

## A FEEDBACK CANCELING REVERBERATOR

Jonathan S. Abel, Eoin F. Callery, and Elliot K. Canfield-Dafilou

Center for Computer Research in Music and Acoustics,  
Stanford University, Stanford, CA 94305 USA  
abel|ecallery|kermit@ccrma.stanford.edu

### ABSTRACT

A real-time auralization system is described in which room sounds are reverberated and presented over loudspeakers. Room microphones are used to capture room sound sources, with their outputs processed in a canceler to remove the synthetic reverberation also present in the room. Doing so suppresses feedback and gives precise control over the auralization. It also allows freedom of movement and creates a more dynamic acoustic environment for performers or participants in music, theater, gaming, and virtual reality applications. Canceler design methods are discussed, including techniques for handling varying loudspeaker-microphone transfer functions such as would be present in the context of a performance or installation. Tests in a listening room and recital hall show in excess of 20 dB of feedback suppression.

### 1. INTRODUCTION

Real-time virtual acoustic/auralization systems have been made possible by advances in signal processing and acoustics measurement. Computational methods for simulating reverberant environments are well developed, and these auralization systems process sound sources according to impulse responses encapsulating the acoustics of the desired space and render them over loudspeakers in the venue or through headphones [1, 2]. In live and recoding settings, close mics or contact mics are commonly used for acoustic instruments and voice to avoid feedback. Such mic'ing can be cumbersome and can affect or restrict performances. In virtual, augmented, or mixed reality settings, the immersive audio possibilities are similarly restricted by the use of headphones. Moreover, in all of these situations, unless the locations of the sound sources are tracked, movement in the virtual space will not be reflected in the experienced auralization.

In recent years, several systems which use room microphones and loudspeakers have been developed to create virtual reverberant auralizations. These include products by Meyer Sound [3] and Lexicon [4] as well as the system designed by Woszczyk [5] at McGill. See [6] for a more extensive review. In such systems, a number of approaches have been used to suppress feedback, including adaptive notch filtering to detect and suppress individual frequencies as they initiate feedback [7], frequency shifting the synthesized acoustics [8], varying the synthesized acoustics over time [4], and decorrelating the various auralization impulse responses [9, 10]. Such processing compromises the original dry signals. In addition, to provide the needed control and to achieve the best possible performance, these systems are typically built from the ground up using proprietary hardware and software. Accordingly, they do not take advantage of existing loudspeaker and microphone arrays already present on site. Ultimately, this makes these systems expensive, involving significant alteration to the installation site, and requiring prolonged calibration and tuning.




Figure 1: Feedback Canceling Auralization System. Room sounds are convolved with an auralization impulse response  $h(t)$ , generating simulated acoustics  $l(t)$  which are projected into the room via a loudspeaker. A room microphone captures both room sounds and simulated acoustics  $m(t)$ , and is processed according to measurements of the loudspeaker-microphone transfer function to remove the simulated acoustics, thus leaving an estimate of the room sounds  $\hat{d}(t)$  to be auralized.

Here, we present a system for real-time auralization that uses standard room microphones and loudspeakers, and employs signal processing tools to cancel the feedback, thus eliminating the need for close or contact microphones. The cancellation method described here is similar to the adaptive noise cancellation approach developed by Widrow [11] for removing unwanted additive noise from a signal. In that approach, a reference signal, which is correlated with the unwanted noise, is used to estimate and subtract the unwanted noise from the primary signal. Related literature also includes echo cancellation and dereverberation [12–14].

The system we describe can also be integrated into existing speaker arrays as it does not require proprietary hardware and can be implemented using inexpensive and readily available software. The system is designed to be easy to configure and straightforward to calibrate. The ease of use and mobility afforded by not requiring close mic'ing creates opportunities for dynamic artistic experiences for performers and audiences in disciplines such as music, theater, dance, and emerging digital art forms [15]. For example, in virtual, augmented, and mixed reality scenarios, the system allows users to dispense with headphones for more immersive virtual acoustic experiences.

In the sequel, the system and cancellation processing are described. Example applications and a performance analysis follow.

## 2. AURALIZATION SYSTEM

We begin by describing the auralization system, which is similar to the recording processing described in [16].

Referring to Fig. 1, a room microphone captures contributions from room sound sources  $d(t)$  and synthetic acoustics produced by the loudspeaker according to its applied signal  $l(t)$ , with  $t$  being the discrete time sample index. One can impart the sonic characteristic of a space,  $h(t)$ , on the room sounds  $d(t)$  through convolution,

$$l(t) = h(t) * d(t). \quad (1)$$

Many auralization systems work this way, using fast, low-latency convolution methods to save computation [17–19]. The difficulty is that the room source signals  $d(t)$  are not directly available. As described above, the room microphones also pick up the synthesized acoustics, and would cause feedback if the room microphone signal  $m(t)$  were reverberated without additional processing.

Here, we auralize an estimate of the dry signal  $\hat{d}(t)$ , formed by subtracting from the microphone signal  $m(t)$  an estimate of the synthesized acoustics. Assuming the geometry between the loudspeaker and microphone is unchanging, we have

$$d(t) = m(t) - g(t) * l(t), \quad (2)$$

where  $g(t)$  is the impulse response between the loudspeaker and microphone. Here, we design an impulse response  $c(t)$ , which approximates the loudspeaker-microphone response, and use it to form an estimate of the "dry" signal  $\hat{d}(t)$ ,

$$\hat{d}(t) = m(t) - c(t) * l(t). \quad (3)$$

This is shown in the signal flow diagram Fig. 1: the synthetic acoustics are canceled from the microphone signal  $m(t)$  to estimate the room signal  $\hat{d}(t)$ , which is then reverberated.

### 2.1. Canceler Design

The question then becomes how to design the canceling filter  $c(t)$ . A measurement of the impulse response  $g(t)$  provides an excellent starting point, though there are time-frequency regions over which the response is not well known due to measurement noise (typically affecting the low frequencies) or changes over time due to air circulation or performers, participants, or audience members moving about the space (typically later in the impulse response). In regions where the impulse response is not well known, the cancellation should be reduced so as to not introduce additional reverberation.

Here, we choose the cancellation filter impulse response  $c(t)$  to minimize the expected energy in the difference between the actual and estimated room microphone loudspeaker signals. For simplicity of presentation, for the moment let us assume that the loudspeaker-microphone impulse response is a unit pulse,

$$g(t) = g \delta(t), \quad (4)$$

and that the impulse response measurement  $\tilde{g}(t)$  is equal to the sum of the actual impulse response and zero-mean noise with variance  $\sigma_g^2$ . Consider a canceling filter  $c(t)$  which is a windowed version of the measured impulse response  $\tilde{g}(t)$ ,

$$c(t) = w \tilde{g} \delta(t). \quad (5)$$

In this case, the measured impulse response  $\hat{g}(t)$  is scaled according to a one-sample-long window  $w$ . The expected energy in the difference between the auralization and cancellation signals at time  $t$  is:

$$\mathbb{E} [(g l(t) - w \tilde{g} l(t))^2] = l^2(t) [w^2 \sigma_g^2 + g^2(1 - w)^2]. \quad (6)$$

Minimizing the residual energy over the window  $w$ , we find

$$c^*(t) = w^* \tilde{g} \delta(t), \quad w^* = \frac{g^2}{g^2 + \sigma_g^2}, \quad (7)$$

a Wiener-like weighting of the measured impulse response. When the loudspeaker-microphone impulse response magnitude is large compared with the impulse response measurement uncertainty, the window  $w$  will be near 1, and the cancellation filter will approximate the measured impulse response. By contrast, when the impulse response is poorly known, the window  $w$  will be small—roughly the measured impulse response signal-to-noise ratio—and the cancellation filter will be attenuated compared to the measured impulse response. In this way, the optimal cancellation filter impulse response is seen to be the measured loudspeaker-microphone impulse response, scaled by a compressed signal-to-noise ratio (CSNR).

Typically, the loudspeaker-microphone impulse response  $g(t)$  will last hundreds of milliseconds, and the window that scales the measured impulse response will preferably be a function of time  $t$  and frequency  $f$  so as to account for changes in impulse response variance over time and frequency. Denote by  $\tilde{g}(t, f_b)$ ,  $b = 1, 2, \dots, N$  the measured impulse response  $\tilde{g}(t)$  split into a set of  $N$  discrete frequency bands  $f_b$  using a filterbank such that the sum of the band responses is the original measurement,

$$\tilde{g}(t) = \sum_{b=1}^N \tilde{g}(t, f_b). \quad (8)$$

In this case, the canceler response  $c^*(t)$  is the sum of measured impulse response bands  $\tilde{g}(t, f_b)$ , scaled in each band by a corresponding window  $w^*(t, f_b)$ . Expressed mathematically,

$$c^*(t) = \sum_{b=1}^N c^*(t, f_b), \quad (9)$$

where

$$c^*(t, f_b) = w^*(t, f_b) \tilde{g}(t, f_b), \quad (10)$$

$$w^*(t, f_b) = \frac{g^2(t, f_b)}{g^2(t, f_b) + \sigma_g^2(t, f_b)}. \quad (11)$$

We suggest using the measured impulse response bands  $\tilde{g}(t, f_b)$  as stand-ins for the actual impulse response bands  $g(t, f_b)$  in computing the optimal window  $w^*(t, f_b)$ . In addition, repeated measurements of the impulse response  $g(t, f_b)$  could be made, with the measurement mean used for  $g(t, f_b)$ , and the variation in the impulse response measurements as a function of time and frequency used to form  $\sigma_g^2(t, f_b)$ . We also suggest smoothing  $g^2(t, f_b)$  over time and frequency in computing  $w(t, f_b)$  so that the window is a smoothly changing function of time and frequency.




Figure 2: Canceling auralizer using multiple loudspeakers and microphones. Multiple loudspeakers and microphones can be accommodated in this auralizer architecture by estimating the matrix of loudspeaker-microphone transfer functions,  $\mathbf{G}(t)$ . Additionally, the room sound estimates may be processed using beamforming or other techniques before being diffused about the space.

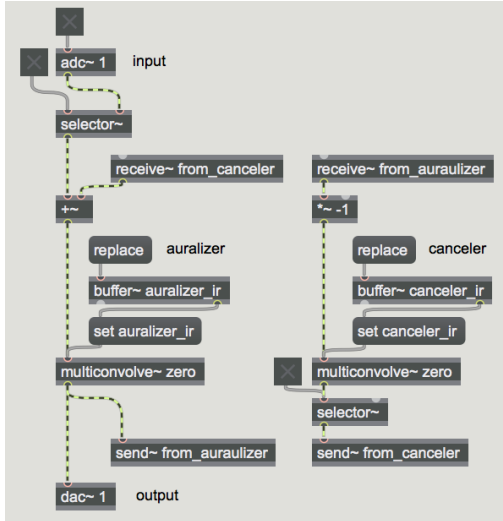


Figure 3: Max/MSP patch showing one possible implementation of the auralization system.

## 2.2. Multiple Microphones and Speakers

In the presence of  $L$  loudspeakers and  $M$  microphones, a matrix of loudspeaker-microphone impulse responses is measured, and used in subtracting auralization signal estimates from the microphone signals. Stacking the microphone signals into an  $M$ -tall column  $\mathbf{m}(t)$ , and the loudspeaker signals into an  $L$ -tall column  $\mathbf{l}(t)$ , our cancellation system becomes

$$\mathbf{l}(t) = \mathbf{H}(t) * \hat{\mathbf{d}}(t), \quad (12)$$

$$\hat{\mathbf{d}}(t) = \mathbf{m}(t) - \mathbf{C}(t) * \mathbf{l}(t), \quad (13)$$

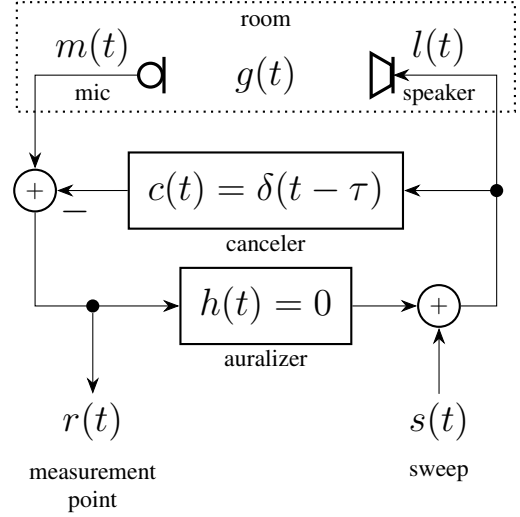


Figure 4: Canceling Auralizer Calibration. The cancellation processing  $c(t)$  may be determined by measuring the impulse response between the loudspeaker and microphone, simultaneously with the response through  $c(t)$ .

where  $\mathbf{H}(t)$  is the matrix of auralizer filters and  $\mathbf{C}(t)$  the matrix of canceling filters. As in the single-speaker single-microphone case, the canceling filter matrix is the matrix of measured impulse responses, each windowed according to its respective CSNR.

Moreover, a conditioning processor,  $\mathbf{Q}$ , can be inserted between the microphones and auralizers,

$$\mathbf{l}(t) = \mathbf{H}(t) * \mathbf{Q}(\hat{\mathbf{d}}(t)), \quad (14)$$

$$\hat{\mathbf{d}}(t) = \mathbf{Q}(\hat{\mathbf{d}}(t)) - \mathbf{C}(t) * \mathbf{l}(t), \quad (15)$$

as seen in Fig. 2. This processor could serve several functions. First,  $\mathbf{Q}$  could act as the weights of a mixing matrix to determine how the microphones signals are mapped to the auralizers, and subsequently, the loudspeakers. For example, it might be beneficial for microphones that are on one side of the room to send the majority of their energy to loudspeakers on the same side of the room, as could be achieved using a B-format microphone array and Ambisonics processing driving the loudspeaker array. Another use could be for when the speaker array and auralizers are used to create different acoustics in different parts of the room. The processor  $\mathbf{Q}$  could also be a beamformer or other microphone array processor to auralize different sounds differently according to their source position. In such a situation,  $\mathbf{Q}$  could change the dimensionality of the  $M$  microphone signals into  $P$  signals which are then auralized. It is worth noting that depending on the purpose,  $\mathbf{Q}$  could be a matrix of weights, a matrix of convolutions, a combination of the two, or other processor.

## 3. IMPLEMENTATION AND EVALUATION

### 3.1. MaxMSP Implementation and System Calibration

The signal flow of Fig. 2 is straightforward to implement in any number of environments. A Max/MSP implementation of a single-




Figure 5: Cancellation Processing Design. The cancellation processor reproduces the impulse response between the loudspeaker and microphone, accounting for the scaling and delay experienced through the canceler convolution.




Figure 6: Example Cancellation Impulse Response. A cancellation impulse response  $c(t)$  (top) and its associated spectrogram (bottom) are shown for the Listening Room at CCRMA, Stanford University, configured with a ceiling-mounted full-range loudspeaker and hanging omnidirectional microphone.

microphone, single-loudspeaker canceling auralizer is shown in Fig. 3. We use [20] for fast convolution.

To calibrate the system, the canceler impulse response  $c(t)$  was set to a delayed pulse and the impulse response of the system configured as shown in Fig. 4 was used to determine the scaling and delay through the Max/MSP patch and to measure the loudspeaker-microphone transfer function. An example result, using a Sennheiser MKH 20-P48 omnidirectional microphone placed about 50 cm from an Adam A8X full-range loudspeaker is shown in Fig. 5. To find  $c(t)$ , the measured impulse response  $\tilde{g}(t)$  is shifted and scaled according to the amplitude and arrival time of the  $c(t) = \delta(t - \tau)$  pulse. An example canceler impulse response is shown in Fig. 6. Finally, note that an optimal window may be

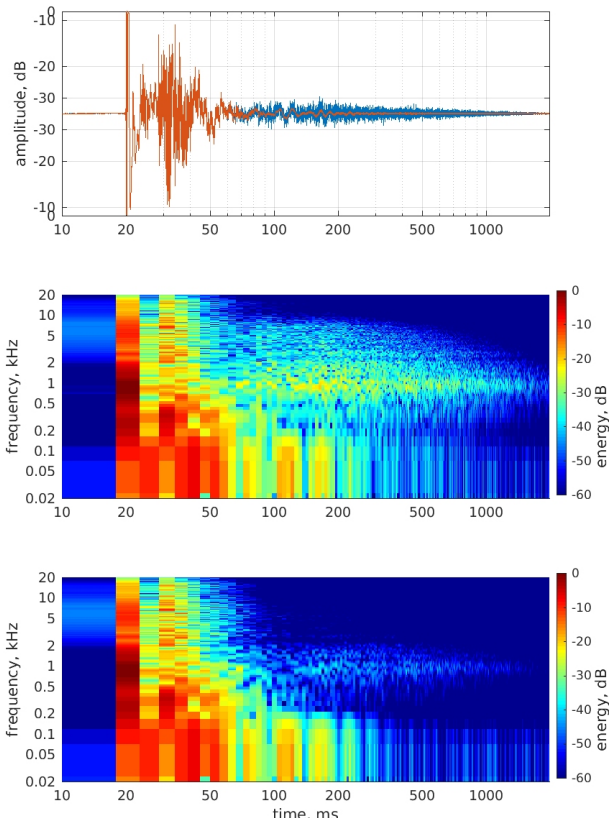


Figure 7: Canceling Auralizer Room Impulse Response. A sine sweep from a separate loudspeaker in the room was used to measure the impulse response between a room source and the canceling reverberator system microphone input (top, blue), and system room source estimate (top, orange). The corresponding spectrograms are also shown (middle and bottom). Note that the room impulse response contains both the “dry” room response and the “wet” synthesized room acoustics (Memorial Church at Stanford University), while the estimated room source response shows a substantially drier signal.

applied according to the discussion above by making a number of measurements, and estimating the variance of the measured impulse responses as a function of time and frequency.

### 3.2. Performance Evaluation

It is useful to anticipate the effectiveness of the virtual acoustics cancellation in any given microphone. Substituting the optimal windowing (7) into the expression for the canceler residual energy (6), the virtual acoustics energy in the canceled microphone signal is expected to be scaled by a factor of

$$\nu = \frac{\sigma_g^2}{g^2 + \sigma_g^2}, \tag{16}$$




Figure 8: Canceling Auralizer Example. A dry source, Suzanne Vega’s “Tom’s Diner,” was played in the CCRMA Listening Room, configured with the canceling auralizer described here. Spectrograms are shown for the microphone signal (top), the room signal estimate (middle), and the synthetic acoustics projected into the room (bottom). The room signal estimate contains little of the synthetic reverberation, and is effectively a mix of the dry Suzanne Vega track and low-frequency ventilation noise present in the room.

compared to that in the original microphone signal. Note that the reverberation-to-signal energy ratio is improved in proportion to the measurement variance for accurately measured signals,  $\sigma_g^2 \ll g^2$ . By contrast, when the impulse response is inaccurately measured, the reverberation-to-signal energy ratio is nearly unchanged,  $\nu \approx 1$ .

To evaluate the performance of the system, we implemented several versions of the system shown in Fig. 2 with one–two microphones and one–four loudspeakers in the CCRMA Listening Room and CCRMA Stage recital hall at Stanford University. We used a single loudspeaker source, playing exponentially swept sinusoid test signals and Suzanne Vega’s “Tom’s Diner” as dry program material. This was selected as it often used to test reverberators and makes for a repeatable test signal.




Figure 9: Room Impulse Response Variation. The mean room impulse response (top) formed from 1145 sine sweep responses measured between a loudspeaker and microphone mounted in the CCRMA Stage, a 120-seat recital hall at Stanford University that was unoccupied during the measurement. The impulse response energy, smoothed over a 10-millisecond-long Hanning window, is also shown (bottom, solid), along with the smoothed sample standard deviation (bottom, dashed). The smoothed sample standard deviation is also shown for a set of 75 measurements made with a dozen occupants near the loudspeaker and microphone, and in different positions for each measurement (bottom, dotted). Note that the impulse response variation is smallest relative to the impulse response energy near the beginning of the impulse response, and that the variation for the occupied room is modestly larger as the room becomes mixed.

In a first test, the impulse response of the room with the system active is measured. As seen in Fig. 7, the room impulse response contains both the “dry” room response and the “wet” synthesized room acoustics of Memorial Church at Stanford University. The 4.5 s reverberation time is plainly visible. Also shown in Fig. 7 is the system dry signal estimate,  $\hat{d}(t)$ . Compared to the virtual room impulse response, the canceler produces a substantially dry signal, canceling in excess of 30 dB of the simulated reverberation.

Fig. 8 shows the response of the system to a dry source, Vega’s “Tom’s Diner.” Spectrograms are shown for the microphone signal, the room signal estimate, and the synthetic acoustics projected into the room. Note that the room signal estimate contains little of the synthetic reverberation, and is effectively a mix of the dry Suzanne Vega track, and low-frequency ventilation noise present in the room. As expected, the room response shows the imprint of the Memorial Church acoustics, as added by the system.

To better understand the practical performance of the system, repeated measurements of the loudspeaker-microphone response were made at the CCRMA Stage in unoccupied and occupied conditions. Fig. 9 shows the mean room impulse response and the impulse response energy, smoothed over a 10-millisecond-long Hanning window. The sample standard deviation is shown separately for the unoccupied and occupied conditions. The impulse response variation is smallest relative to the impulse response energy near the beginning of the impulse response. Also, the variation for the occupied room is modestly larger as the room becomes mixed. As




Figure 10: Canceler Performance Example. The smoothed energy of the mean loudspeaker-microphone impulse response is shown (blue), as is the residual energy of suppressed loudspeaker signals for the unoccupied (yellow) and occupied (orange) rooms. Note that the cancellation is most effective at the impulse response start, during which there is little variation, cf. Fig. 9. Note also that the occupied room has a slightly larger residual energy as the room is becoming well mixed.

seen in Fig. 10, the canceler residual energy is small near the beginning of the response, and increases relative to the decreasing impulse response energy throughout the response, consistent with the notion that the beginning of the impulse response shows little variation. As seen in Fig. 11, the canceler residual energy is also small for frequencies below about 2 kHz. Over the speech band of 200 Hz–3200 Hz, the residual simulated acoustics energy present in the room signal estimate  $\hat{d}(t)$  was 16.4 dB for the occupied CCRMA Stage with moving participants, and 24.3 dB for the unoccupied CCRMA Stage.

Finally, we present an example of the ability of the system to suppress feedback resulting from creating a wet synthetic acoustic environment. Fig. 12 shows a spectrogram of a recording of the canceling auralizer simulating Memorial Church, along with the spectrogram of the same recording, but with the canceler component of the system switched off, and then switched back on. Note the rapid build up and subsequent suppression of feedback with the temporary removal of the cancellation processing.

#### 4. CONCLUSIONS AND FUTURE WORK

In this paper, we have shown a real-time auralization system capable of generating multiple auralizations while canceling synthetic reverberation with greater than 20 dB of feedback suppression. We have also shown that the system can be calibrated and integrated into an existing speaker array, using currently available room microphones to pick up live sounds, and function in real-time by running with off-the-shelf software. Importantly, our system allows flexible and dynamic experiences for performers, audiences, and other users. In theatrical, musical, or other live performance situations, this system does not require performers to wear microphones, transmitters, or battery packs in order to be processed through artificial reverberation, thus expanding performance pos-




Figure 11: Canceler Performance Example, Residual Energy. The loudspeaker-microphone impulse response spectrogram (top) is shown along with the root-mean-square canceler residual for the unoccupied CCRMA Stage (middle) and occupied CCRMA Stage (bottom). Note that a substantial amount of the loudspeaker energy has been canceled, particularly at the impulse response beginning and for frequencies below about 2 kHz.

sibilities. Similarly, in emergent virtual, augmented, or mixed reality settings, such as those one might find in industrial simulations, home entertainment systems, and artistic installations, our system does not require use of headphones to facilitate immersive auralizations.

We have tested our system in several rooms at CCRMA, Stanford University. Additionally, the system has been used in a series of network-audio concerts between Stanford University and Stockholm, Sweden [21]. We are planning to continue to develop our system for further electroacoustic music works, for virtual reality and virtual acoustic research, music and theatrical rehearsals and performances, art installations, and for other academic and industrial research projects at Stanford University. In particular, we are installing a larger system (4 microphones and 8–16 loudspeakers) for a study of performance practice using vocal repertoire written for Rome’s *Chiesa di Sant’Aniceto* using impulse responses recorded in Rome during 2017–18, [22, 23].




Figure 12: Feedback Example. A spectrogram of a recording of the canceling auralizer simulating the 5-second-long reverberation of Memorial Church at Stanford University is shown (bottom), along with the spectrogram of the same segment, but with the canceler component of the system switched off near 500 ms, and switched back on a little after 3000 ms (top). Note the rapid build up and subsequent suppression of feedback near 1800 Hz with the temporary removal of the cancellation processing.

## 5. REFERENCES

- [1] Michael Vorländer, *Auralization: Fundamentals of Acoustics, Modelling, Simulation, Algorithms and Acoustic Virtual Reality*, Springer, 2007.
- [2] Mendel Kleiner, Bengt-Inge Dalenback, and Peter Svensson, “Auralization—an overview,” *Journal of the Acoustical Society of America*, vol. 41, no. 11, pp. 861–75, November 1993.
- [3] Meyer Sound, “Constellation acoustic system,” <https://meyersound.com/product/constellation/>, 2006.
- [4] David Griesinger, “Improving room acoustics through time-variant synthetic reverberation,” in *Proceedings of the 90th Audio Engineering Society Convention*, 1991.
- [5] Wieslaw Woszczyk, Tom Beghin, Martha de Francisco, and Doyuen Ko, “Development of virtual acoustic environments for music performance and recording,” in *Proceedings 25th Tonmeisterstagung VDT International Convention*, 2008.
- [6] Mark A. Poletti, “Active acoustic systems for the control of room acoustics,” in *Proceedings of International Symposium on Room Acoustics*, 2010.
- [7] Shang Li, Nobuaki Takahashi, and Tsuyoshi Takebe, “Fast stabilized adaptive algorithm for IIR bandpass/notch filters for a single sinusoid detection,” *Electronics and Communications in Japan (Part III: Fundamental Electronic Science)*, vol. 76, no. 8, pp. 12–24, 1993.
- [8] Manfred R. Schroeder, “Improvement of acoustic-feedback stability by frequency shifting,” *The Journal of the Acoustical Society of America*, vol. 36, no. 9, pp. 1718–24, 1964.
- [9] Mark A. Poletti and Roger Schwenke, “Prediction and verification of powered loudspeaker requirements for an assisted reverberation system,” in *Proceedings of the 121st Audio Engineering Society Convention*, 2006.
- [10] Jonathan S. Abel, Nicholas J. Bryan, Patty P. Huang, Miriam Kolar, and Bissera V. Pentcheva, “Estimating room impulse responses from recorded balloon pops,” in *Proceedings of the 129th Audio Engineering Society Convention*, 2010.
- [11] Bernard Widrow, J. R. Glover, J. M. McCool, J. Kaunitz, C. S. Williams, R. H. Hearn, J. R. Zeidler, J. Eugene Dong, and R. C. Goodlin, “Adaptive noise cancelling: Principles and applications,” *Proceedings of the IEEE*, vol. 63, no. 12, pp. 1692–716, Dec 1975.
- [12] Emanuël Habets, “Fifty years of reverberation reduction: From analog signal processing to machine learning,” AES 60th Conference on DREAMS, 2016.
- [13] Patrick A Naylor and Nikolay D Gaubitch, Eds., *Speech Dereverberation*, Springer, 2010.
- [14] Francis Rumsey, “Reverberation... and how to remove it,” *Journal of the Acoustical Society of America*, vol. 64, no. 4, pp. 262–6, April 2016.
- [15] Paul DeMarinis, “Personal communication,” 2014.
- [16] Jonathan S. Abel and Elliot K. Canfield-Dafilou, “Recording in a virtual acoustic environment,” in *Proceedings of the 143rd Audio Engineering Society Convention*, 2017.
- [17] William G. Gardner, “Efficient convolution without latency,” *Journal of the Audio Engineering Society*, vol. 43, 1993.
- [18] Guillermo Garcia, “Optimal filter partition for efficient convolution with short input/output delay,” in *Proceedings of the 113th Audio Engineering Society Convention*, 2002.
- [19] Frank Wefers and Michael Vorländer, “Optimal filter partitions for real-time FIR filtering using uniformly-partitioned FFT-based convolution in the frequency-domain,” in *Proceedings of the 14th International Conference on Digital Audio Effects*, 2011, pp. 155–61.
- [20] Alexander Harker and Pierre Alexandre Tremblay, “The HISSTools impulse response toolbox: Convolution for the masses,” in *Proceedings of International Computer Music Conference*, 2012.
- [21] Leif Handberg, Ludvig Elblaus, Chris Chafe, and Elliot Kermit Canfield-Dafilou, “Op 1254: Music for neutrons, networks and solenoids using a restored organ in a nuclear reactor,” in *Proceedings of the Twelfth International Conference on Tangible, Embedded and Embodied Interactions*, 2018.
- [22] Jonathan Berger, “Reanimating the music of *La Chiesa di Sant’Aniceto a Palazzo Altemos*, in Rome,” in *QMUL School of Electronic Engineering and Computer Science Distinguished Lecturer Series*, 2017.
- [23] Jonathan Berger, Jonathan S. Abel, Talya Berger, and Elliot K. Canfield Dafilou, “Timbre, texture, space and musical style: Rome’s *Chiesa di Sant’Aniceto* and its music,” in *Timbre is a Many-Splendored Thing*, 2018.