Reverberated Averages

Marc Evans
University of California, Santa Barbara
Santa Barbara, CA, USA
marc.p.evans@gmail.com

I. Introduction

Jim Campbell's "Illuminated Averages" (2000-2009) is a series of works exploring the compression of temporally ordered images into a single static image [1]. The most well-known pieces in the series ("Hitchcock's Psycho", "Welles' Citizen Kane Breakfast Table Sequence", and "Fleming's The Wizard of Oz") use feature-length films as their source, averaging all of the frames of the film together to form a hazy, gestural image that nevertheless reveals much about the action and visual palate of the film. Other pieces in the series explore the averaging of different sorts of subject matter: a series of still photos of a schoolyard, the pages of "Wuthering Heights", and a performance of a Bach Cello Suite. The series as a whole explores the compression of time in fundamentally temporal media. In doing so, the pieces probe the nature and contribution of time in the original work, as well as the "relationship of time to space" [2].

Other artists have worked along similar lines, notably Hiroshi Sugimoto and Douglas Gordon. Sugimoto's "Theaters" is a series of long-exposure photographs of the inside of movie theaters, in which the area of the screen becomes an undifferentiated, over-saturated white, while the inner space of the cinema comes clearly into view. In this way, the works in the series invert the movie-going experience, emphasizing space over time [3]. Gordon obliterates time in a different way; instead of condensing it, he stretches it out to a near stand-still in works like "24-hour Psycho" and "Five Years Drive By" [4].

Similar work has been done in the audio domain, notably Leif Inge's "9 Beet Stretch," which turns Beethoven's 9th symphony into a slowly evolving ambient soundscape. However, while this extreme temporal stretching has been explored, the complete compression of time exhibited by Campbell's "Illuminated Averages" has not been explored in the audio domain. "Reverberated Averages" represents such a venture.

II. Exposition

Reverberated averages uses a similar approach to the now-famous algorithm for Extreme Audio Streching called "Paulstretch" [6]. This consists of taking many FFT windows with a small hop-size, randomizing the phase, performing an IFFT to resynthesize the sound, and then overlap-adding the resulting windows at a standard hop-size of half the FFT length. The key element of this from the point of view of the work presented here is the randomization of phase.

In this work, a "reverberated average" is created by taking overlapping FFTs throughout an input audio file and averaging the resulting amplitudes for each bin. Large FFT windows of at least 8192 samples are used for good spectral resolution, and Hann windows are applied. The resulting "spectral average" is resynthesized by performing IFFTs on the averaged spectrum, randomizing the phase of each bin as in the Paulstretch algorithm, and overlap-adding the resulting windowed waveforms. The result is a sonic texture that represents a kind of spectral fingerprint of the input recording. Although static in terms of the magnitude spectrum, the output shimmers as a result of the randomized phase.

III. Elaboration

The result of the above process can be enhanced by considering questions of brightness and contrast. By default, quiet parts of the music will have little effect on the resulting average. Also, any low-level noise that is present for every FFT frame gets compounded. A solution to this is to raise the magnitude of each bin to a power according to whether it lies above or below a threshold percentage of the maximum magnitude for that frame (Fig. 1). Since bin magnitudes are in the range [0, 1], an exponent less than 1 is used to increase the bin magnitude for those bins above the threshold, and an exponent greater than 1 decreases the magnitude of bins below the threshold. Not only does this increase dynamic contrast within a given frame, but quiet frames are also more affected by the exponentiation than loud frames, thereby boosting their weight in the final average.

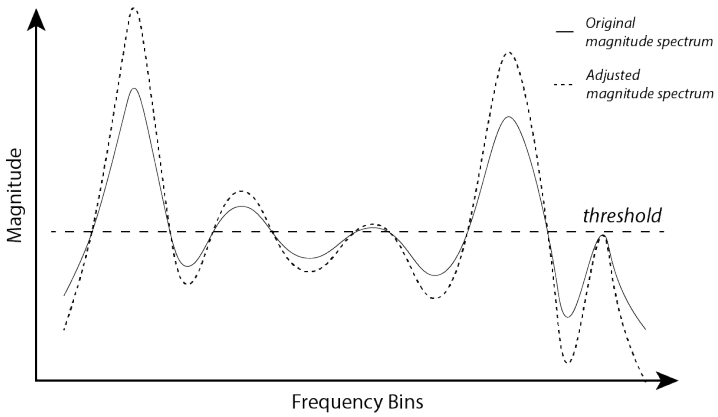

Fig. 1. Exponentiation of bin magnitudes to increase contrast.

Another striking enhancement occurs when the process is done in stereo, simply by performing the averaging in parallel on the two input channels. The result is much richer than that which results from a mono source file, as it adds an extra dimension of information. For instance, if the original input features hi-hat panned to the right, the resulting average will feature the frequency components of the hi-hat more strongly

on the right channel. Of course, phase-based spatial cues are thrown out in the process of averaging.

IV. Results

Two different kinds of installations have been created from this spectral averaging technique, one audiovisual, and one purely audio. The audiovisual installation consists of a looped, stereo "reverberated average" played through headphones or loudspeakers, while a visual component flickers in time with the original recording. As an hommage to the inspiration of Jim Campbell, the visual consists of a traditional illuminated average, sourced from a video of the performer, or from several portaits, in the case of a composer (Fig. 2). The flickering is colored according to the frequency content of the original recording as it evolves over time. In this way, after removing the role of time in the music, it is reintroduced as fragile apparition in the accompanying visual.

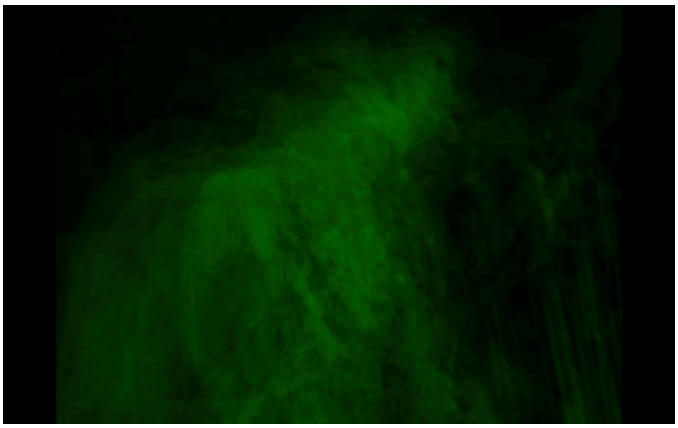

Fig. 2. Illuminated Average of Jimi Hendrix that flickers in response to the reverberated average of "All Along The Watchtower"

The second installation turns time into a spatial dimension. Each of the 32 minutes of a recording of the Rite of Spring (Seiji Ozawa, Chicago Symphony) are averaged in mono according to the above process. These averages are played back (in chronological order from left to right) on 32 highly directional speakers arranged in an arc. The result is that if one stands close to the speakers, one experiences the local spectrum of the given portion of the music, but if one stands back at the focal point of the arc, one experiences the "Reverberated Average" of the entire work. This installation format has some commonality with Tristan Perich's "Microtonal Wall" [7].

V. Conclusion

One major motivation for this work was to consider the relationship between art that acts through the dimension of time and art that does not, and to do so by converting temporal art to an atemporal form. As a composer, I have often thought about how listeners at a concert are trapped in a particular flow of time. If a piece is 20 minutes long, the listener is obliged to sit there for 20 minutes. The resulting experience can range from rapturous to torturous, depending on the sensitivities of the composer and the listener. By contrast, in an art museum, one can curate one's own experience of time, standing longer and going deeper if desired, or simply walking past if not.

Both installations presented here take pieces of music that would have originally been presented in the context of a concert and translate them to a form in which the listener can curate their own experience of time. In so doing, part of the hope is that it causes reflection on the role of time in the original work.

References

[1] http://www.jimcampbell.tv/portfolio/still_image_works/illuminated_averages/index.html

[2] http://www.medienkunstnetz.de/themes/overview_of_media_art/perception/7/

[3] https://www.sugimotohiroshi.com/new-page-7/

[4] http://www.medienkunstnetz.de/themes/overview_of_media_art/perception/8/

[5] http://www.læyf.com/

[6] http://hypermammut.sourceforge.net/paulstretch/

[7] https://vimeo.com/45225412