

PC: With the many ways artists are able to work now in digital environments beyond coding, what have you considered is the most enticing way people are interacting with music visually, in concerns to production? Have you ever considered other options beyond coding your visuals?

SAM: I live in relative isolation—no television, no newspaper, only subscribe to one technical periodical, don't surf the web much, don't play video games, don't attend many concerts or other art-related events. So, I'm afraid I don't know much about how people are interacting with music visually these days. What little I have seen has not enticed me particularly. As part of working with Björk in 2011, I saw a video in which she played a musical instrument with a visual interface; it was cool to watch, but it didn't strike me as very interesting musically. And I've seen some music apps for the iPod and iPad that I thought were clever and amusing—and fun—but again, I didn't come across anything where my reaction was "I want to spend some serious time with this." (Am I becoming a sourpuss in my old age?)

As for my own work, there are lots of directions I'd like to explore. One that I've done a little work on in the past and think has a lot of potential is: interfaces for expressive performance. In the pre-electronic era, the physical design of musical instruments was constrained by acoustics: you couldn't produce a bassoon-like sound without having an instrument very much like a bassoon. In some cases, like with the pipe organ, it was possible to isolate the performer from the sound-producing parts of the instrument, but only in the piano was a significant expressive control—dynamics that could be changed note-by-note—provided in a way that was decoupled from the physics of the sound production (the hammers hitting the strings). But now, we can configure any physical interface to produce any sound. I started thinking about this in the context of the conductor program (<http://www.musanim.com/tapper/>), and I did some experiments (some working prototypes, some mock-ups, some conceptual work), but only scratched the surface.

One type of visual interface for music that I'd like to develop is one in which the instrument is dynamic—that is changes depending on how it's used. For example, if you were playing a piece that had repeated notes, you might want to build a "virtual plectrum" that could pluck each note multiple times with a single gesture.

Or, the instrument might change in response to tonality. In the Harmonizer (<http://www.musanim.com/harmonizer/>), pitches are arranged in a pattern that reflects tonal relationships; in this mock-up

https://www.youtube.com/watch?v=O_7DoDCRiuM

a design that's similar to that changes in response to the pitches that are being played: the more you play a particular pitch, the bigger its "key" in the instrument becomes. In this way, the instrument itself "modulates" to the key you're in, as a result of playing it. This could be used in a multi-player context; for example: each player could have their distinctive instrumental sound, but all would be contributing to the tonality picture seen

by all players. You could see, through the instrument, what the other players were doing, tonality-wise, which would make it easier to play in the same key, follow modulations, etc.

PC: One reason I moved to After Effects was because I could not properly render what I created in a coding environment as I did not have the experience to build a way to read and write the MIDI data to and from a file. Most of what I did make had to be made in real-time, which did not, of course, leave any way to show the future of notes while they were being displayed. This was a fairly hurtful limitation as you can imagine. The MAM client had this ability, and I'm curious how you went about it. Did you just create a MIDI wrapper within the program to create the data for reference? Does the program convert it in real time? Did you use something external?

SAM: The first IBM PC version of the MAM was a fully real-time instrument, complete with an editor (see Performance Editor at 11:47 in <http://www.youtube.com/watch?v=yq0iMqQpxbA>). You could play MIDI into it, which it would visualize immediately and record for future use.

By contrast, the version of the software I'm using now is completely based on MIDI files. The animation is produced in real-time, but the MIDI data is all created before the program is run.

PC: Would you mind telling me more about switching to a standing desk? I read in your history that it was enough of a milestone to include there, and I'm curious if you're still for it and why you prefer it. There's a lot of "buzz" about standing desks recently and I'd like to hear from someone that has specific experience.

SAM: I liked my standing desk, but in the end, I found it too tiring; it was a fixed height, so I had to either stand or sit on a high, relatively uncomfortable chair. There are cool versions now that can be raised and lowered electrically (so that you can alternate standing and sitting), but my setup has grown to the point where single desk wouldn't work; this composite photo gives you some idea (it's about 190 degrees left to right):



My fantasy is to develop an interactive, immersive, full-body musical interface—in which you'd perform (conduct, interpret, improvise) by doing something like dancing. I

love playing music, and if practicing were something that were as demanding physically as, say, an obstacle course, I'd be in great shape!

PC: In regards to animations, I've started to draw a line between "academic" representation and "stylistic" representation. While I know a large extent of MAM is to show the relationship between notes academically - and it works in its own way stylistically - have you ever considered what kind of role the stylistic motivations have had on the relation between music and visuals in MAM? One thing I'm trying to experiment with is to find ways to show note expressiveness, slight changes in timbre for example. While I'm far from creating work that exemplifies what I'm after that could possibly be represented in an academic way as well as purposefully aesthetic, I realize you must have had a reason for the way you represented it beyond academically. It's apparent you've started trying to change how the music is represented in your pieces - what is your thought process behind the need for change from the previous? What was your reasoning for the ... "traditional" MAM's aesthetic?

SAM: There are lots of ways to characterize the differences between approaches to (and aspects of) music visualization. For me, the word "academic" has many connotations—formal, theoretical, conventional, impractical, institutional—too many, for my taste. And "stylistic" has problems of its own. So, let me answer using different terms.

One distinction I make is between interpretive and non-interpretive methods of visualization. This is not a 100% clear distinction, since every choice about how to do a visualization is a kind of interpretation, but we can agree that there is a spectrum, and that a single-color bar-graph (piano-roll) version of a MIDI file is doesn't involve as much interpretive activity as, say, the dancing hippopotamuses in *Fantasia*.

When I listen to a piece of music, I am discovering things, and appreciating the music more as a result of what I'm learning about it. What I'm trying to share with the viewers of my videos is that experience—the experience of discovering things about the music. I'm not trying to foist my own discoveries on people. So, even if a piece of music reminded me of dancing hippopotamuses, I wouldn't put them in a video. Rather, I try to make the visualization as neutral as possible; I try to stay out of it. I think of what I'm doing as being like giving the viewer a "musical microscope"—something that let's them see aspects of the music that would otherwise be hidden—and saying "see for yourself."

However, to continue the microscope analogy, there are some aspects of the music that are not obvious, even with a microscope. To bring out features in a specimen, you might cut it thinner, or stain it, or put it between polarizing filters, etc. These techniques are not as heavy-handed as dancing hippopotamuses, but they are still somewhat interpretive. Along these lines, I might choose to use a different shape for different types of instrument (string, woodwind, brass, percussion).

Even something as innocuous as a change to the horizontal scale results in different aspects of the music being highlighted: in a compressed view, you see large-scale features, and the view changes slowly, while an expanded view makes individual notes more dramatic, and gives more of a sense of motion through the piece.

Sometimes, I go further. In my visualization of James Boyk's performance of Debussy's *Reflets dans l'eau* ...

<https://www.youtube.com/watch?v=lJDlqCX8Qrk>

... each note is depicted in a method that's chosen based on the type of figuration it's part of: parallel chords, fast-moving solo lines, slow regular motion, slow irregular motion, sustained notes. There are several kinds of interpretive decision involved. First, there is a kind of "compositional" or "analytical" decision: what kinds of figuration make up the piece? This involves the same kind of thinking that a performer would do when approaching a score—asking "what is this piece about? how is it made? what features are important? what do I want to bring out?"

Then, for each note, there is the matter of deciding what type of figuration it's part of. There isn't always a single, clear, unambiguous answer to this. In some cases, I chose to depict a note more than one way, since it functioned in more than one way. In some cases (though rarely), my decisions were different from those that Boyk made.

And finally, there's the choice of what visualization method to use for each type of figuration. In this, there are both interpretive and practical considerations. For example, sustained notes have an effect for a long time; if you don't reflect their duration visually, then their function is not clarified/enhanced/reinforced by the visualization. Or, for notes that are fast-moving in a complex passage, it may be hard to see the motion unless you draw lines between notes to show the connections and the trajectories. Those are practical considerations, and one could, in theory at least, make such decisions algorithmically (based on principles of cognitive psychology). But these decisions can spill over into the purely interpretive; for example, I might think "this piece is very relaxing, and feels like flowing water" and pick visualization methods that express that.

Where do I draw the line, and why? I don't have a simple or complete answer to this (and even if I did, it would be changing from day to day), but one important part of the answer is: I am a composer. I don't actually write music any more, but I still think like a composer, and I am, for better or worse, "on the composer's side" in my approach to music visualization: my main focus is on what the composer did, helping the viewer perceive and appreciate that. I don't want to add things to my visualizations that might distract the viewer from what's going on in the music.

I'm not saying that there isn't a place for a more interpretive approach to visualization—there most certainly is. Just as a ballet can be added to a piece of music, music visualization can add its own layer of content, expressiveness, beauty and meaning. But, for the moment at least, that isn't my job. I want to let the music speak for itself.

PC: What is your opinion of most of the work being done in music visuals? While I'm starting to see many interesting usages, more and more involving motion graphic oriented software and MIDI, most of what I come across is usually fairly simple changes to the methods you used (the scrolling bars, pitches linearly distributed, etc) or particle-laden, ambiguously synced visuals (such as those found in default music visualization options for some audio playing software). Would you mind sending me a few examples of where you think music visualization is heading? What kind of work really interests you?

SAM: With music visualization, as with music, my response is often more from the point of view of a composer than as a consumer or an historian. The result of this is that when I see something completely satisfying, my reaction is: that's perfect—no more can be done there. For example, Bach's *Air* ("on the G-string"). When I listen to that, I marvel, I am moved, I am in awe. I don't think "where does this suggest music is heading?" Nor do I think "I want to write something like that." Bach did it; it's done; I'm not needed here. On the other hand, when I attend a concert of contemporary music, my reaction is almost invariably "good idea, but the composer didn't take very good advantage of it ... now, if it were *my* piece ..." and I'm back into composer mode.

And the same with music visualization. My reaction to all the work I see is "that's neat, but if I _____ (fill in: had the chops, had the tools, had the time, had the inclination to do something like that), I'd _____ (fill in: whatever the visualization suggested to me)." For example, my reaction to Esteban Diácono's wonderful *Ljósið* ("Let Yourself Feel") was: I wish I could drive that kind of animation with my data. My desire to do that came both from the beauty of his images and the fact that the forms he was using had so many axes of freedom—that I thought could be tied to aspects of the music. Is Diácono's work "where music visualization is heading"? I don't know.

Because I'm a composer, I'm more interested in tools, techniques and ideas than in any particular type of work. When I see something new, my attention is on the part that I've never used myself, and I'm thinking "what is this? what is its essence? what could it be used for?" If the piece that it's part of is good, then I am moved and full of admiration, but that doesn't translate directly in interest.

PC: Would it be possible to request a MIDI and piano recording for one of your pieces so as to have a better position for experimentation? I am not much of a classically trained musician, so my ability to create the needed subtle nuances in audio to find ways to utilize more extensively through MIDI is meek, and I'd much rather have a classical piece to test with (When I started in Processing I had used a Vivaldi Concerto No. 4 MIDI extensively for the range it offered for testing, which you can see here with a hastily applied non-MIDI soundtrack). I would of course give you full due credit (and link musanim.com) if I chose to publicize the results of any experimentation online. If you'd permit me to request, I'd pick Chopin's Nocturne in E-flat Major, opus 9 no. 2.

SAM: I'm going to give you the stock "give a man a fish" answer: I'd much rather help you develop the skills to do this yourself. Let's discuss this off-line.

PC: How have you considered the MAM visualizations in regards to showing beat? More specifically, have you considered pieces that are accompanied with strong rhythms, or perhaps "popular music" and drum sets? Or perhaps more interestingly, say a drum set in a jazz piece - what thoughts have you had in response to how something like a strong beat would be used graphically? It seems like the majority of music videos (in pop music anyway) tend to emphasize the beat more than anything else, which makes sense for that medium, but beyond syncing movement to a beat and instead representing it graphically, say in MAM, what considerations have you come across?

SAM: I haven't seen anything that addressed this directly. It was one of the first things I thought about once I'd done my first bar-graph scores, and I'm still thinking about it. I've done a few experiments (e.g. see "Rhythm" at 22:00 in <http://www.youtube.com/watch?v=yq0iMqQpxbA>), and I've added barlines to some of my videos, but I don't think that anything I've done yet is significant. I have a better sense of the way forward on this than I used to, and so whereas a few years ago I would have said "I have no idea how to approach this," I wouldn't say that now. But I'm not ready to work on it yet; I'm waiting for some idea to gel into a form that's simple enough to implement.

What I appreciate now (that I wasn't as clear about before) is the importance of expectation and anticipation in our experience of rhythm. Having a light flashing in time to a musical rhythm may reinforce the beat, but it does not explain it—which is what I'd want to do. In my visualization of Tenny's percussion piece *Wake*, I took a step in this direction, but it was an almost completely cerebral (abstract, geometrical) explanation. For a visualization of rhythm to work the way I want, it needs to both show aspects of the structure that you might not otherwise understand, and make you feel like moving in a way that expresses (or would result in the performance of) that structure.

PC: You can see some of what I've been experimenting with on Vimeo. I am currently trying to find a way to work on new music with visuals that are designed synchronously, though in different environments so each can have their full potential, so as to tighten the relationship of the two. As I'm working towards more digital production musically I feel this would be the best way to tightly created synchronization. Before that happens I'm spending some time researching this subject more fully, as I felt my last project became too separated from trying to create a link between music and visual and that's what I'd like to experiment with moreso.

SAM: On the subject of rhythm, your video "Midi Animation Test #09 - Vivaldi, Violin Concerto No. 4, parts 1 and 3" (<http://vimeo.com/30363314>) does something with rhythm

that's related to what I was saying about rhythmic structure, expectation and anticipation. In the animation, the typical phrase length is equal to the length of time it takes for the notes to make one revolution on the circle. When the texture changes from sparse to dense, the "head" of the dense passage, its "leading edge," functions as a visual object, which the viewer can track around the circle. When it arrives at the "now" point, that phrase has completed. This explains something to the viewer (there is a length of time, a unit, that is a fundamental aspect of this piece, and that passage was one unit long), and sets up an expectation: since the passage is over, whatever happens next is going to be something else. It might be a repetition, it might be a change, but regardless of what it is, it is another chunk, another "unit" of the piece.

PC: With the prevalence of work done with visuals without MIDI guiding it, and the specificity inherent in MIDI, would you consider merit in something along the lines of a digital file format that includes MIDI tied into it so as to automate visuals with more direct expressiveness? (Disregarding issues of having a "source code" to the music included with the music itself)

SAM: Sure. In fact, there are lots of people who are doing that. For example, lighting equipment can be controlled by MIDI, which means that people can create MIDI files to "express" the lighting changes for stage shows.

PC: I can imagine a future where automation of direct musical data would be fantastically entertaining, though at the same time it's obvious that developers are finding new ways to directly utilize audio in more specific ways. Do you feel MIDI should be utilized more - and if so why is it not? Or have you seen other options that intrigue you? What limitations and benefits do you feel MIDI usage incurs? Ultimately do you feel that engineers will be able to create automated visuals based off of audio alone that will be comparable to the expressiveness that MIDI offers?

SAM: MIDI files contain a rich combination of information that's in a score and information that's in an audio file. It is precise in many of the ways a score is precise (pitch is precise, instrumentation is precise) and in many of the ways an audio recording is precise (timings are precise, dynamics are precise), and it has certain advantages over both scores and audio: it is simple and compact, and its format is well-defined (not to mention widely used). To be sure, it has limitations and deficiencies, but to do better, you'd want to include at least as much as what MIDI has in it.

Polyphonic pitch detection, functionality upon which an audio-to-MIDI (or audio-to-somethingLikeMidi) depends, is still in the early stages of development (coincidentally. I worked on it briefly at the beginning of my employment at Audience, Inc.). As engineering challenges go, it is in the "hard but not impossible" range, and the reason that it isn't as far along as, say, speech recognition, is that relatively little work has been done on it. This means that it's only a matter of time.

What is much more feasible, technically—and would give you the kind of data which you'd ideally want (though not in an automated way)—would be to develop software that took a MIDI file, a corresponding audio file (which would be assumed to have most if not all the notes in the MIDI file), and some simple "hints" (like information about what instruments are actually present in the recording, and generated from these a rich representation of the music, one that included all the data in the MIDI file, together with a complete description of timbre, spatial attributes (i.e. stereo position), etc.