Cena concertante alla maniera di Vivaldi: Considering the Restaurant as a Musical Interface

Ben Houge
Berklee College of Music
1140 Boylston Street
Boston, MA 02215, USA
bhouge@berklee.edu

ABSTRACT
In recent years, I have been conducting research into the links between music and gastronomy by collaborating with chefs to develop multisensory dining experiences that I call "food operas." These events incorporate real-time music techniques adapted from the world of video game development to respond to the unpredictable events and timings of the dining room.

This paper details an event I developed in collaboration with the Boston Symphony Orchestra, chef David Verdo, and designer Jutta Friedrichs that took place on January 5 and 7, 2017. *Cena concertante alla maniera di Vivaldi* was a four-course meal with real-time musical accompaniment deployed from a seventy-channel speaker array comprised of sixty-four iPads and six near field studio monitors, all coordinated via the same network. The iPads were positioned in custom-built acoustic resonators placed at each seat in the restaurant, presenting a unique audio channel to each diner, synchronized to the rhythms of each diner's meal and sited as close as possible to the food. The music was based on Vivaldi's Piccolo Concerto in C Major, RV 443, drawing from archival BSO performances, with the objective of enhancing diners' appreciation of a live performance of the work on a concert following the meal. The menu was based on the music, drawing on research in the field of crossmodal psychology that identifies links between the senses of taste and hearing.

This paper discusses the background of the project, its musical organization, the infrastructure and control techniques required to execute it, and relevant research in the field of crossmodal psychology, concluding with a discussion of areas for future work.

Author Keywords
mobile devices, audiogustatory experiences, game audio

1. BACKGROUND
In a series of events dating back to 2010, I have been investigating the links between music and food by collaborating with chefs to develop multisensory dining experiences that I call "food operas."

The project draws from my extensive experience developing audio for video games since 1996, working for companies including Sierra, Microsoft, and Ubisoft. I considered a meal as a time-based art form, and I posited that the reason there is no culinary equivalent to ballet, theater, or opera combining music with gastronomy is due to the challenges inherent in synchronizing and siting music with a meal. I framed it as a technological problem, and used techniques from the world of video games to solve it. I developed a speaker array that presents a unique sound source at each seat in the restaurant and a music deployment system that renders a unique stream for each diner in real-time.

The focus of this project is not to arbitrarily introduce technology into the dining room, but to enable new forms of multisensory expression with the emergent activity of the dining room driving the music. Previous projects have explored crossmodal perception, engaged in sensory ethnography to tell the stories of local farmers, and limned the poetry of William Carlos Williams. The Saint Paul Food Opera, a commission from St. Paul-based new music ensemble Zeitgeist, presented a meal as an eight-movement chamber suite.

I of course acknowledge many antecedents to the sonic aspects of this project in work such as Golan Levin, Gregory Shakar, and Scott Gibbons's mobile ringtone piece *Dialtones (A Telesymphony)* [7]; the Princeton Laptop Orchestra (PLOrk) [13] and its mobile phone spin-offs [11]; and more recently Tristan Perich's *Microtonal Wall* [8] and Dajuin Yao's *Garden of Buddhahood* [4]. In addition, chefs such as Heston Blumenthal [1] and Paul Pairet [6] have been incorporating sound into the dining room, while artists such as Marina Abramovich [5] and Rirkrit Tiravanija [2] have been exploring the aesthetic possibilities of a meal. See Janice Wang's paper for more information on the history of multisensory dining practices. [14]

2. WORK DESCRIPTION
Cena concertante alla maniera di Vivaldi was developed for the Boston Symphony Orchestra as a special pre-concert dining event at Boston's historic Symphony Hall. I composed music based on Vivaldi's Piccolo Concerto in C Major, RV 443. Different aspects of the piece were highlighted at different times, sometimes overt themes, sometimes harmonic progressions, sometimes structural or textural details or ornamental ideas. The title of the event evokes the historical *sinfonia concertante* genre, while also referencing Johann Sebastian Bach's *Aria variata alla maniera italiana* BWV 989.

The goal of the event was threefold: to enhance appreciation for the Vivaldi work to be performed later that evening, to give people a sense of what it's like to be surrounded by sound inside an orchestra, and to explore the unique timbre of the piccolo.

The dining area was Symphony Hall's Choir Room. We set up six round tables, with about ten diners per table. An iPad was positioned at each place setting, set in a custom-built, walnut acoustic resonator designed and constructed by Jutta Friedrichs, which functioned as a stand while also refocusing the audio output of the built-in iPad speakers towards the diner.

The work exists in four movements. The first three movements correspond to the first three movements of the Vivaldi composition, and the fourth movement extends some ideas from the third movement while also introducing new material based on the overall tonality of the piece. A version of the fourth movement also serves as a kind of prelude as diners enter the room.

3. CROSSMODAL PSYCHOLOGY
As we slyly told guests at our event, Vivaldi designed the menu. More specifically, the menu was based on the character of Vivaldi's original composition as well as my conception of how the music could be adapted and manipulated, drawing out and focusing on specific attributes of the music.

There is a growing body of research in a field known as crossmodal psychology that investigates the connections between the senses. The central figure in this field is Dr. Charles Spence, head of the Crossmodal Research Laboratory at Oxford University and co-author, with Betina Piqueras-Fiszman, of *The Perfect Meal: The Multisensory Science of Food and Dining* [9], a compendium of work done at the intersection of gastronomy and sensory perception. Janice Wang is a researcher in Spence's group at Oxford, and prior to moving to Oxford, she assisted in the production of some of our Boston-area food operas. I visited Oxford in February 2015 to give a presentation on my work to Spence's group, and while I was in town, Janice and I organized an event pairing music and wine at a local wine bar with wine expert Alistair Cooper.

I began my research into music and food pairing before becoming aware of this stimulating area of psychology, but I have observed that many of the associations I made via intuition and observation have been corroborated through more rigorous psychological studies. In my current work, I draw from this emerging research as well as from my own experience and investigations.

I would characterize my work as incorporating a scoring approach, rather than a mapping approach. There are many aspects of a dish that can be emphasized or supported with sound, and it is an aesthetic choice to determine which aspect should be highlighted. My perspective is that music and sound can be used not merely to translate a dish into sound, but to highlight, enhance, comment, or even subvert it, just as music functions in cinema, for example. In this, it is related to the principles of wine pairing, in which pleasing contrast, rather than similarity, is often the desideratum.

Several associations formed the basis for the music and food pairings in this work. Sweet tastes are associated with high frequencies, and the piccolo is the instrument with the highest notes in the orchestra. In an oft-cited study, people rate food as tasting sweeter when listening to high frequency sounds and more bitter when listening to low frequency sounds. High frequency sounds are conveniently also well matched to the frequency response of an iPad's built in speakers. Sour tastes are also associated with high frequencies, as well as with fast tempos, and this piccolo concerto, like much of Vivaldi, contains many extremely rapid, virtuosic passages. [10]

Therefore I proposed a menu that based on the interplay of sweet and sour tastes. I suggested the primary ingredients and preparations for each course (e.g., salad with contrasting elements, pumpkin soup with mushroom, a fish-based dish with lemon and capers, creamy panna cotta with sour berries and mint), and chef David Verdo of Gourmet Caterers turned my general ideas into actual (very tasty) dishes. We had a round of revision on the written menu, then a tasting with a few tweaks to arrive at the final version of the menu. In earlier food opera events, I would often respond to what the chef proposed but in recent events, I have taken on the challenge of architecting the progression of a meal, considering it more as a succession of musical movements; this was especially important in my recent collaboration with Zeitgeist, as there were five chefs involved, none of whom had a perspective on the whole meal until very late in development.

A description of each course and its musical accompaniment follows:

The first course was a two-grain salad with quinoa, faro, goat cheese, olives, kumquats, cherry tomatoes, pine nuts, and herbs, with a mustard-citrus vinaigrette. As the meal begins, the lively opening phrases of the first movement of Vivaldi's piccolo concerto are played intermittently on diners' iPads. The timings and choices of phrases are different for each diner but synchronized to the same underlying pulse, resulting in a brisk, pointillistic texture of spatially deployed sound. To support this heterophonous musical behavior, I suggested a dish that had many different elements jumping out at the diner, asserting themselves at different times. The bright, sour elements of the tomato, kumquats, goat cheese, olives, and vinaigrette support the strongly articulated, high frequency piccolo phrases.

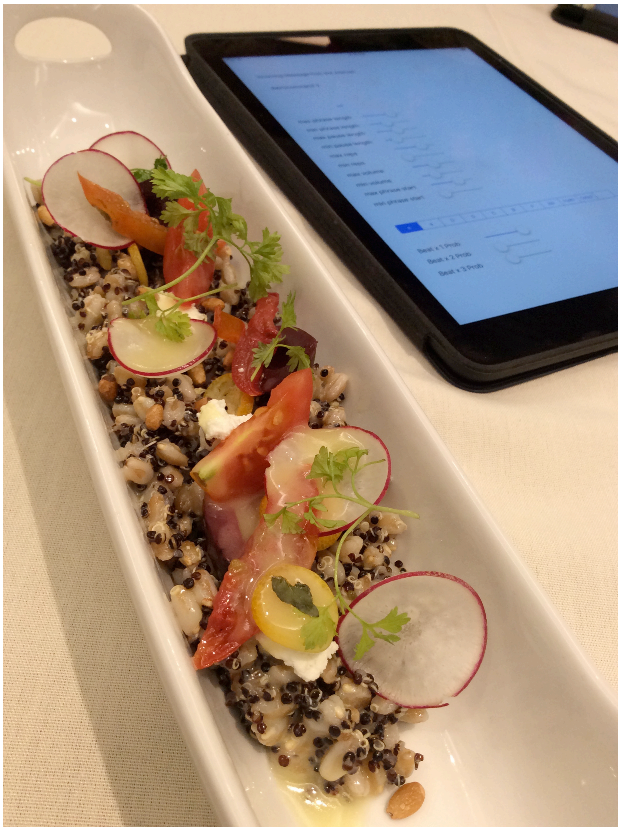

Figure 1. First Course.

The second course was a cream of pumpkin soup with spiced pepitas, pickled shimiji mushrooms, and toasted coconut. The second movement of the Vivaldi is quite different in character from the first, featuring long, legato, freely flowing, highly ornamented piccolo phrases over a hushed and sustained string accompaniment with a very slow harmonic rhythm. I generated new phrases on the fly by concatenating excerpts of a recording of the Boston Symphony Orchestra's previous performance of RV 443 in 1996, retaining this florid quality, which I sought to parallel with the smooth and creamy pumpkin soup, not only in its fat content (as fat causes taste to linger longer on the palette), but also in its flowing, liquid properties. These contactenated phrases were accompanied by intermittent, rolled harpsichord chords.

The third course was a roasted halibut fillet with crispy artichokes, bok choy, lemon caper butter, shaved radish, and petite soft herbs. The fish was reflected in a string chorale that played intermittently throughout the course, stable, steady, lightly voiced. Over the top, the iPads chirped a dense, percolating cloud of sharp, staccato phrases in the piccolo's upper register, evoking the bold, bracing tastes of the various accents, especially the lemon caper butter. The surround speakers played ostinato figures derived from Vivaldi's third movement and intermittent viola sforzandos based on the current chord.

The final course was a vanilla bean panna cotta with sour cherry, micro mint, and shaved espresso bean. The music for this course featured slow moving, consonant harmonies, stable chords articulated by intermittent piccolo notes played at a much lower pitch than the original recordings, resulting in a lush, mellow, hollow sound. The surround speakers played a low, subtly evolving drone. Over this background, short phrases excerpted from one of the third movement's rapid piccolo passages were played intermittently at a slower tempo using viola and flute key clicks samples, reflecting the sour berry and mint accents.

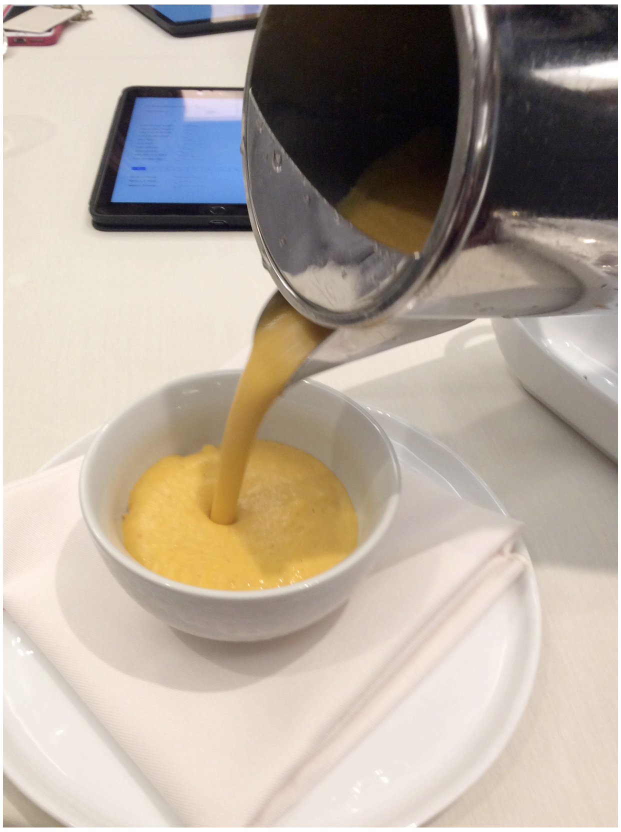

Figure 2. Second Course.

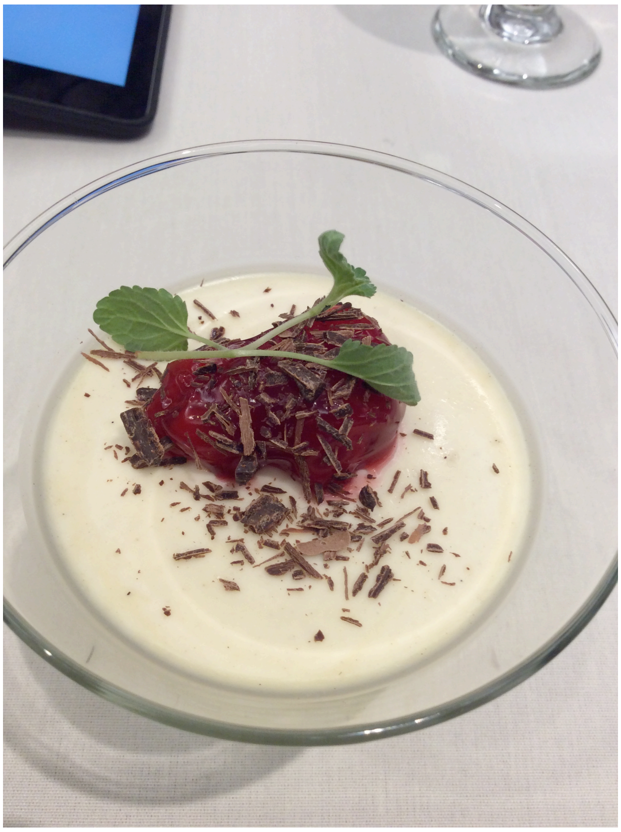

Figure 4. Fourth Course.

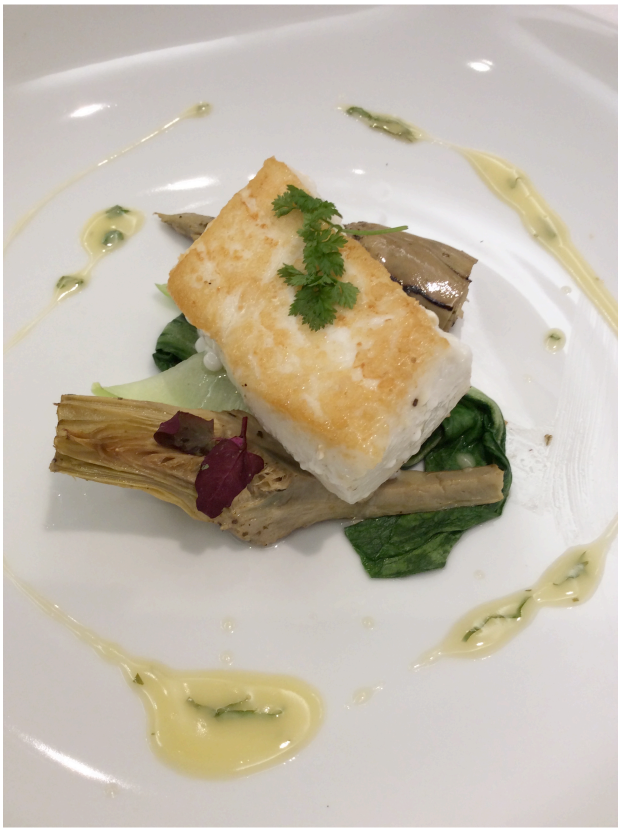

Figure 3. Third Course.

4. SETUP, INFRASTRUCTURE, AND CONTROL

The infrastructure for previous food opera projects consisted of several computers (networked together for synchronization) running custom software developed in Max. A speaker was positioned at each seat in the restaurant on a custom-built speaker stand by designer Jutta Friedrichs. The computers were connected to multichannel audio interfaces, which were connected to amplifiers, and hundreds of feet of speaker wire connected the amps to the speakers.

For *Cena concertante alla maniera di Vivaldi*, I wrote a completely new system from scratch, enabling us to take advantage of the phalanx of iPads owned by the Boston Symphony Orchestra. Switching to iPads allowed us to scale up to accommodate a larger number of guests than any previous project, while also eradicating the wires leading to each seat in the dining area. (Previously the largest group we had accommodated was thirty-six diners plus four additional speakers around the perimeter of the dining area at Studio Z in St. Paul, MN, for the *Saint Paul Food Opera*, a collaboration with new music ensemble Zeitgeist in October 2016.)

This new system was developed using three different programming languages. I developed a custom iPad app in Objective C, and I worked with the Boston Symphony Orchestra's IT department (led by Tim James) to deploy this app to the devices using the Meraki device management system. Once the devices were properly provisioned, this process took only a few minutes. I developed the server in JavaScript using Socket.IO and hosted it on Digital Ocean. And I used Max to control the sound of the six surround speakers (Mackie MR5's, connected the old fashioned way, via balanced cables to my MOTU UltraLite-mk3 audio interface).

The iPads were synchronized using an algorithm developed for multiplayer computer games that is related to the Simple Network

Time Protocol (SNTP). While testing, at the press of a button, all of the iPads could click together in synchronization with a common metronome. Max was also synchronized with the iPads using the same algorithm.

Once the iPads were launched, there was no further interaction with the devices. They functioned primarily as speakers but also displayed two different types of messages on the screen throughout the event. As each dish was set down, the screen would display an array of different words designed to encourage diners to think about the links between the food and the music. After a short pause, the screen would display the menu and a description of the music. The iPads also displayed the title of the event prior to the first course and a list of credits after the dessert course was cleared.

I originally designed the system to be controlled from my laptop from a station at the side of the room, but after running through the piece, I realized the room was too large for me to control everything from a stationary location, so I adapted the interface to run from my own iPad as I walked the floor, listened and observed. Via a custom web app, I was able to balance volumes per table and advance the music for each iPad individually, synchronizing a change in music with the new dish being set down in front of each diner. This resulted in a gradual crossfade from one musical texture to the next as servers made their way through the room.

In addition to changing the currently playing music for each diner, I could send global commands that would affect all diners at once. Pressing buttons on my web app could change the current chord, for example, or the current set of phrases from which to select. My web app also had a set of sliders that I could use to change various parameters, such as the minimum and maxim volume, pitch range, or number of repetitions of a certain musical parameter.

I attempted to keep a clear distinction between the commands sent from the server and the interpretation of the commands on the devices. I made two exceptions for data that was generated on the fly that needed to be consistent across all devices. In the second and fourth movements, harmonic progression was determined on the server and sent to all devices at once. In addition, in the second movement, new phrases were also generated on the server and broadcast to devices.

In other papers, I've written about the drawbacks of using native apps for distributed performance, but in this case, with sixty-four identical devices and the Meraki device management system (allowing us to bypass the Apple App Store with an app file signed for enterprise deployment), it was an ideal scenario. It should be noted that the system would not scale so well to a public event in which everyone brought their own device.

The key concept is that the emergent activity of the dining area is driving the music on both the macro and micro levels, as I observe the actions of diners and trigger musical change in response to my observations. From this perspective, due to the control scheme mapped out in advance, the interface onto the musical behaviors of the system is not the spontaneous desire or intuition of the person operating the control software, but the unpredictable actions of the entire room of participants.

5. ALGORITHMIC EXCERPTS

The types of algorithms and procedures employed in *Cena concertante alla maniera di Vivaldi* are quite simple but well suited to this type of work. Much of my work involves deploying multiple instances of simple behaviors with a bit of randomness permitted within certain bounds, and this work is a typical example.

These techniques derive directly from my work developing audio for video games since 1996. Video game audio is largely built around sample playback and manipulation (as opposed to pure synthesis), and consequently so is most of my music. Common game audio manipulations include varying the playback speed or amplitude of a sample within minimum and maximum bounds, allowing for each instance of a sound to play back slightly differently, providing a rough but efficient simulation of the type of subtle variation we are accustomed to hearing when interacting with the physical world. [3, 12] I often characterize the primary challenge of composing game music as first being able to respond to unpredictable events, and then being able to continue a behavior indefinitely between events.

These are the primary types of algorithmic techniques employed in this piece:

For each note to be played, choose from among a set of possible samples. For example, I had several different recordings of a piccolo playing a high G, so if I wanted to play a piccolo note, I would first choose one of these samples to use as a basis for subsequent manipulation.

For each note, choose a volume within minimum and maximum bounds. Also choose a start point in the sample buffer within minimum and maximum bounds. Depending on the sample, this might result in (very computationally efficient) changes to pitch, timbre, or volume.

To play a phrase (consisting of one or more notes in sequence), choose from among a set of valid phrases, play the phrase, and then choose a pause time between minimum and maximum bounds before choosing another phrase. These pauses may be free (measured in fractions of seconds) or synchronized to a pulse (measured in beats). Phrases may be repeated a certain amount of times (within minimum and maximum bounds). Phrases in a set may have different likelihoods of being played, and it is possible to interpolate between different sets of phrases over time in response to a command from the server.

The list of valid phrases may be fixed (as during the first course, which highlights selected phrases of Vivaldi's first movement), or it may be determined dynamically. For example, in Vivaldi's third movement, I identified a characteristically Baroque sequence of chords arpeggiated with a regular sixteenth-note pulse. The first sixteen notes outline e minor, the next sixteen a minor, the next sixteen D major, then G, then C, then B. So to play a desired chord, a phrase would be selected from the corresponding section of this sequence (e.g., selecting a subset of notes 17 through 32 for a minor, again within minimum and maximum bounds for phrase length). Each iPad would perform this behavior independently, resulting in a room full of the desired harmony, but built from different phrases deployed at different times. This same technique was applied at different points throughout the piece with different sample data and using different excerpts of the Vivaldi as source material to achieve a range of effects, sometimes as foreground material, sometimes as a background accent.

In the second movement, there are two mutually agnostic processes in play; each iPad is choosing which phrase to play from a set of six possible phrases, while at the same time, a separate process is intermittently choosing one of the six phrases at random to replace it with a newly generated phrase. This is a technique I have employed in a number of works; it allows for a degree of consistency, even though the material is randomly generated, and imbues the piece with a greater sense of coherence.

In the second and fourth sections of my work, first order Markov chains are employed to control the harmonic progression, allowing the piece to move through different key areas intermittently and avoid monotony. In the second movement, this progression is derived from an analysis of the harmonic of Vivaldi's second movement, while in the fourth movement the progression is a more arbitrary interpretation of the overall harmonic materials of Vivaldi's composition.

In the fourth movement, an underlying, low frequency drone is based on this harmonic motion, determining which note of the current chord to play, with different weightings for the root, third, or fifth.

The string chorale in the third movement plays on the six surround speakers, while the iPads chirp short phrases excerpted from Vivaldi, corresponding to the currently selected chord. It chooses a different string instrument sample for each note in the chord, and each note is played on a randomly determined speaker. As most chords contain only four notes, two notes are randomly duplicated to result in a note for each of the six speakers. The variation that results as a chord is randomly permutated in space, timbre, and doubling is rich and substantial.

As indicated at the outset, these behaviors are very simple, but as they are applied independently for each of the 70 speakers in the room, the result is a musical texture of great richness and complexity that is nonetheless consistent and controllable.

Much of the composition of this work involved solving questions of coordination between independent instances of these algorithms. If a phrase is repeated, should its randomly determined volume be the same for each repetition, or different each time? (I decided they should be different.) Should all of the iPads have the same set of randomly generated phrases, or should they each generate their own set of phrases? (I decided they should all use the same set.) In all of these decisions, I attempted to strike a balance between clear coherence and textural richness. (I am fascinated by questions of coordination between multiple instances of a behavior in aleatoric systems, and in this I derive great inspiration from the types of coordinated processes outlined by Christian Wolff in works like *Burdocks* from 1971.)

6. FUTURE WORK

The point of departure for my project was to use technology to achieve a new level of synchronization linking music and a meal. My current research, funded by a Berklee College of Music Newbury Comics Faculty Fellowship, seeks to increase the granularity of this synchronization to respond to individual gestures and even bites of each diner. I am investigating ways to meaningfully embed sensors into the implements of the dining room: silverware, plates, bowls, serving utensils, etc. I seek not only to respond to individual diners' immediate input, but also to extract information about the overall status of the dining room as it changes from moment to moment. This could have practical application in allowing restaurants to know when diners have stopped eating, either because they are finished or because they are dissatisfied, while also adapting the sound and visual environment to match the level of activity in the dining area.

By allowing diners to become the agents of musical expression in creating personalized soundtracks to their own meals, I hope to foster a heightened sense of awareness to the environment via all of the senses. Furthermore, I have observed that the unique speaker deployment system of these food opera events engenders a sense of community among diners, since the actions of each diner impacts the sonic experience of everyone else in the room; every participant is complicit in the experience. Especially given the current, factious state of world politics, I hope, in all of my activities, to engender communication and empathy between different people, and I feel that inviting people to come together around a meal, in some small way, works towards accomplishing this goal.

7. ACKNOWLEDGMENTS

Thanks to Jutta Friedrichs, my design collaborator; to Tim James, Sarah Manoog, Kim Nolthemy, and everyone else at the Boston Symphony Orchestra; to Chef David Verdo and Stephen F. Ponchak of Gourmet Catering for a fantastic meal; to Miles Gordon for additional audio editing and Matt Azevedo for acoustics advice; to Team Piccolissimo from Classical Music Hack Day at MIT, where portions of the second course music were prototyped: Julien Heller, Sophia Sun, Miles Gordon, Walter Werzowa, CheHo Lam, and Seda Röder; to Michael Bierylo and the EPD Department at Berklee; and to Berklee College of Music's Faculty Development Office.

8. REFERENCES

[1] Blumenthal, Heston. *The Fat Duck Cookbook*. London: Bloomsbury, 2009.

[2] Bourriaud, Nicolas. *Relational Aesthetics*. Dijon: Les presses du réel, 2002.

[3] Collins, Karen. *Game Sound: An Introduction to the History, Theory and Practice of Video Game Music and Sound Design*. Cambridge, The MIT Press, 2008.

[4] Delany, Ella. "The Power of Sound as an Art Form." *New York Times*, October 3, 2013. Accessed January 31, 2017. http://www.nytimes.com/2013/10/04/arts/international/The-Power-of-Sound-as-an-Art-Form.html

[5] Fabricant, Florence. "Marina Abramovic's Art Doubles as Dessert." *New York Times*, January 11, 2011. Accessed June 24, 2014. http://www.nytimes.com/2011/01/12/dining/12art.html

[6] Kessel, Jonah. "Ultra Dining at Ultraviolet." *New York Times*, October 15, 2013. Accessed June 24, 2014. http://www.nytimes.com/video/dining/100000002498301/ultra-dining-at- ultraviolet.html

[7] Levin, Golan. "An Informal Catalogue of Mobile Phone Performances, Installations and Artworks." Accessed May 25, 2015. http://www.flong.com/texts/lists/mobile_phone/

[8] The Museum of Modern Art. "Tristan Perich. Microtonal Wall. 2011." Accessed January 31, 2017. http://www.moma.org/explore/multimedia/videos/276/1331

[9] Spence, Charles, and Betina Piqueras-Fiszman. *The Perfect Meal: The Multisensory Science of Food and Dining*. Oxford: Wiley Blackwell, 2014.

[10] Spence, Charles, and Qian Janice Wang. "Wine and Music (III): So What if Music Influences the Taste of the Wine?" *Flavour* 4 (2015): 36. doi: 10.1186/s13411-015-0046-9.

[11] The Stanford Mobile Phone Orchestra. "Stanford Mobile Phone Orchestra (MoPhO)." Accessed May 25, 2015. http://mopho.stanford.edu/

[12] Sweet, Michael. *Writing Interactive Music for Video Games: A Composer's Guide*. Upper Saddle River, NJ: Addison-Wesley, 2015.

[13] Trueman, Dan. "Why a Laptop Orchestra?" *Organized Sound* 12/2 (August 2007): 171-179.

[14] Wang, Qian (Janice). "Music, Mind, and Mouth: Exploring the Interaction between Music and Flavor Perception." MS diss., Massachusetts Institute of Technology, 2013.