*Glimmer* for chamber orchestra and audience

Jason Alexander Freeman

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Musical Arts
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2005

ABSTRACT

*Glimmer* for chamber orchestra and audience

Jason Alexander Freeman

*Glimmer*, a composition for chamber orchestra and audience, uses novelty light sticks, video cameras, computer software, multi-colored stand lights, and projected video animation to create a continuous feedback loop in which audience activities, software algorithms, and orchestral performance together create the music. This thesis establishes the aesthetic background and motivations behind *Glimmer*, explores how the constraints imposed by the orchestral performance environment influenced the design of the work, and evaluates *Glimmer*'s successes and failures in meeting its goals.

TABLE OF CONTENTS

VOLUME I (COMMENTARY)

LIST OF FIGURES

ACKNOWLEDGEMENTS

*Glimmer* was commissioned by the American Composers Orchestra, Steven Sloane, music director, Robert Beaser, artistic director, Dennis Russell Davies, conductor laureate. This work was also funded in part by the Composer Assistance Program of the American Music Center and the Creative Connections Program of Meet the Composer.

Most of *Glimmer* was written while in residence at Akademie Schloss Solitude (Jean-Baptiste Joly, director), and I am grateful for the opportunity to work in such a supportive environment. I am also indebted to the Columbia Computer Music Center (Brad Garton, director) for loaning much-needed equipment and studio space for the project, and to the Fritz Reiner Center at Columbia University (Fred Lerdahl, director) for its partial support for video documentation of the premiere performance.

*Glimmer* is an enormously complex piece, and there are many additional people whose ideas, support, and hard work made it possible: at Akademie Schloss Solitude, Kristjan Varnik, Sekhar Ramakrishnan, Rodney LaTourelle, and Hannes Gnad; at Columbia University, Douglas Repetto, Terry Pender, Luke Dubois, Michael Klingbeil, Keith Moore, Paul Hogan, Hiroya Miura, Brahim Kerkour, Joe Dubiel, and George Lewis; at the American Composers Orchestra, Alex Johnston, Michael Geller, Daniel Brodney, Jeff Milarsky, Debra Kreisberg, Keith O'Brien, Barbara Burch, and Paola Prestini; at Carnegie Hall, Tom Eirman; at Color Kinetics, Brian Chemel; at iPROMOTEu, Mary Staubitz; and my family, Stephen and Sheila Freeman, Brian Bosworth, Lauren Freeman-Bosworth, and Andrew Bosworth.

in memory of Jonathan Kramer

# INTRODUCTION

Strange things were happening in Zankel Hall. Audience members were laughing and cheering and doing the wave while wielding novelty light-up cocktail stirrers. The orchestra was playing, without a conductor, while lights on their music stands flashed, faded in and out, and changed color. A video screen showed an animation of the audience seating area, with a white circle occasionally jumping from one section to another. And in a control room backstage, four laptop computers were keeping watch over it all. What was everyone doing? Where was the music coming from?

Strange things had actually been happening for a while. One afternoon a few months earlier, I spent a quiet afternoon in the Carnegie Hall house manager's office, flipping through a thousand pages of architectural plans for Zankel Hall in search of some measurements. I needed to construct a three-dimensional computer simulation of the hall to test camera placements and lens angles, since my testing time in the actual hall was so constrained. How did I get here? And why were courses in architecture and closed-circuit television systems not part of the composition curriculum?

Of course, both of these events were connected to the same piece of music. *Glimmer* (2004), for chamber orchestra and audience, engages the concert audience as musical collaborators who do not just listen to the performance but actively shape it. Each audience member is given a battery-operated light stick which he or she turns on and off over the course of the piece. Computer software analyzes live video of the audience and sends instructions to the orchestra via multi-colored lights on each player's stand.

This brief description merely hints at the underlying details, which are thoroughly explored in the accompanying score. If you are not yet familiar with *Glimmer*, I encourage you to peruse the score before reading this document.

This paper does not describe what *Glimmer* is and how it works. Instead, it seeks to explain where the piece came from and to understand what transpired at its premiere. In the first chapter, I outline my motivations for creating interactive music in general and *Glimmer* in specific, considering examples of works by myself and others within the framework of software art. In the second chapter, I consider the context for which *Glimmer* was created — an orchestral concert — and the ways in which that context, along with my own high-level goals, influenced the design and realization of the piece. In the third chapter, I evaluate the successes and failures of the premiere performance and suggest ways in which the experience could be improved if the work were performed again. In the last chapter, I offer some final thoughts about the goals of interactive music and my interest in pursuing it.

To explore *Glimmer* even further, you can peruse additional performance materials online at http://www.jasonfreeman.net; the Max/Jitter source code and the supplementary video simulations referenced in the score are available for download. Unfortunately, stipulations of various unions at Carnegie Hall and the American Composers Orchestra prohibit me from making the video of the premiere performance available, but you may borrow an archival copy of the video by contacting me through the web site.

# 1. Background and Motivations

## 1.1 Composer, Performer, and Listener

> Tellers of stories with ink on paper, not that they matter anymore, have been either swoopers or bashers. Swoopers write a story quickly, higgledy-piggledy, crinkum-crankum, any which way. Then they go over it again painstakingly, fixing everything that is just plain awful or doesn't work. Bashers go one sentence at a time, getting it exactly right before they go on to the next one. When they're done they're done. (Vonnegut 1997: 137)

When writing music, I have always been a swooper. I get the first draft out of the way as quickly as possible so I can get to my favorite part: editing. For me, editing is more than just tidying up; it is a process of gradually discovering the music I want to hear and determining how best to notate it. I imagine hearing what I have written, then critique and revise it, then imagine hearing the new version, then critique and revise it again, and on and on until I am satisfied. The final score often bears little resemblance to the first draft.

Many performers do something similar when they practice. In the initial swoop, they focus on basic technical mastery of a score. Then comes the editing: they play the piece or small passages from the piece over and over again, each time critiquing their interpretation and revising their approach.

And even listeners can be swoopers. As they listen to a piece of music for the first time, they quickly develop ideas about what they are hearing — whether those ideas be subconscious expectations of what is to come or full-blown theories explaining their experience as a listener. And as they continue to hear the piece, or as they hear it again, or as they remember hearing it, they constantly evaluate and revise those ideas.

All three of these editing processes are feedback loops which gradually transform ideas over many iterations. They often operate independently of each other, even though

they are all linked to the same piece of music — like three ice skaters each circling a rink on a different day. Figure 1 shows a simple representation of these processes.

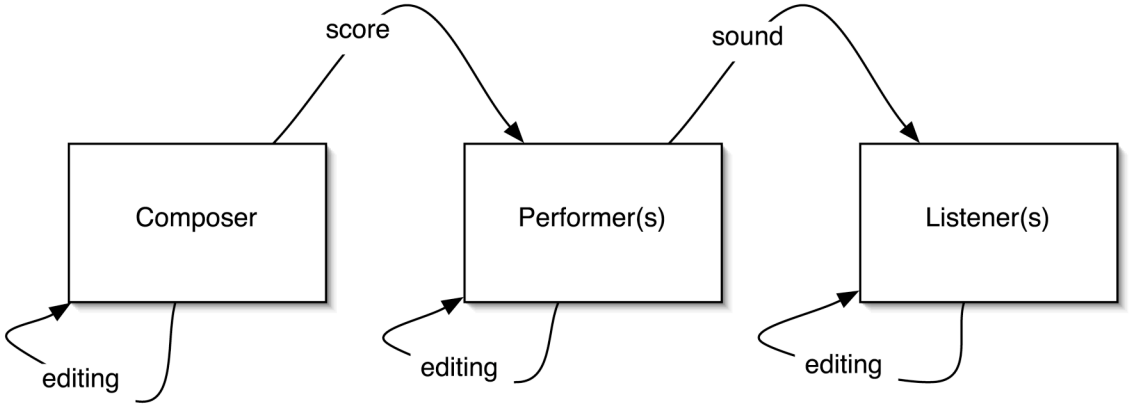


**Figure 1.** One type of relationship between composers, performers, and listeners.

The independence of these processes is often inevitable. Imagine, for example, that you are listening to an old historical recording of a Beethoven symphony. The composer wrote the score, *then* the performers interpreted the score, *then* you reacted to the performance. There is no way that your activities as a listener could influence the composition or the performance to which you are listening.

But now, imagine that you are sitting in a concert hall, listening to the world premiere of a new orchestra piece. The composer, the performers, and the listeners are all sitting in the hall together, yet their activities remain largely independent. The composer probably finished the score weeks or months before the concert; the performers rehearsed it beforehand as well. It is mostly just you, along with the rest of the audience, who is actively developing ideas about the music as you hear it for the first time.

You do get a chance to respond to the orchestra and the composer, applauding them and perhaps even talking to them, but these opportunities come too late to affect the

performance or the score. Short of an errant cough or cell phone ring or a riot, the musicians neither see nor hear you as they perform: the audience sits quietly in darkness, and the players focus their attention on the conductor, who has his back turned toward you.

John Cage put it this way: "Composing's one thing, performing's another, listening's a third. What can they have to do with each other?" (Cage 1961: 15).

Figure 1 does not accurately describe all of the ways in which we experience music; improvisation, for example, cannot be so neatly categorized into these boxes. But it does capture many of our experiences with music. And while there is nothing inherently wrong with this design, one of my primary interests in recent years has been to rethink and restructure these relationships. The following sections describe some motivations and strategies for doing so.

### 1.2 Drawing the Line Between Composition and Performance

"I'm not good at telling people what to do." — Terry Riley (Strickland 1991: 120).

Think again about listening to an historical recording of a Beethoven symphony. Because composer, performers, and listeners are separated from each other in time, there are limited channels for communication from one stage to the next. The composer communicates with the performers via a written score, and the performers communicate with listeners via sound. Since no one can ask Beethoven how he wanted the oboe to play measure 63, the score becomes quite important.

From early in my compositional training, my teachers emphasized the importance of that score. They urged me to make my notation as clear and precise as possible,

writing down every detail and leaving little room for doubt and ambiguity. I am grateful for this training, because it helped me learn to develop musical ideas clearly in my own head and to express them clearly to others through notation.

But sometimes my attention to detail has gone too far. Figure 2 shows an excerpt from my piece *Pantoum* (2000), for violin, viola, cello, clarinet, and trombone. Practically every note has an articulation and dynamic assigned to it, and measures 23-25 nest two levels of durational information: tuplet types change during an *accelerando* and *ritardando*. Yet the entire section is marked "Freely (quasi-cadenza)." How can a clarinetist interpret this passage freely when so little is left to her interpretation?


**Figure 2.** Excerpt from *Pantoum* (2000), for violin, viola, clarinet, cello, and trombone.

The answer, at least for the clarinetists in Speculum Musicae, the Columbia Composers Ensemble, and Le Nouvel Ensemble Moderne, was to ignore some of the

details of the written notation when it conflicted with their freer interpretations. Strangely enough, these deviations from the score have never bothered me, which forces me to question why those markings exist at all.

At the time I wrote *Pantoum*, I thought that this level of detail proved to myself and to others that I was in complete control of my materials and thus demonstrated that I was a good composer. The contemporary music community largely supports that logic, as evidenced by a recent *New York Times* review titled "Atonal, Edgy, and In Control":

> During an informative preconcert discussion with the composer Augusta Read Thomas…[she said]…"If you hate everything you hear…if you leave thinking 'What a crazy lady,' I hope you will at least say that 'we know she heard what she wrote.'"
>
> Ms. Thomas should rest assured…it's hard to imagine that anyone who heard the four works performed…left the hall doubting that Ms. Thomas was in control of every nuance in these vividly colorful pieces. (Tommasini 2005)

I no longer believe that such nuanced control is necessary or sufficient to demonstrate compositional skill; in fact, it is not always desirable to be in complete control of one's materials. For me, the level of specificity of a score has become just another compositional parameter: something which can change from piece to piece, or from moment to moment within a piece. The line where composition ends and performance begins need not be frozen.

In my piece *Prior Art* (2004), for chamber ensemble, the conductor and the players make important decisions regarding the timing of events and the pacing of the music. The score alternates among three different notational styles, regulating the ensemble's interpretive freedoms. Some of the music is notated in regular meter with precise metronome markings, leaving room for only small variations. Some of the music is notated with indeterminate measure lengths; the conductor gives only downbeats, and

the length of measures varies within a specified range (e.g. 2-5 seconds). And some of the music is notated in independent time; the conductor gives sporadic entry cues at time intervals entirely of her choosing, and musicians play independently of the rest of the ensemble once they enter.

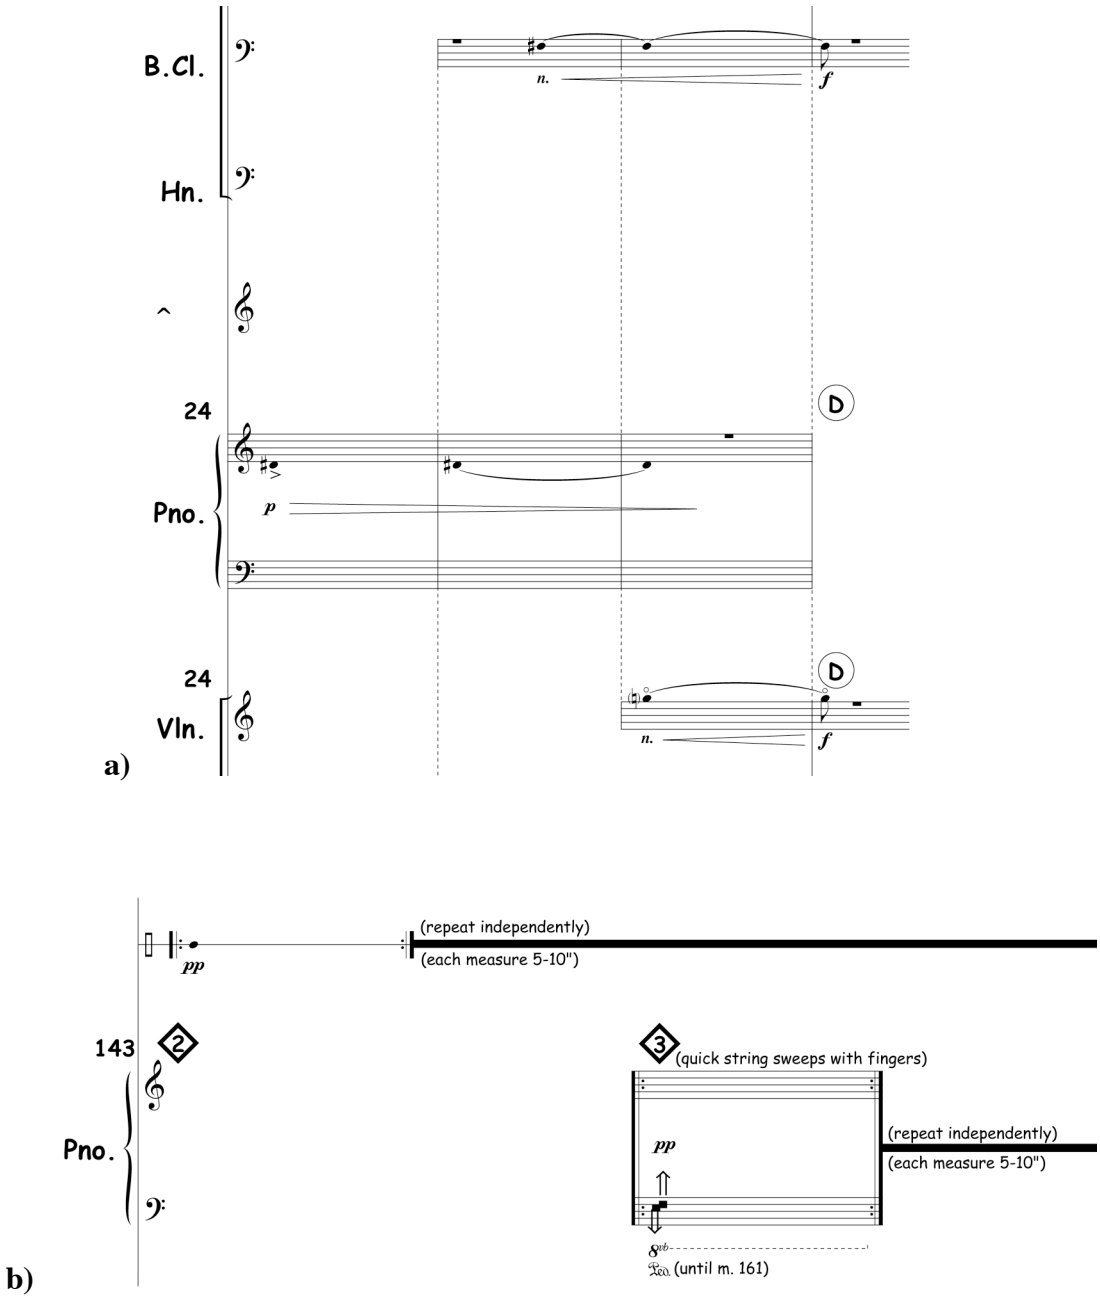


**Figure 3.** Score excerpts from *Prior Art* (2004) showing **a)** indeterminate measure lengths and **b)** independent time.

These notational techniques are nothing radical, as one glance at *Notations* (Cage 1969) would show. But the notation is not the point. It is merely a tool which enables me to give performers the flexibility to respond in the moment of performance — to the acoustics of the space, to the resonance of their instruments, to the arc of their gestures, to the length of their breaths — and to allow the music to develop fluidly and organically. It enables me to move the line between composition and performance but to remain clear and precise about the location of that line.

The location of that line in my scores is ultimately about risk and trust. By giving performers greater interpretive freedom, I define a larger space of possibilities for them to explore, and I trust their musical skills and intuition to guide them through that space. Sometimes performers surprise me with wonderful decisions which take the music in directions I had not envisioned. And sometimes performers disappoint me with decisions I dislike. But if I am not willing to take that risk, then why would I write music for human performers at all?

### 1.3 Listening and Participating

Listeners, like performers, can make interesting and creative contributions to a musical performance when given the opportunity, taking the music in directions I alone could not have envisioned. I am also driven to involve listeners by a kind of evangelical motivation to encourage them to approach music in different ways and to experience the process of creating and performing music firsthand.

### 1.3.1 Reasons for Listening

In our daily lives, most of us have become accustomed to hearing music which functions as background to something else: the muzak on the elevator or on the telephone, the soundtrack accompanying a film or television show or advertisement, the tracks on our iPods being drowned out by the subway. It is unusual for music to be the sole focus of our attention.

For some people, listening to music on its own can be difficult and disorienting; they grow bored and turn the music into background for daydreams, missing so many of the rewards which it can offer to more attentive listeners. But by drawing listeners into active creative roles, works engage them more directly and demand more of their attention. They also give listeners a new reason to listen closely, since they have a personal stake in the sounds they hear.

### 1.3.2 Personal Reflection and Shared Experience

For me, composition is a form of reflection about important experiences in my life. For example, *Prior Art* (2004) was influenced by an experience I once had playing an arrangement of Verdi's overture to *La Forza Del Destino* (1861). As a saxophonist attending a small high school with an even smaller band program, my opportunities to play in large ensembles were mostly limited to band festivals. At one such festival, I was shocked to hear the opening measures of the Verdi during our first rehearsal. A full concert band's brass section was new to my ears, and I was so taken by the sound that I missed my entrance. I thought the conductor was giving each beat, but in fact he was only

giving the downbeats to each measure. Part of me wished I had been correct, so that the sound I enjoyed so much would have lingered a bit longer.

The opening measures of Verdi's overture are a prominent motive in *Prior Art* (see Figure 4), but I would never expect listeners to connect my piece with Verdi's, let alone with my personal childhood experience. In fact, I do not even mention Verdi in my program note for the piece, dwelling instead on a story involving lost sheep. Just as my own experiences influenced the score I created, I expect listeners to understand the music in relationship to their own experiences.

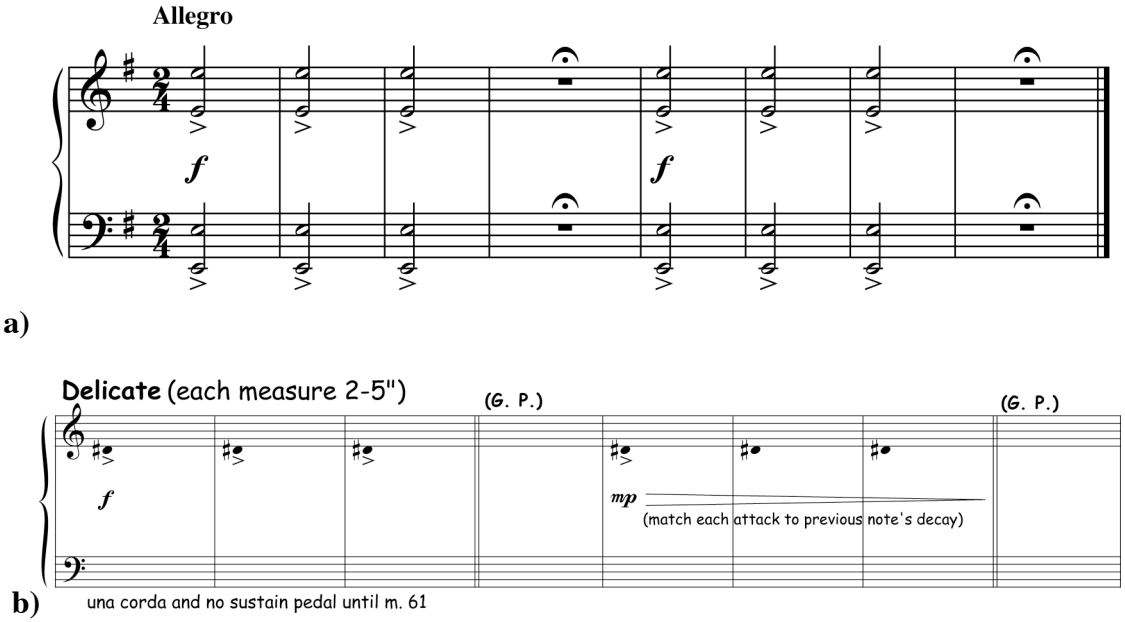


**Figure 4.** Opening measures of **a)** Verdi's overture to *La Forza Del Destino* (grand staff reduction) and **b)** *Prior Art* (piano part).

But sometimes, that is not enough. Much of my music, including *Prior Art*, reflects on my experiences creating and performing music. What made those experiences special was not the musical product, but the process by which it was produced. Yet many listeners have little experience creating or performing music themselves. Tod Machover explains it this way:

> …music's around all the time, but fewer and fewer people actually participate in it themselves. It's not a mystery to say there's some disconnect there, and anything we can do to make those ends meet is I think a really good thing. (Oteri 1999)

If I can involve them in my own creative process, I can create an experience for us to share. Then composition is no longer just a personal, private reflection. It is a transformation of a personal reflection into a shared experience.

### 1.4 The Challenge of Involving Listeners

Listeners may lack musical training, skill, and talent, but they can still contribute meaningfully and creatively to a musical experience. Charles Ives explains with a story about his father:

> Once a nice young man...said to Father, "How can you stand it to hear old John Bell...sing?" (as he used to at Camp Meetings) Father said, "He is a supreme musician." The young man (nice and educated) was horrified — "Why, he sings off the key, the wrong notes and everything — and that horrible, raucous voice — it's awful!" Father said, "Watch him closely and reverently, look into his face and hear the music of the ages. Don't pay too much attention to the sounds — for if you do, you may miss the music." (Kirkpatrick 1972: 132)

Few people, however, possess the self-confidence to sing like John Bell. Listeners who are not trained performers, cannot play an instrument, and do not read music usually have little faith in their musical intuition. Especially when surrounded by their peers, they become too shy to contribute creatively to a musical performance.

I witnessed such fear at a piano recital by Kathleen Supové at the Cutting Room in 2001. During an improvised cadenza in Frederic Rzewski's *Which Side Are You On?* (1982), Supové invited members of the audience to come onstage and improvise along with her at the piano or with their voices. Even though the audience consisted primarily of professional musicians, none of us dared go up on stage. In the final moments of the

cadenza, a poet finally came up and spoke a few words. But no one else accepted the invitation to join the performance.

At the opposite extreme, listener participation can be disappointingly shallow if it is not structured carefully. I learned this myself in 2001, when I was commissioned to write a piece for the fourth-grade band, chorus, and orchestra at an elementary school in Richmond, Virginia. As part of their "Be a Composer" project, I did a short residency, meeting with all of the students (whether musicians or not), presenting my work to them, and soliciting ideas from them for the piece. By actively involving them in my compositional process, they would be motivated to listen for the realization of their ideas, and they would share the experience of writing a piece of music.

I had never done anything like this before and took a naïve approach. At the end of my meeting with each group, I simply asked: "What ideas do you have for the piece?" There was never a shortage of ideas, but they were not particularly helpful. Most were hopelessly vague: one student would say he wanted a loud piece, then another student would raise her hand and say she wanted a soft piece, and then I would suggest a piece that was loud in some parts and soft in others. The class would be delighted, but I felt like an idiot. After all, I probably would have written loud and soft passages without the benefit of their advice.

Those suggestions which were not hopelessly vague were not really about music. The school had just completed a unit on rainforests, so I received many requests for rainforest sounds and jungle animals. I complied by using a text about the jungle, a few special percussion instruments, and an occasional animal howl. The students were

delighted that they had helped me to write the piece, but I felt like I had cheated them. How much did they really contribute, and how much was just an illusion?

This problem, needless to say, is not limited to elementary school students. Musical scholars often struggle to effectively describe music through language. For a lay public without knowledge of a specialized vocabulary, the task is even more difficult. So discussion turns toward vague generalities and extramusical elements.

## 1.5 Technology and Software Art

For me, technology offers a powerful means to address some of the challenges of soliciting and incorporating the creative contributions of listeners. It is a tool through which to facilitate collaboration without resorting to conventional musical instruments or descriptive language. With technology, I am able to create simple interfaces through which people can address musical content and structure and through which they can express their creative ideas.

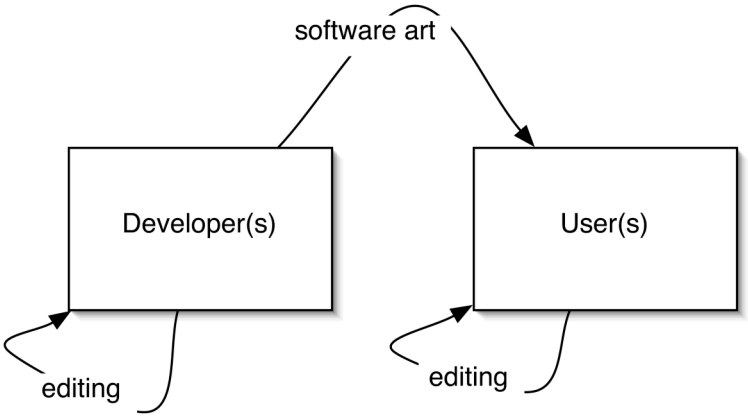


**Figure 5.** The relationship between developers and users in software art.

To effectively describe the use of technology in my music, I often discard the notions of composer, performer, and listener altogether, thinking instead about a software developer and software users. Instead of a score, the developer creates software which processes data from input interfaces and sends the results to output interfaces. Instead of listening to a performance, users operate those input interfaces and respond to the output interfaces, creating a feedback loop of interaction between them and the software.

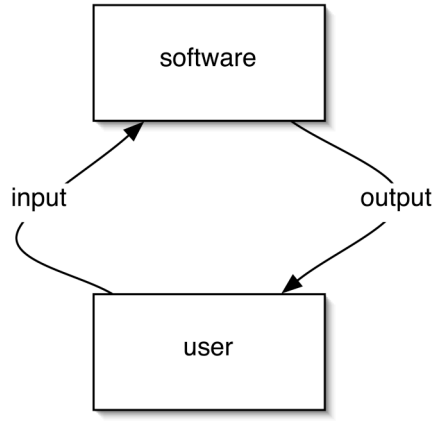


**Figure 6.** Interactive feedback loop linking users and software.

What differentiates this software from a general-purpose program? This special subset of software is often described as software art, and is distinguished from other software as follows:

> ...at the basis of each piece of software there are definite algorithms, but if conventional programs are instruments serving purely pragmatic purposes, the result of the work of artistic programs often finds itself outside of the pragmatic and the rational. (Goriunova and Shulgin 2002)

It is possible to create art with general-purpose software; a word processor can be used to write this document, but it can also be used to write a novel. But with software art, the practical application is absent: the environment is too narrowly constrained and too idiosyncratic in its design to be useful for anything other than creative ends.

Software art also occupies a middle ground between a passive and active experience, and its users lie somewhere between witnesses and participants. Software art users do not merely listen to a piece of music or look at a painted canvas. Nor do they follow specific instructions in the same way that performers follow a score; their activities remain rather undirected and open-ended. But at the same time, they can only exercise their creativity within narrow confines, and software art applications tend to emphasize the creative process over the final product. In this sense, software art differs not only from productivity software, but also from creative "pro-sumer" applications, such as Apple's Garageband, which provide general-purpose interfaces for users to create a wide range of artistic products.

While much software art (and software) uses conventional input and output interfaces — keyboards, mice, monitors, and speakers — it is by no means limited to these devices. In fact, in many situations, the software and the computer on which it runs remain invisible to the user. In extreme cases (e.g. the Turing test), the user may not even realize she is interacting with software at all.

One of my first works of software art was *Telephone Etude #1: Shakespeare Cuisinart* (2001). Users call a toll-free phone number, say their favorite quotation from a Shakespeare play or poem, and then listen as their voice is sliced and diced to create a short piece of *musique concréte*. The work is extremely accessible; users need not have any musical or technical training, and even the most timid user can whisper some famous lines of *Hamlet* into the phone. But at the same time, the creative space which users are able to explore is rather limited: every instance of *Shakespeare Cuisinart* is immediately

recognizable as such. And while users have some control over the content, they have little influence over its form.[1]

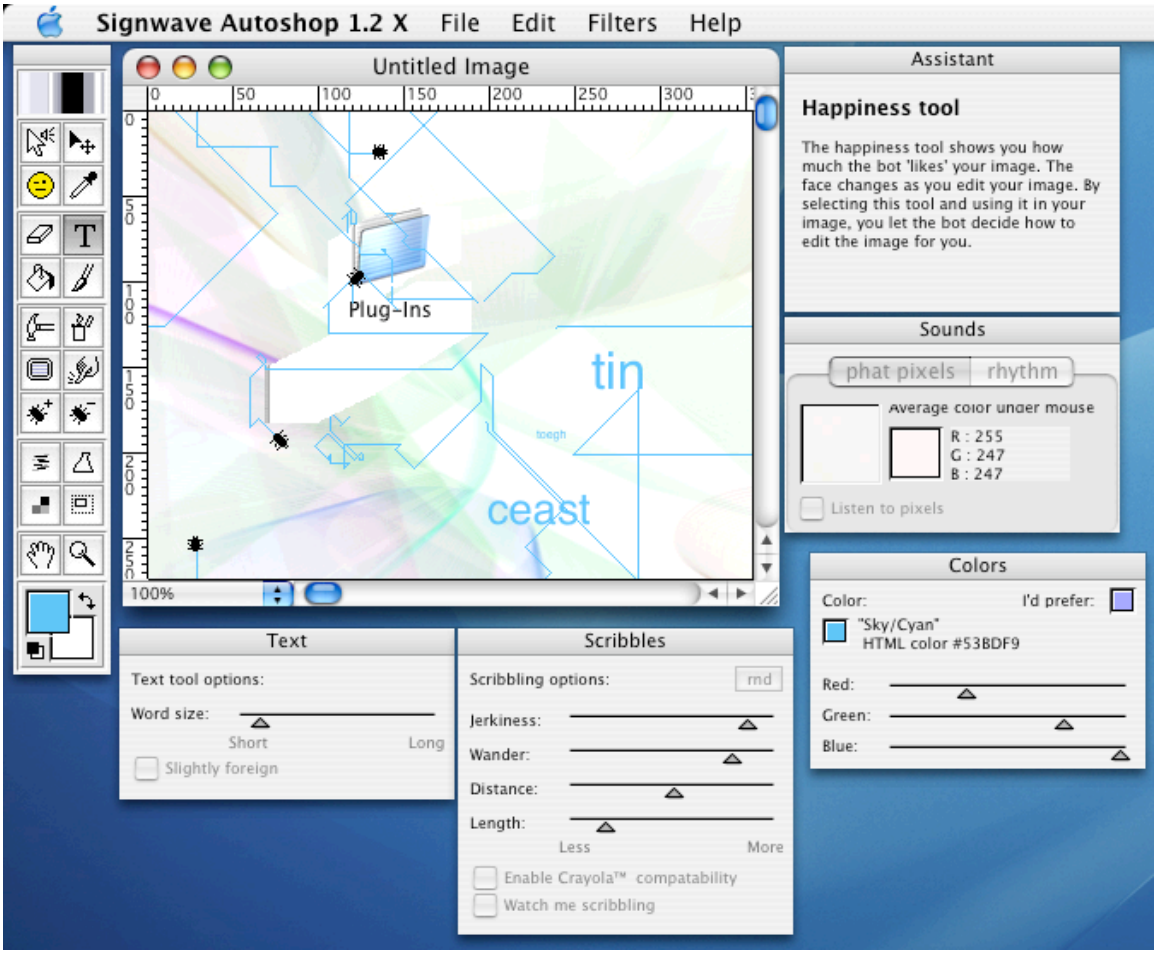


**Figure 7.** Adrian Ward's *AutoShop* in action.

Adrian Ward's AutoShop (1999) is software art which deliberately resembles a conventional software application but systematically thwarts the user's attempts to utilize it in a conventional manner:

> …Autoshop is an explorative parody of professional bitmap graphic manipulation software. By asserting it's own creative agendas upon the user, it raises awareness of authorship and the position of the creator when digital systems are in use. It's also good for destroying your images. (Ward 1999)

[1] A hierarchy of random decisions controls the slicing and dicing operations.

Traditional toolbar items do not operate as expected: the eraser tool is as likely to add pixels to the image as to remove them; circles and rectangles are distorted into odd squiggles; and the text tool inserts seemingly random words in various sizes and styles. And user activities are frequently interrupted with actions initiated by the software, discouraging the methodical creation of a finished image.

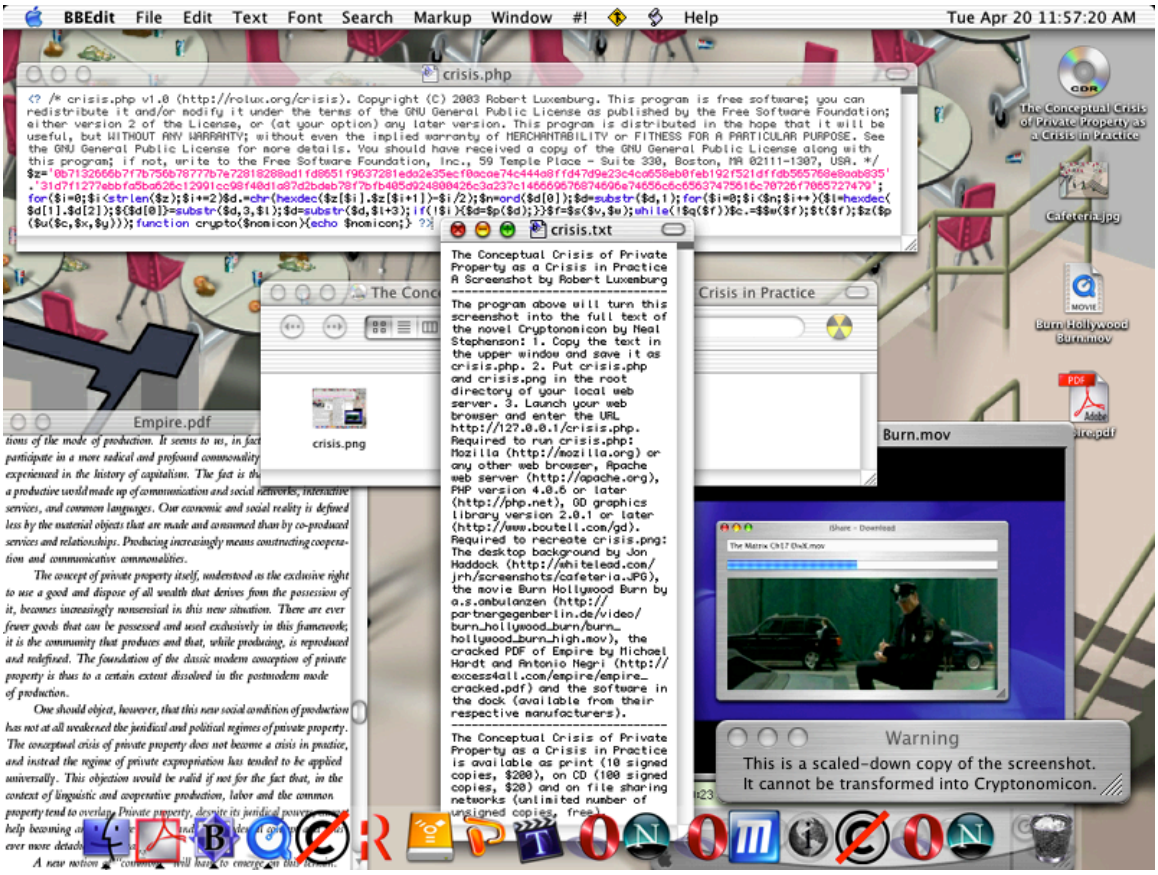


**Figure 8.** *The Conceptual Crisis of Private Property as a Crisis in Practice* by Robert Luxemburg (Luxemburg 2003).

And while still clearly a work of software art, Robert Luxemburg's *The Conceptual Crisis of Private Property as a Crisis in Practice* (2003) engages a different user base in a different manner. The software is a short script which transforms a

screenshot file of the artist's desktop (Figure 8) into the complete text of Neal Stephenson's novel *Cryptonomicon* (1999). The software itself is not actually interactive; it processes an input file and creates an output file, and it only functions properly with one specific input file. But it still engages the creativity of a technically-qualified user within the narrow confines defined by the work:

> F.A.Q.
> ------
>
> Q: How does it work?
> A: The source code is open, so find out yourself.
>
> Q: What's up with the last 50 bytes of $z?
> A: If you are that far, you are very close.
> (Luxemburg 2002)

And like much other software art, the emphasis is on process over product, on engagement over utility. It is difficult to imagine a user running the software (and spending $20 to purchase the necessary input file) just in order to read *Cryptonomicon* (whose retail price is $8).

## 1.6 Large Scale Listener Participation

Like the examples in the previous section, *Glimmer* is a work of software art. Its input interfaces are light sticks and video images, and its output interfaces are multi-colored lights (and by extension the orchestra) and a video projection. The primary users are the audience members.

But unlike the previous examples of software art, *Glimmer* is performed during an orchestral concert in a conventional hall. It is a collective, not an individual, experience; there are 600 audience members simultaneously participating. And the roles of composer,

performers, and listeners are still relevant, because the work must still operate within many of the conventions of an orchestral concert.

So *Glimmer* also follows in the tradition of works which facilitate real-time participation by a large audience. Many such works are not software art; they do not use technology, and listeners function more as performers than users, closely following instructions dictated by a pre-composed score. For example, in Jean Hasse's *Moths* (1986), the audience whistles as directed by a conductor and a graphical score to perform the piece; there are no musicians on stage. And during *La symphonie du millénaire* (Chénard 2000), an outdoor event in Montreal written collaboratively for multiple ensembles, two thousand audience members rang handheld bells at designated times.

Many Fluxus scores also involve the audience as performers, but in a less directed manner which more closely resembles software art. For example, the score for Tomas Schmit's *Sanitas No. 35* reads: "Blank sheets are handed to the audience without any explanations. 5 minutes waiting" (Schmit 1962).

In other works, the audience contributes input which affects the musical performance rather than creating sounds which are part of the performance. As with much software art, these works create simple input interfaces to control musical decisions. For example, in Thomas C. Duffy's *The Critic's Choice* (1995), a film-music takeoff for concert band, the audience votes for one of three alternate endings to the piece. And in a performance of Terry Riley's *In C* (1964) staged by the Eos Orchestra, audience groups seated at tables tap electronic dome-shaped centerpieces to advance a MIDI instrument to the next musical motive (Bianciardi, Igoe, and Singer 2003).

Several recent projects use technology to create new interfaces for live concert audiences but never ask listeners to take on the active role of performers or users. Golan Levin's *DialTones: A Telesymphony* (Sheridan 2001) triggers audience mobile phones to play pre-composed ringtones; the phones are used as an output interface, but there is no audience input interface. And the Concert Companion software (Mirapaul 2003), which has been tested with audiences at the New York Philharmonic and the Aspen Music Festival, delivers real-time multimedia program notes to audience members via wireless PDAs.

*Glimmer* also follows in the tradition of several mass-audience gaming projects staged in conventional performance venues. Cinematrix has produced events in which audience members hold up the red or green side of a paddle to collectively navigate objects on a video screen and play simple competitive games (Carpenter and Carpenter 1999). The same system has also been used to create interactive educational planetarium shows (Fisher et al 1997). Maynes-Aminzade, Pausch, and Seitz (2002) have created similar environments for classic arcade games by tracking video of audience members as they shift left and right in their seats. And the NYU Movement Group's *SquidBall* (2004) tracks the motion of giant weather balloons; the audience throws the balls through the auditorium to clear virtual targets on a giant screen and trigger corresponding video and sound effects.

These projects all faced daunting technical, logistical, and conceptual challenges in organizing mass-audience participation within conventional performance venues. The strategies and solutions they devised provided a valuable foundation for the development of *Glimmer*.

# 2. Context and Design

## 2.1 Design Constraints

From the beginning, the environment for which *Glimmer* was conceived — an orchestral performance in a concert hall — influenced the direction of the project's development. While I was interested in rethinking the conventions associated with that environment, I also needed to address the practical limitations the situation imposed.

The American Composers Orchestra gave me three broad directives for the work: it must use technology in some way, it must be fun, and it must last no longer than ten minutes. Beyond these high-level constraints, numerous practical details imposed further limitations: there was just a single, one-hour offsite rehearsal and a thirty-minute dress rehearsal; technical setup in the hall lasted just six hours, along with a ninety-minute testing session a few months earlier; and stage setup time during the concert was just fifteen minutes. To further complicate matters, the project was produced with a shoestring equipment budget (about $2500) and with no funds to hire additional personnel for technical assistance. And some of the equipment on rental was not received until just a few days before the performance.

But perhaps most frightening of all, there was no way to rehearse the piece, and there was only a single performance by the orchestra. Union regulations at Zankel Hall prohibited the orchestra from bringing more than a handful of audience members to the dress rehearsal, which was not nearly enough to understand how the piece would work with a full audience. The only chance to truly test the technical and conceptual design

was at the premiere, with a full 600-member audience and 25-player orchestra in the hall.

## 2.2 Design Goals

Similar design goals inform most of my works of software art. The following sections describe these goals and the ways in which they combined with design constraints to influence the course of development for *Glimmer*.

### 2.2.1 Accessibility

My software art is designed for users who do not have any specialized musical or technical training. It must be easy for users to understand the functionality of the input and output interfaces, and the environment must put users at ease so that they are not shy about contributing.

With *Glimmer*, the short length of the piece made accessibility even more critical than usual. The interface — a single light stick for each audience member — was chosen partially because of its ease of use. The audience's role in the performance is straightforward, and important information is explained verbally and displayed on a video screen. And the audience's activities do not directly produce sound, but instead affect the music played by the orchestra: a Greek chorus more than a *Messiah* sing-in chorus. This puts shy audience members at ease, who might otherwise be hesitant to participate in front of 600 peers.

Because of the limited rehearsal time, accessibility was also a concern for the orchestral musicians. A simple color system communicated instructions to them, with only four color families and corresponding pitches for each player. Musicians were asked

to exercise only a limited degree of interpretive freedom, mostly with regards to exact dynamics, balance, and accent strength. And to facilitate individual practice, each musician received a video DVD which showed a simulation of her light.

### 2.2.2 Transparency

In my software art, users must quickly understand the relationship between their activities and the software's response. Otherwise, they can become frustrated, convinced that their actions have no impact on the music they hear.

I did several things in *Glimmer* to make this relationship transparent. The music is extremely simple — slowly changing sustained notes and clusters inspired by works such as John Cage's *Four*[2] for chorus (1990) — and so events such as accents and note changes are easily identifiable. Furthermore, the large-scale structure of the piece, which moves from a thin texture to a dense mush and back to a thin texture, enables listeners to identify these events most clearly at the moments when transparency is most important.

Audience members must also be able to aurally identify their own group of musicians from within the ensemble. While the existence of seven different musical groups might seem to push these limits, no audience member must isolate all seven; she need only recognize her own group within the ensemble. To that end, the seven groups are distinct in both pitch and timbre. Each group plays a different set of adjacent pitches[2] which are never used by other groups in the orchestra. And each group includes instruments which are similar in timbre; for instance, Group A consists of four violins and Group G includes three pitched percussion instruments. In situations where two

[2] These are pitch sets, not pitch-class sets, and they are adjacent within the 28-note set used in the piece.

groups are timbrally similar — such as with Groups A and B, which both include violins — their pitch sets are registrally distant from each other.[3]

Finally, the activities of both the audience and the orchestra are reinforced visually. Audience members directly see the light sticks of their peers, but they also see an abstract video animation of their group's activities and its competitive rank. And audience members directly hear the music, but they also see the colors of the lights on the musicians' stands.

### 2.2.3 Sustained Interest

My software art must be able to sustain the interest of users over time. With *Glimmer*, the duration of the work is short (ten minutes), and so this was generally less of an issue than accessibility and transparency.

I was concerned, though, that the piece could stagnate if audience members left their light sticks on for extended time periods, or if the simple harmonic, rhythmic, and textural language became too static. To address these issues, I developed a competitive aspect to the audience's participation; groups which changed their on-off percentages more dramatically over time were rewarded. If an audience group simply left their lights on for the entire piece, their musicians would recede into the background or stop playing altogether, while more interesting groups would become more musically prominent.

I also designed a fixed large-scale structure for the piece. At regular time intervals throughout the performance, subtle changes to the software's algorithmic

[3] For a complete listing of instrumental groupings and pitch sets, refer to the Pitch Map section of the score.

parameters gradually transform the musical texture.[4] Each individual change is barely noticeable, and so these events never seem to stem directly from audience activities. But cumulatively, they create dramatic musical transformations which shape the piece and help keep the music interesting.

### 2.2.4 Reliability

At the risk of stating the obvious, it is incredibly important that software art, along with the hardware on which it runs and the interfaces which it uses, actually works. Before *Glimmer*, most of my experiences with reliability issues centered around my Internet-based software art. The challenge with these works is to monitor and maintain compatibility with a variety of computer platforms over a long period of time. If a user in Moscow downloads the software but is unable to launch it, she is much more likely to give up than to send a detailed report. How would I even know that anything had gone wrong, let alone be able to diagnose and fix the problem?

With *Auracle* (2004), a networked sound instrument I developed in collaboration with Max Neuhaus, Phil Burk, Sekhar Ramakrishnan, Kristjan Varnik, and David Birchfield, we solved this problem. The client-side Java applet connects to a server-side MySQL database, where it logs detailed information about the user's system configuration and any errors which she encountered. Each evening, an automated e-mail report is sent to the developer mailing list, summarizing the problems which users encountered during the previous day. The reports often trace errors to a specific line of source code, making it easy to quickly identify and resolve problems (Varnik et al 2004).

---

[4] The complete structure is documented in the Timeline section of the score.

With live-performance works, there is a different kind of challenge. During the performance, it is usually quite easy to identify problems with the system, but it is too late to fix them. With *Glimmer*, where access to the hall and to some of the equipment was so limited, and where there was no true rehearsal for the piece, the software and hardware could not be tested in a performance situation until the performance itself. Under these conditions, how could I build safeguards against problems?

My solution was to maximize redundancy and flexibility, build numerous simulation, monitoring, and diagnostic tools, and leverage industry-standard hardware and communications protocols whenever possible. Following from Madden et al (2001), two redundant backup computers ran the same software as the main machines, receiving state information over UDP. If a primary machine crashed, the backup could be switched into use with a single button press. The software itself included graphical and text-file interfaces through which hundreds of algorithmic parameters could be quickly modified during setup, rehearsal, and even the performance. And numerous simulation, logging, and visualization mechanisms were built into the software for development, testing, and rehearsal purposes. Musician light data, for example, could be imported into Microsoft Excel and Quicktime for frame-by-frame analysis of problems.

But the Color Kinetics iAccent LED light tubes, which sat on each musician's stand, were the most dramatic example of the reliability measures taken for *Glimmer*. These lighting units, enclosed in a rugged, water-resistant plastic housing, were daisy-chained together using power/data cables with over-molded, airtight-locking connectors, and attached to power supplies which safely sat on the far corners of the stage. The lights were serially configured so that they could be connected in any order and to any power

supply, so long as the correct light was placed on the correct music stand.[5] And the computer software communicated with the lighting power supplies via a simple UDP protocol over a standard Ethernet LAN.

Certainly a water-resistant lighting system certified for 100,000 hours of outdoor operation was overkill for this project. There were other important features which led to this choice of hardware: its 24-bit color space, low latency, high bandwidth, compact size, and silent operation. But when given the choice between hacking together a proprietary solution or leveraging an industrial-grade product, the answer was obvious. Unsurprisingly, the Color Kinetics lights were the easiest component to integrate into the system, and they were the least cause of concern during setup, rehearsal, and performance. And it actually cost less to rent these lights than it would have cost to build a proprietary solution.

## 2.3 System Design

All of these ideas about the roles of composer, performer, and listener and about software art, and all of the design constraints and goals, ultimately led to the high-level system design outlined in Figure 9.

Figure 9a is essentially a combination of Figure 1, which showed the relationship between composer, performer, and listener in three discrete stages, and Figure 6, which showed the feedback loop linking users and software. Here, the feedback loop takes place during the performance and includes composer, performers, and listeners: audience activities are the input to the composer's software, which outputs instructions to the

[5] This reduced the chance for errors during the quick stage setup.

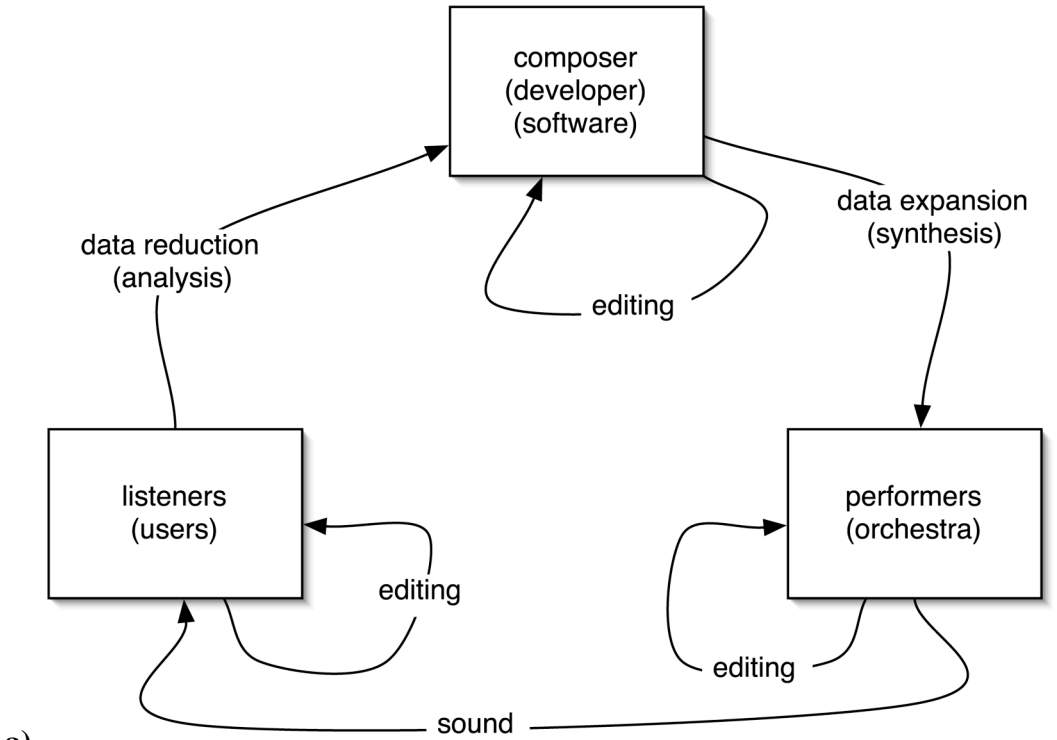


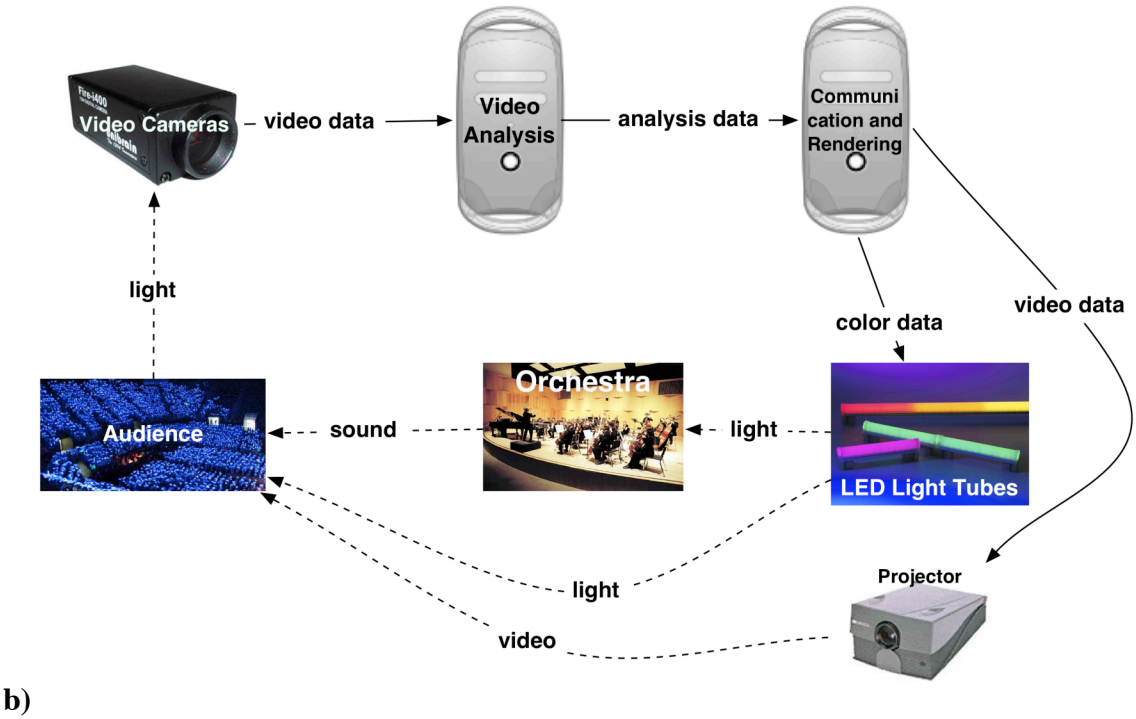


**Figure 9.** High-level overview of system design, on **a)** a conceptual level, and **b)** a concrete level. In Figure 9b, solid lines indicate connections by physical cables, while dotted lines indicate transmission by other means.

orchestral musicians. The musicians respond to those instructions, generating sounds which the audience hears and to which it responds, thus beginning another iteration through the loop.

Because this feedback loop occurs during the performance, portions of the editing loops of the composer, performers, and listeners also take place during the performance. Listeners edit their participation strategies based on the musical results of their previous activities and the responses of their peers. The composer monitors the progress of the software, tweaking parameters in real time to improve its response to the audience. And the performers listen to the musicians around them, modifying dynamics, balance, and accent strength to blend with the unique music being played by the ensemble during that particular performance.

The conventional roles of composer, performer, and listener are clearly identifiable, but the users and developers of software art are also clearly relevant. The audience members are the users: the light sticks are their input interface and the orchestral musicians are the main output interface; the orchestra takes a role similar to that of computer-based audio processing and synthesis in other works of software art.

The flow of data through the system illustrates another important aspect of the design. As audience activities are input into the software, they are analyzed and compressed down to their bare essence: the on-off percentages for each audience group. As data flows out of the software to the musicians, it is expanded in two stages. The software first maps the data onto instructions for each player in the orchestra, and then the orchestra transforms those instructions into actual sound.

# 3. Assessment and Discussion

The previous two chapters outlined an aesthetic vision for *Glimmer* and its high-level design. This chapter evaluates the work in the context of its premiere performance.

The American Composers Orchestra asked for a piece which used technology and was fun; in this sense, the performance was a success! Technologically, *Glimmer* met all design goals for reliability. The only significant technical problems were the result of human errors; for example, the video threshold filters were calibrated using light sticks with weak batteries and had to be readjusted in the opening minute of the performance. And some contingencies had simply not been anticipated. At the rehearsal, one of the violinists reported that he was color blind, but fortunately he was able to differentiate the colors with only minor difficulties.

And the performance was fun. The audience even gasped and laughed at moments in the performance, and audience members found creative ways to contribute, including a version of the ballpark wave which raised and lowered light sticks from the camera's view.

But the other high-level design goals — accessibility, transparency, and sustained interest — did not fully meet my expectations; these goals are evaluated in depth below. And beyond these criteria, two more fundamental questions remain. Did audience members feel that their contributions were important and that the performance would have been different without them? And did I feel that because of their activities, the music proceeded in surprising directions which I had not imagined it would?

## 3.1 The Audience

In informal interviews after the performance, some audience members claimed they had a tremendous influence over their group's activities and the performance, while others were frustrated that none of their actions seemed to make any difference.

The input and output interfaces were sufficiently accessible; audience members had no trouble understanding how to participate in the piece, learning in which group they were located, and identifying the musicians in the orchestra whom their group controlled.

The primary problem was with the transparency of the system. While almost all audience members could clearly hear their musicians within the orchestra, only some of them were able to understand the connection between their individual actions and the music they heard. And groups had great difficulty working together to effect more dramatic musical change. The following sections explore these issues in detail.

### 3.1.1 Variation In On-Off Percentages

The system's response to audience activities was subtle because of the small variation in group on-off percentages during the performance. Groups ranged from approximately 30% to 70% activation, but as Figure 10 indicates, each individual group only varied by a smaller amount over the course of the performance. As a result, the direct musical effect of each group on its musicians was often difficult to hear. And while the musical effects of the competitive component of the piece were clearer, the

competitive rankings sometimes seemed mysterious and arbitrary, since they were based on relatively small changes in group on-off percentages. [6]

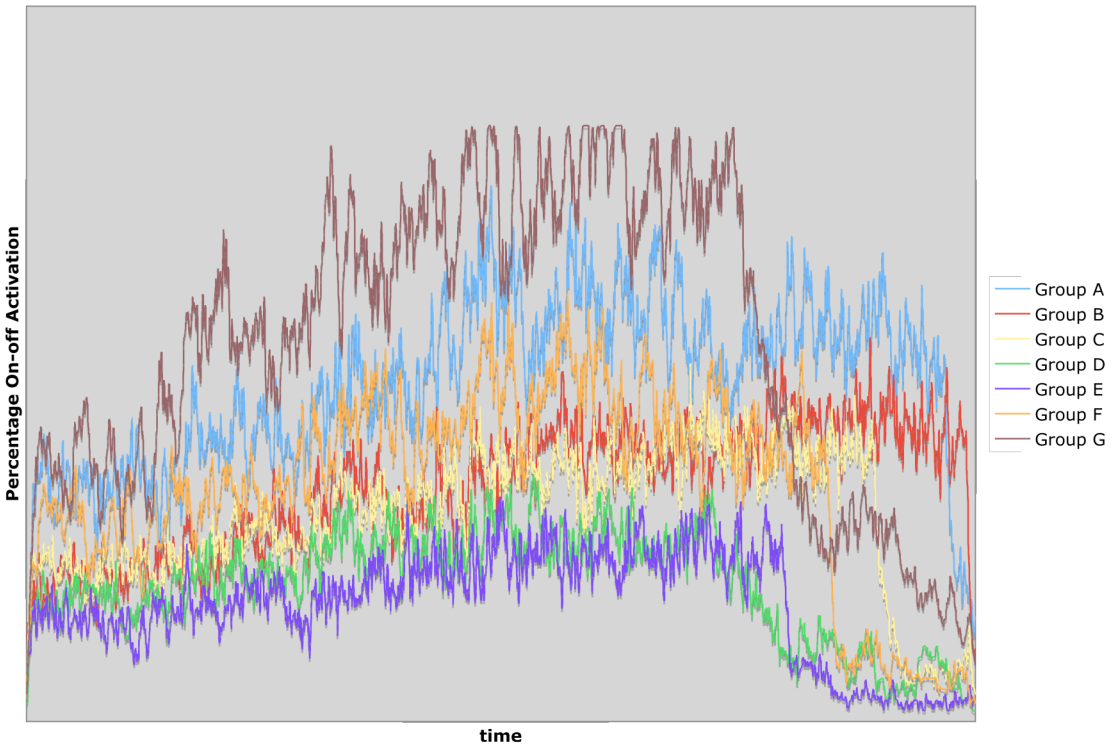


**Figure 10.** Audience group on-off percentage data from the premiere performance. When groups suddenly drop towards 0% towards the end of the piece, they have been eliminated from the piece and have turned off their light sticks.

While my larger goal would be to facilitate greater group cohesion and greater variation in audience data over time, some small modifications could improve the transparency of the system even when audience data variation is minimal. Adaptive normalization algorithms could be implemented to maximize the direct control exerted by audience groups over their musicians. And competitive ranking decisions could be made by a human judge rather than a computer algorithm.[7] Groups could then be rewarded for

[6] The order of magnitude of each group's variation was not much larger than the order of magnitude of the expected error in the video analysis algorithms.
[7] In fact, this functionality already exists in the software.

activities which are not anticipated by the algorithm, such as the burst of creativity exhibited by Group G when it started the wave.[8]

### 3.1.2 Self-Organizing Group Behavior

The transparency of *Glimmer*'s interactive system was problematic because of the relatively small variation in group on-off percentages. In turn, those percentages varied so little because groups were unable to work together cohesively. When many group members switched their lights on and off quickly — but out of sync with their neighbors — their activities simply cancelled each other out in the software's analysis.

My hope had been that even in the absence of pre-appointed or self-appointed group leaders, interesting group behavior would emerge over time in a manner similar to cellular automata (Wolfram 2002). The simple rules which governed the competitive aspect of the piece were designed to encourage such behavior, but while the competition added an exciting dimension to the experience, it failed to elicit the group behavior I had hoped it would.

In informal discussions with audience members, I learned of several reasons why groups had failed to cohere. Some people complained that the piece was too short for them to develop a group sensibility; they felt they would have done better had the piece been longer, or had it been performed a second time. Others had trouble seeing all the

[8] Even though Group G was working together cohesively to perform the wave, the number of light sticks showing at any given moment did not change very much, so the group was not rewarded by the ranking algorithm for this activity.

people in their group, so it was difficult to respond to what peers were doing.[9] A more detailed, less abstract video animation could give audience groups a better sense of current on-off percentages and help individuals make more informed decisions based on group state.

But despite these problems, some interesting group behavior *did* emerge; it was just not the kind that I had expected, nor the kind that I had tried to elicit, nor the kind that the software was designed to recognize. Audience members enjoyed waving their light sticks around much more than switching them on and off, even though they knew that such activity had little effect on the music. Not only was it more fun for them to do, and not only was it more pleasing for them to watch, but it also gave them the *feeling* of more communication with and control over their peers. In an on-off system, listeners could essentially send only a single message to their group: *I am turning my light on*. People rarely noticed when a single light was switched off within a sea of light sticks, and once it was off, that audience member effectively lost her voice until the light went back on.

In a waving system, though, audience members would be able to leave their lights on for long periods of time and create gestures in space with them, sending continuous information to their peers. Suddenly, the velocity of their light stick would become important. And if velocity, rather than on-off state, were analyzed, then it would also be more difficult for individuals to remain in an activated state; it would require continuous energy to keep the stick moving. There would also be a continuum of data on both the

---

[9] While this limited visibility is a characteristic of most cellular automata, it would be helpful in *Glimmer* for audience members to respond to changes in the total on-off percentage of their group, since this data is so critical to the performance.

individual and group levels; individual actions could express a range of values through velocity. It is little wonder that the most cohesive group activity in the performance was the wave.

A simple change to the video analysis algorithm to calculate pixel differentials between successive frames, instead of blob counting, could have accounted for stick movement in the video analysis, captured a greater amount of variation, reflected this activity in the instructions sent to the musicians, and ultimately improved the transparency of the system and the cohesiveness of groups.

### 3.1.3 Transparency and Sustained Interest

Did the audience remain interested and engaged for the duration of the piece? A cursory analysis of the performance video offers a clue. During the middle of the piece, audience members become less active in using their light sticks for a few minutes, suggesting a fading interest in participation. While audience members do still use their light sticks, they do not switch them on and off as often or wave them around as aggressively. In the closing minutes of the piece, as groups are taken out of the music one by one, the audience's interest revives, and remaining groups participate more actively and aggressively than ever.

So why did the audience's energy fade for a few minutes? I believe it was connected to the transparency issues discussed in the previous sections. By the middle of the performance, the audience's difficulty in understanding the effects of their individual actions may have reached its peak. The music had also reached its densest textures, making the identification of a group's musicians more challenging. So with less

understanding of the role they played in the music, the audience participated less. Later, when the musical texture became thinner, and when the competitive aspect of the interaction surged in importance, the system became more transparent again, and the audience participated more energetically.

But the audience's relative lack of activity in these middle minutes could also indicate an *increase* in a different kind of engagement. As the music reached its densest and most complex textures, audience members may have listened more carefully, dedicating less energy to participating in the performance.

If this explanation is correct, then there is an apparent contradiction. In section 1.3.1, I explained that I actively involve listeners in musical performance to encourage them to listen more closely. Yet in *Glimmer*, I transformed an environment in which music is usually the sole focus of attention into an environment which is full of distractions. Can audience members both participate actively and listen carefully at the same time?

I believe that they can, and that they did in the opening and closing minutes of *Glimmer*. In the middle, their multi-tasking became more difficult because the individual tasks became more difficult: the music's growing complexity required more concentrated listening. And since the system's decreasing transparency made participation less directly rewarding, audience members chose to focus on listening at the expense of interacting.

But even when audience members focused on listening, they listened in a qualitatively different way than they would have if they had they not been participants. With a personal stake in the outcome of the performance, audience members directed their attention towards their own musicians in relationship to the rest of the ensemble.

The music sounds different to an audience participant than to a passive observer, and it sounds different to participants in different audience groups, because each person searches for different things in what she hears.

And regardless of what they did in the middle of the piece, every audience member (except for those in the last group standing) made the transition from active participant to passive observer in the closing minutes of the performance. Once their group is permanently eliminated from the music, they lose the personal stake that persuades them to focus on particular parts of the music. For me, this transition is one of the most powerful aspects of *Glimmer*, because it forces the question: how should I continue to listen to this music now that I am no longer involved in making it?

### 3.1.4 The Role of Competition

Of all the interactive aspects of *Glimmer*, the competitive component engaged the audience the most. This surprised me, especially since the ranking of audience groups was not particularly transparent in the performance. But especially in the work's final minutes, as groups were taken out of the piece one by one in a strange cross between Haydn's "Farewell" Symphony (1772) and a reality television series, the excitement of the audience was obvious and audible.

I had originally conceived the competitive component of *Glimmer* as a means to subtly encourage group cohesion, but once that competition was represented both aurally and visually, the audience's obsession with it was probably inevitable. This pleased me because it helped engage the audience and make the piece fun, but it also concerned me because it detracted from other aspects of the feedback loop (see Figure 9). Some

audience members were so intent on "winning" the piece that their decisions were based more on competitive urges than musical responses.

Unless aspects of the transparency of the system were improved, the competition remains the most immediately accessible and engaging aspect of the audience's participation. But if those aspects were improved, then the competitive aspect might be diminished by removing some visual information about rankings, such as the circle on the video projection highlighting the first-place group. It could also be downplayed in the verbal introduction to the performance.

### 3.1.5 Undirectedness and Context

With *Glimmer*, it was extremely important to me that the audience's activities remain relatively undirected, that they creatively explore an interactive space rather than merely follow instructions. To ask them to follow a linear score, to conduct their participation, or to otherwise guide their actions would undermine the original motivations and goals of the project.

Such undirected creative activity comes naturally in the context of much software art. When users are sitting in front of a computer, the distinction between developer and user is clear, and undirected interaction is the dominant paradigm. But when software art comes to the orchestral concert hall, it enters a dramatically different context. The computers are hidden backstage, the interface is transformed, and the experience is collective. And the program lists a musical composition, credited to a single composer,

thus concealing the importance of the developer / user relationship and devaluing any contribution audience members might make. [10]

Furthermore, *Glimmer*'s premiere was situated in the middle of an evening of music; it was the only work on the program requiring the audience to assume an active role. As *Glimmer* began, audience members had to instantly leave behind the conventions established not only by previous orchestral concerts they had attended, but also by the previous pieces on the evening's program.

Under these circumstances, the undirected nature of their interaction can compound an already challenging task. Not only must audience members suddenly shift from a passive to an active role, but they must make this shift with no specific instructions about what to do and when to do it. My introductory remarks simply outlined the basic structure of the piece and provided deliberately vague explanations of possible strategies.

In future performances, short directed practice exercises could better prepare the audience for their role in the piece; such practice has been an effective tool in similar projects (Fisher et al 1997). During my introductory remarks, I had each audience group, one at a time, switch on their light sticks. I should have taken this exercise further, giving each group time for "solo" practice before the performance began, providing them visual and aural feedback during the practice period, and directing their practice activities so that they could learn to produce the full range of visual input with their light sticks. Then they would have been better equipped participate during the undirected interaction of the performance itself.

---

[10] To partially remedy this, I listed the instrumentation of the piece as "chamber orchestra and audience" in the program and offered a brief explanation in the program note.

## 3.2 The Orchestra

### 3.2.1 Accessibility

Of all of the high-level design goals outlined in Chapter 2, accessibility was the most important with respect to the orchestral musicians, and unfortunately, it did present some problems in rehearsal and performance. As expected, few musicians practiced their parts individually with the DVD they had been given. In rehearsal, the players still quickly learned how to respond to the colors displayed by their lights, but many of them had difficulty responding accurately to quick accents and note changes. And because their attention was so focused on lighting cues, they were unable to concentrate on issues such as balance and intonation, and the quality of the performance suffered as a result.

A few small changes could have addressed many of these problems: increased rehearsal time; a rehearsal venue in which overhead lights could be switched off to make the colors of stand lights clearer; more active involvement by the conductor; and a small increase in the length of preparatory cues for accents and note changes.

### 3.2.2 Interpretation

In *Glimmer*, there is a fundamental inequality between the audience and the orchestra. The audience works within the framework defined by the piece but follows no score, interacting with the software to shape the music. The orchestra, on the other hand, closely follows the instructions indicated by their lights, with only limited interpretive

freedom. The musicians function less as users of *Glimmer*'s software than as a part of its output interface.

Given that the audience is able to respond to the music they hear, should the orchestra have a greater opportunity to respond to the audience's activities? In line with the ideas presented in section 1.2, should the musicians be given increased freedom to interpret the music they play? Should they feel that they are making a unique contribution to the performance, and that the music would be dramatically different if played by someone else?

While this idea is conceptually appealing, it does not make practical sense in the context of *Glimmer*. The orchestral musicians struggled just to familiarize themselves with their cues during the limited rehearsal time available. How could they also be expected to make far-reaching, unfamiliar interpretive decisions in this environment?

Giving musicians greater interpretive freedom in *Glimmer* could also undermine the transparency of the system. The music is carefully constructed so that perceptually salient events always originate from audience activity. Were musicians to alter these events or add their own, it would be even more difficult for audience members to establish the relationship between the things they did and the music they heard.

Finally, an aspect of this inequality is actually quite appealing. Orchestral audiences typically have no control over the performance they attend. Finally, some poetic justice: the audience is empowered at the expense of the orchestra.

## 3.3 The Music

Throughout this paper, I have emphasized the importance of the process of making music in *Glimmer* over the resulting musical product. When I have discussed aspects of the music, it has always been in connection with the design of the work's interactive system. Nevertheless, there are some questions which remain about the music and its role in the work. How did it relate to the nature of the audience's interaction? How did it relate to other works of mine? How (and why) did the music depart from my expectations?

### 3.3.1 The Character of the Music and the Audience's Activities

Audience members at the performance felt excitement, tension, amusement, and surprise, but these feelings were aroused by the interactive experience, not by the music itself. The music focused on slowly-changing textures and harmonies, subtle timbral transformations, and gradual dynamic sweeps. This contrast between the character of the audience and of the music was most vivid in the closing minutes of the performance; the audience's excitement audibly rose as only a few groups remained, but the music grew increasingly sparse as only a few musicians continued to play.

While this discrepancy was disconcerting to a few of the audience members with whom I spoke, it does not concern me. As previously discussed, I wanted to keep the music as simple as possible so that it would be easy for audience members to identify their musicians within the ensemble and to connect their activities to the sounds they heard.

Equally important, I wanted to create music which would counter-balance the novel mode of interaction and encourage audiences to take their role seriously. The slowly-moving music helped subdue the audience, urging them to listen more carefully to what they heard. While there was a strong element of novelty and humor to the experience, it was impossible to treat it as merely that.

### 3.3.2 Details and Pure Music

Stylistically, the music of *Glimmer* is similar to my recent non-interactive instrumental works such as *Prior Art*. Yet listening to both as pure music, I find *Prior Art* to be by far the better piece. What makes that music more satisfying to me?

*Prior Art* is slow, simple, and deliberate music. It is not so much the broader form which makes it effective, but the small details: the fragile construction of a harmony, the irregular breathing of the meter, the slow decay of a percussive sweep. Some of these details are specified in the score, and some of them are decided by the performers, but they are all controlled by a small group of people who have invested a great deal of time, energy, and musicianship in the performance.

In *Glimmer*, these kinds of details are determined not so much by the composer or the performers as by the six hundred audience members who know nothing about the piece until moments before its performance. It was important to me to give this level of control to the audience, because I feared that their role would otherwise become superficial and banal. Yet under these circumstances, it would be absurd to expect those details to be controlled with the same degree of subtlety as in pieces such as *Prior Art*.

Opening up the creative process means giving up control, and lowering the barriers of training and commitment to enter that process usually leads to a more exciting process but less exciting results. No interactive system can do everything (see Chapter 4), so instead, I am content to write many different kinds of works which fall at different points along a continuum, each piece finding its own way to distribute control among composer, performers, and listeners. And still, incredible, unexpected musical details sometimes do emerge from works at the far interactive extreme of the continuum.

### 3.3.3 Musical Surprises

At the beginning of this chapter, I stated that *Glimmer*'s success hinged on the belief of audience members that their contributions mattered. For this to be true, the performance cannot proceed exactly as it did in my imagination (or in my software simulations). I should be surprised — sometimes for better and sometimes for worse — by the directions it takes.

In the performance, one passage particularly surprised me, because for a moment, the subtle musical details combined as effectively as in any music I have ever composed. The passage begins approximately eight minutes into the piece, when four, and then just three, groups remain. At the same time that Group F (upper winds and brass) exits the piece (Gb5 in Figure 11), Group A (first violins) leaps up a perfect fourth to an Eb6 (the highest note it plays) and Group C (violas) leaps down an augmented fourth to Gb3 (the lowest note it plays). Group B (second violins) remains on a Db5. Group A soon moves down a step from Eb6 to Db6, melodically resolving its leap, doubling the Db5 at the octave, and momentarily leaving us with the sound of a bare perfect fifth. Group A's

crossfade from Eb6 down to Db6 draws out the drama of this melodic and harmonic resolution; it almost sounds like a slow glissando. I could not myself have written this passage better, nor could the orchestra have played it better.

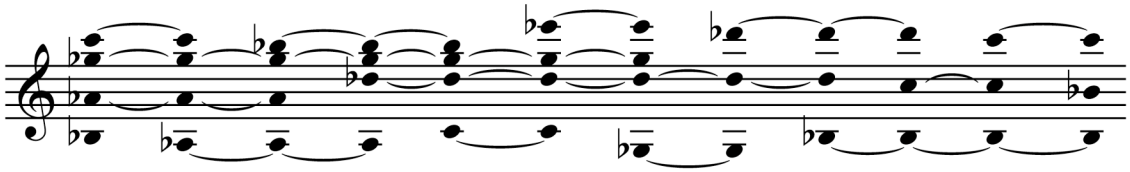

**Figure 11.** Reduction of the premiere performance, beginning approximately eight minutes into the performance.

On one level, it is easy to understand how this moment transpired. There is no element of chance in the software; given the same sequence of inputs, it will always produce the same outputs. I can understand exactly how my own pre-composed cues combined with the competitive rankings of the audience to create this succession of events, even though I could not have predicted in advance what would take place then, or even which instruments would be playing.

But in truth, I have no idea why this moment transpired. I wish I could claim that audience members intended it to happen, but even had they been working together more cohesively, I seriously doubt they could have planned such a subtle sequence of events so precisely. Yet I am also reluctant to label it mere chance, because in my hundreds of tests and simulations, nothing similar ever happened. I developed the software, and the audience used the software, but still, some events proceeded with a logic which none of us could control or even follow. If software art "finds itself outside of the…rational" (Goriunova and Shulgin 2002), then this was its opportunity to do something seemingly irrational. And it was our opportunity to enjoy it.

# 4. Final Thoughts

## 4.1 Towards A Perfect Interactive System?

In the previous chapter, I outlined several ways in which the interactive system at the heart of *Glimmer* could be incrementally improved if the piece were ever performed again. But for a moment, forget about real-world updates and imagine instead a perfect, fantasy-world system. Some yet-to-be-invented technologies would enable each of the 600 audience members to have a clear, identifiable role in shaping the music. Each of the 25 orchestral players would have a clear, identifiable role in further transforming the music. No one would need any experience or training to effectively participate. Everyone would be able to listen to the music and understand their critical role in creating it.

What would the music produced by such a system sound like? How would it combine all the disparate ideas of these people? Would anyone actually like the music it produced? Or would all of the input be averaged out to the blandest common denominator?

Dave Soldier and Komar & Melamid's *The People's Choice Music* (1997) offers an answer to these questions. Soldier uses the results of a web survey to write the "most wanted" and "most unwanted" songs. The former is a boring, generic love ballad, while the latter is an incredible mix of elements ranging from opera stars rapping about Wittgenstein to children singing holiday jingles advertising Wal-Mart.

*The People's Choice Music* must be taken with a grain of salt, of course; the project seems to have been designed from the beginning to lead to certain results. The survey questions focused on very general aspects of music (e.g. instrumentation,

duration, and dynamics) and extramusical elements (e.g. the subject matter of the lyrics), giving the interaction a level of superficiality[11] and giving Soldier considerable freedom in incorporating the survey results. What makes the most unwanted song so fascinating are not the responses to the survey questions, but rather Soldier's response to the survey results; he is able to cleverly include *everything* that people dislike about music in a single song. The focus is on the product; the process is secondary.

What would the music produced by a perfect interactive system sound like? Maybe this is not the right question to ask. Unlike *The People's Choice Music*, works such as *Glimmer* focus more on the process by which the music is created than on the finished product. Not even the perfect interactive system will give users who lack training, talent, and experience the instant ability to create masterpieces. And no system will magically resolve differences of opinion among a group of collaborators.

What the perfect interactive system can do is invite users into a space, encourage them to explore and push against the boundaries of that space, and, along the way, help them discover new things about themselves and about the way they perceive and participate in creative activities. The perfect interactive system, then, is not a system which creates perfect (or even necessarily good) music, but rather one which leads its users towards more fulfilling creative lives.

## 4.2 What's In It For Me?

Throughout this paper, I have justified the work I do through largely evangelical motives. I want to leverage the creativity of other people. I want to share the incredible

[11] The exchange reminds me, in fact, of my experience at the Virginia elementary school which I recounted in section 1.4.

experience of creating and performing music. I want to encourage people to approach music and creativity in different ways.

These are all important motivations for me, but they alone are not enough. Jonathan Harvey once said to me that writing music should be a joyful experience, and that simple advice has resonated with me. So where, then, do I find joy in writing music?

More than anything, I enjoy uncertainty and surprise in music. With every piece I write, there is a gap between the notated score and the experience of its performance. With training and experience, I have become better and better at using my imagination to bridge that gap. What I imagine in my head gets closer and closer to what I experience sitting in the concert hall.

But if that gap becomes too small, I no longer enjoy writing music. I want my imagination to take me to the point where I give control over the music to other people, and then I want to be amazed and surprised by the directions in which they take it. If things turn out exactly as I expect, then either I have become too controlling or I have chosen to work with people who lack initiative, creativity, and self-confidence. Neither situation is pleasant.

This is the most fundamental reason that I search for ways to give performers broad interpretive freedom, that I have never written a piece of tape music, and that I only began to really use technology in my work once I embraced the idea of software art. I enjoy being surprised, I enjoy being amazed, and I even enjoy being disappointed. But I do not enjoy making music alone.

# BIBLIOGRAPHY

Bianciardi, David, Tom Igoe, and Eric Singer. 2003. "Eos Pods: Wireless Devices for Interactive Musical Performance." *Proceedings of the Fifth Annual Conference on Ubiquitous Computing*, Seattle.

Cage, John. 1961. *Silence*. Middletown, CT: Wesleyan University Press.

Cage, John. 1969. *Notations*. New York: Something Else Press.

Cage, John. 1990. *Four*$^2$. London: Edition Peters, No. EP67368.

Carpenter, Loren, and Rachel Carpenter. 1999. "Audience Participation." *Ars Electronica: Facing the Future*, ed. T. Druckrey. Cambridge: MIT Press.

Chénard, Marc. 2000. "The Millenium Symphony: A Work for the Beginning of Time; Part I: The Musical Challenge." *La Scena Musical*, Vol. 5, No. 8, May 2000.

Duffy, Thomas C. 1995. *The Critics Choice*. New Haven, CT: self-published.

Fisher, Rob, Paul Vanouse, Roger Dannenberg and Jeff Christensen. 1997. "Audience Interactivity: A Case Study in Three Perspectives." *Proceedings of the Sixth Biennial Symposium for Arts and Technology*. Connecticut College.

Goriunova, Olga, and Alexei Shulgin. 2002. "Artistic Software for Dummies, and, by the way, Thoughts about the New World Order." *Read_Me Festival 1.2: Software Art Games*, Moscow. http://www.macros-center.ru/read_me/teb1e.htm.

Hasse, Jean. 1986. *Moths*. Euclid, OH: Visible Music.

Haydn, Franz Joseph. 1772. *Symphony No. 45*. Boca Raton, FL: Kalmus, No. A1556.

Kirkpatrick, John, ed. 1972. *Charles E. Ives: Memos*. New York: W.W. Norton and Company.

Luxemburg, Robert. 2003. *The Conceptual Crisis of Private Property as a Crisis in Practice*. http://rolux.net/crisis/index.php.

Madden, Timothy, Ronald Bruce Smith, Matthew Wright, and David Wessel. 2001. "Preparation for Interactive Live Computer Performance in Collaboration with a Symphony Orchestra." *Proceedings of the 2001 International Computer Music Conference*, Havana.

Maynes-Aminzade, Dan, Randy Pausch, and Steve Seitz. 2002. "Techniques for Interactive Audience Participation." *Proceedings of the IEEE International Conference on Multimodal Interfaces*, Pittsburgh, PA.

Mirapaul, Matthew. 2003. “A Hand-Held Portal to Musical Delights.” *The New York Times*, July 17, 2003.

NYU Movement Group. 2004. *SquidBall.* http://movement.nyu.edu/commotion/squidball.

Oteri, Frank. 1999. “Technology and the Future of New Music (interview).” *NewMusicBox*. New York: American Music Center, October 1999.

Riley, Terry. 1964. *In C*. San Francisco: SRI Moonshine Music.

Rzewski, Frederic. 1982. *Four North American Ballads*. Liege: self-published.

Schmit, Tomas. 1962. *Sanitas No. 35*. http://www.artnotart.com/fluxus/tschmit-sanitasno.35.html.

Sheridan, Molly. 2001. "Composer Calling: Cell Phone Symphony Premieres.” *NewMusicBox*. New York: American Music Center, October 2001.

Soldier, Dave, and Komar & Melamid. 1997. *The People’s Choice Music* (compact disc). New York: Dia Center for the Arts.

Stephenson, Neal. 1999. *Cryptonomicon*. New York: Avon Books.

Strickland, Edward. 1991. *American Composers: Dialogues on Contemporary Music*. Bloomington, IN: Indiana University Press.

Tommasini, Anthony. 2005. “Atonal, Edgy, and In Control.” *The New York Times*. April 9, 2005.

Varnik, Kristjan, Jason Freeman, Sekhar Ramakrishnan, Phil Burk, David Birchfield, and Max Neuhaus. 2004. “Tools Used While Developing Auracle: A Voice-Controlled, Networked Instrument.” *Proceedings of the 12th ACM International Conference on Multimedia*, New York.

Ward, Adrian. 1999. *AutoShop*. http://www.signwave.co.uk/go/products/autoshop.

Wolfram, Stephen. 2002. *A New Kind of Science*. Champaign, IL: Wolfram Media.

Verdi, Giuseppe. 1861. *La Forza Del Destino: Overture to the Opera*. London: Edition Eulenberg, No. 1107.

Vonnegut, Kurt. 1997. *Timequake*. New York: Berkley Books.