means that every piece using Pure Data comes with its own Wasm-compiled version of libpd. To reduce bundle size and load times, it may be worth allowing composers to simply reference common templates (provided by Alternator) rather than always bundling them with the composition. At a minimum, the player could cache `.wasm` files in `localStorage`.

Finally, a major limitation is lack of tooling. Currently the only way to determine if a piece will play smoothly on a given device & browser is to try it and see. Profiling tools and benchmarks on common devices could give composers more confidence in distribution. Also, the existing workflow for creating bundles is straightforward but technical. It would be simple enough to make a graphical, web-based version of the bundler (possibly embedded in the Alternator client). It would be better to fit into artists’ existing workflows, integrate with DAWs, and provide access to generative possibilities without requiring composers to code. We intend to explore these challenges in future work.

## 6. CONCLUSION

Behind all the technical details and design discussion, a question lurks in the background: will Alternator, or something like it, have an impact? Will generative music reach a wider audience? Is there an appetite for distributable music that changes?

The popularity of dynamically-generated playlists (as on Pandora and Spotify) and endless YouTube music streams<sup>17</sup> suggests an appetite for familiarity with variety: *give me more like this, but different*. Generative approaches in other media have proven popular: procedural generation is used in games such as Minecraft, Dwarf Fortress, and countless roguelikes to make the game more fun and improve replayability.<sup>18</sup> And musicians, composers, and songwriters—including those in the “mainstream”—are not shy to try out and adopt new technologies on their own terms.

Ultimately, the appeal of generative music, like any music, depends on its content. Much computer music has historically been “pure” computer music, generated in-the-box with techniques, timbres, and textures that set it far apart from its contemporaries. Or else it has stuck firmly to certain genres, such as electronic and ambient music. We have the opportunity to broaden the horizons of computer music, blur the boundaries, and explore hybrids: music with both conventional and computational aspects. For example, a pop musician might record several good takes in a DAW, and then create a computational bundle that chooses a path through them dynamically at the time of playback. Uses like these may serve to bring generative music to a wider

audience and enable more artists to take advantage of its possibilities, and we hope to explore the space of generative hybrids in future work.

As Eno put it in the essay we started with [5], “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?’” Realizing this future requires bridging the gap between the composer and would-be listener of generative music; it requires a distribution channel which opens up the possibilities of generative music to both.

## 7. REFERENCES

- [1] A. Bainter. Generative.fm. In A. Xambó, S. R. Martín, and G. Roma, editors, *Proceedings of the International Web Audio Conference, WAC ’19*, page 148, Trondheim, Norway, December 2019. NTNU.
- [2] P. Burk. Jsyn - A real-time synthesis API for java. In *Proceedings of the 1998 International Computer Music Conference, ICMC 1998, Ann Arbor, Michigan, USA, October 1-6, 1998*. Michigan Publishing, 1998.
- [3] J. Cage. Composition as process: indeterminacy. *Christoph Cox, Daniel Warner, Audio Culture: Readings in Modern Music*, pages 176–187, 2004.
- [4] I. Clester and J. Freeman. *Composing the Network with Streams*, page 196–199. Association for Computing Machinery, New York, NY, USA, 2021.
- [5] B. Eno. *A Year with Swollen Appendices: Brian Eno’s Diary*. Faber and Faber, 1996.
- [6] K. Essl. Algorithmic composition. In N. Collins and J. d’Esquivan, editors, *The Cambridge Companion to Electronic Music*, Cambridge Companions to Music, page 107–125. Cambridge University Press, 2007.
- [7] J. Gorson, N. Patel, E. Beheshti, B. Magerko, and M. Horn. Tunepad: Computational thinking through sound composition. In *Proceedings of the 2017 Conference on Interaction Design and Children, IDC ’17*, page 484–489, New York, NY, USA, 2017. Association for Computing Machinery.
- [8] C. J. Holbrow. *Fluid Music*. PhD thesis, Massachusetts Institute of Technology, 2021.
- [9] F. Hollerweger. Streaaam: A fully automated experimental audio streaming server. In *Audio Mostly 2021, AM ’21*, page 161–168, New York, NY, USA, 2021. Association for Computing Machinery.
- [10] D. Iglesia. The mobility is the message: The development and uses of mobmuplat. In *Pure Data Conference (PdCon16)*, New York, NY, USA, 2016.
- [11] B. Magerko, J. Freeman, T. Mcklin, M. Reilly, E. Livingston, S. Mccoid, and A. Crews-Brown. Earsketch: A steam-based approach for underrepresented populations in high school computer science education. *ACM Trans. Comput. Educ.*, 16(4), sep 2016.
- [12] J. Maloney, M. Resnick, N. Rusk, B. Silverman, and E. Eastmond. The scratch programming language and environment. *ACM Trans. Comput. Educ.*, 10(4), nov 2010.
- [13] M. S. Puckette. New public-domain realizations of standard pieces for instruments and live electronics. In *Proceedings of the 2001 International Computer Music*

<sup>17</sup>Such as the eternal “lofi hip hop radio - beats to relax/study to”: <https://www.youtube.com/watch?v=5qap5aO4i9A>

<sup>18</sup>Indeed, some games feature generative soundtracks. In this sense, the most successful distribution platform for generative music thus far might be Steam.



*Conference, ICMC 2001, Havana, Cuba, September 17-22, 2001.* Michigan Publishing, 2001.

- [14] S. Reich. Music as a gradual process. *Writings on music*, 2000:34–36, 1965.
- [15] The Pyodide development team. pyodide/pyodide. <https://doi.org/10.5281/zenodo.5135072>, July 2021.
- [16] E. Scheirer. The mpeg-4 structured audio standard. In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP '98 (Cat. No.98CH36181)*, volume 6, pages 3801–3804 vol.6, 1998.
- [17] G. Wang. A history of programming and music. *The Cambridge Companion to Electronic Music*, pages 55–71, 2007.
- [18] S. Yi, V. Lazzarini, and E. Costello. Webassembly audioworklet csound. In J. Monschke, C. Guttandin, N. Schnell, T. Jenkinson, and J. Schaedler, editors, *Proceedings of the International Web Audio Conference, WAC '18, Berlin, Germany, September 2018.* TU Berlin.