Audio Engineering Society

Convention Paper 9794

Presented at the 142nd Convention
2017 May 20–23 Berlin, Germany

Optimization of temporally diffuse impulses for decorrelation of multiple discrete loudspeakers

Jonathan B. Moore and Adam J. Hill

Department of Electronics, Computing and Mathematics, University of Derby, Derby DE22 3AW, United Kingdom

Correspondence should be addressed to Jonathan Moore (j.b.moore@derby.ac.uk).

ABSTRACT

Temporally diffuse impulses (TDIs) were originally developed for large arrays of distributed mode loudspeakers to achieve even radiation patterns. This initial investigation evaluates the performance of TDIs in terms of the reduction of low frequency spatial variance across an audience area when used with conventional loudspeakers. A novel variable decay windowing method is presented, allowing users control of TDI performance and perceptibility. System performance is modelled using an anechoic and an image source acoustic model. Results in the anechoic model show a mean spatial variance reduction of 42%, with a range of source material and using the optimal TDI generation methodology. Results in the image source model are more variable, suggesting that coherence of source reflections reduces static TDI effectiveness.

1 Introduction

One of the primary goals of any live sound system is to deliver a consistent audio experience to each member of the audience. The narrowing dispersion patterns of conventional pistonic loudspeakers at medium to high frequencies necessitates the use of multiple drivers to cover a wide audience area. Furthermore, multiple low frequency units are necessary to deliver the required sound pressure level of low frequency sound, without breakup or distortion. These requirements lead to a multiple loudspeaker setup, usually symmetrically spaced, at any medium to large scale live event delivering coherent audio signals. If left unaddressed this can potentially result in severe incongruities in listening experience across wide audience areas. This is due to constructive and destructive interference (i.e. comb filtering), which is by nature, position and frequency dependent.

Live sound engineers are familiar with the concept of "power alley". This is any position in the audience area equidistant from two halves of a symmetrical horizontal array. Where path length is equal, acoustic signals sum coherently. However, any listening position offset from the line of symmetry will be subject to destructive interference. The impact of this depends on the path length difference between coherent sources. The larger the difference, the lower the fundamental frequency of the destructive interference. At high frequencies, the interference results in a "phasey" response. At low frequencies, areas of the audience (often multiple meters wide) will experience significant inconsistencies in the desired sound pressure level.

This variability at low frequencies is not only an issue for the audience, but also leads to issues for the sound engineer who is tasked with providing an appropriate mix, and must do so whilst potentially receiving a

different acoustic response to that of much of the audience.

Decorrelating the discrete sources of a live sound reinforcement system will lead to a reduction in this variability, as sum zero displacement is only achieved when areas of peak rarefaction and compression overlap spatially. Applying random phase changes at each discrete source as a method of decorrelation reduces the likelihood of this occurring across large areas. If the phases are changed sufficiently, areas of constructive and destructive interference will still occur, but on a much smaller scale, and without a recognizable spatial distribution.

Methods of audio signal decorrelation have been previously investigated in [1,2,3], primarily with a focus on perceptual effects of coherent signals in multichannel systems and their impact on audio perception.

Kendall suggests that decorrelation of a monophonic signal with itself can be achieved with finite impulse response (FIR) filters, where the correlation measure of the two output signals will be similar to that of the two FIR filters' coefficients [1]. The filter coefficients are derived from a frequency domain specification, whereby the magnitude is set to unity and the phase is constructed from random number sequences.

Boueri and Kyriakakis address some of the timbral coloration issues that arise with this method, and describe a variation where random time shifts are applied to critical bands discretely and band-pass filters are used to maintain the original signal's energy [2].

Both methods achieve decorrelation, but there are associated impacts on timbre and transients. Molle et. al. discusses the use of decorrelation in multichannel surround systems to tackle the issue of sound objects' apparent width collapsing when outputted by multiple speakers simultaneously. Decorrelation processing is performed only on steady state sounds with a novel transient extraction method [3]. The decorrelation method described by [2] is attempted, but found to be too computationally expensive. Decorrelation is achieved by random phase shifts applied to the frequency bins of a fast Fourier transform (FFT). This previous work suggests that further testing needs to be done to assess the subjective effects of this processing [3].

Diffuse signal processing (DiSP) was described in [4] as a means of decorrelating outputted audio signals to standardize listening experiences across wide audience areas using synthetic, temporally diffuse impulses (TDIs).

In this work, the application of TDIs specifically to live sound reinforcement settings will be investigated, with a primary focus on the reduction of low frequency amplitude nulls caused by destructive interference. Comparisons are made between the performance of a two-speaker system, with and without DiSP. The system is tested in both an anechoic and image source acoustic model, with the metric of spatial variance used to quantify the variability of frequency response across multiple audience positions.

2 Diffuse Signal Processing (DiSP)

The method of signal decorrelation chosen for this work is DiSP, first described by [4]. Whilst [4] does assess TDI performance for several potential applications, this work seeks to assess the suitability of the processing specifically regarding live sound reinforcement systems, and low frequency spatial variance reduction across an audience area. The performance of TDIs when used with a variety of musical sources, steady state pink noise and an impulse source will therefore be evaluated. The objective of this testing is to establish the optimal methodology of TDI generation for the previously described application.

2.1 Perceptual Effects of DiSP

DiSP involves the use of multiple TDIs as all-pass FIR filters, applied prior to amplification. Each transducer in a system must have a unique TDI to achieve optimal decorrelation. To be considered perceptually transparent, the application of DiSP must be imperceptible for single sources, as well as when multiple sources are summed.

There are two primary concerns regarding the perceptual effects of TDIs. These are related to spectral notch width and the length of TDI window.

When viewed in the frequency domain prior to equalization, a TDI exhibits multiple random sharp notches in its magnitude response. Since each TDI is randomly generated, individual TDIs demonstrate notches of different magnitude and frequency. It is this variation that leads to diffuse behavior when multiple sound sources convolved with unique TDIs are summed.

The length of TDI window is directly linked to notch density, depth and width. To achieve the required density of notches at low frequencies, long windows are required. However, the decay rate of high frequency components need not be dictated by this, as they cause sufficient notch density in less time [4]. A frequency dependent exponential window is therefore used, where each frequency component has an appropriate exponential decay function applied to it. In this way, the perceptual effects of exceedingly long decay times can be avoided.

To avoid changes to source timbre, it is essential that the resulting spectral notches are narrow enough (have a high enough Q) to be imperceptible.

Previous research has investigated human perception of spectral notches [6, 7, 8]. The audibility of spectral notches decreases with increasing Q, where a Q factor of 30 or higher is inaudible at all frequencies [6]. A reduction in notch detection with decreasing bandwidth was also found in [7] and [8]. Notches below 125 Hz with a Q of 10 or more were found to be inaudible and a Q value of 30 was undetectable at any frequency [8].

Prior work was also done by one of the authors in this regard in a subjective listening test. Seventeen subjects (10 expert, 7 naive) assessed four audio clips including pink noise, pulsed noise, a piece of rock music and a piece of classical music. Each subject was asked to decrease the Q of a notch filter of unknown center frequency until the spectral notch was perceptible. Comparisons between the unfiltered and filtered audio were allowed and made with a bypass switch.

It was found that spectral notch detection decreased with increasing Q as in the previously mentioned works, but the threshold Q for detection varied throughout the audio spectrum. Table 1 summarizes the mean overall results of this subjective test.

f_0 (Hz)	75	125	500	2k	8k	16k
Q threshold	17	21	14	19	14	11

Table 1 Q detection thresholds for notch audibility

Regarding TDI generation, windowing must be sufficiently long to allow for spectral notch Q factors to approximate Table 1, but short enough to avoid persistent ringing or other audible artifacts.

2.2 Noise Generation Methods

When viewed in the time domain, TDIs consist of an initial impulse followed by a random phase tail. Three probability density functions are described by [4] for random phase generation. They are defined as triangular, uniform and binary random. Each method must be examined to evaluate performance in terms of spatial variance reduction across a wide audience area. The MATLAB code (adapted from [4]) for each PDF is given in Fig. 2.1, and described by Eqs. 1-3 respectively.

```
% create initial phase value matrix
% (nS = number of sources, m = FFT resolution)
phaseMat = zeros(nS, m/2);

% generate phase values based on PDF type
switch pdfType
    case 1  % triangular
        for i = 1 : nS
            phaseMat(i, 1 : m/2) = ...
                pi*(rand(1, m/2) + rand(1, m/2) - 1);
        end
    case 2  % uniform
        for i = 1 : nS
            phaseMat(i, 1 : m/2) = ...
                2*pi*(rand(1, m/2) - 0.5);
        end
    case 3  % binary random function
        for i = 1 : nS
            phaseMat(i, 1 : m/2) = ...
                pi*(round(rand(1, m/2) - 0.5));
        end
end
```

Fig. 2.1 MATLAB code for random phase response generation with each PDF variety

Each discrete source requires a unique vector of phase values. Therefore, *phaseMat* is a matrix where each row holds the phase values for an individual source.

The rows are filled using a *for* loop, depending on which PDF type is selected.

Equations 1-3 describe the different functions for triangular (Eq.1), uniform (Eq.2) and binary random (Eq.3) distribution, where *m* is the length of the TDI, and *r* is a vector of uniform random values.

$$p(n) = \pi \sum_{n=1}^{m/2} (r(n) + r(n) - 1) \quad (1)$$

$$p(n) = 2\pi \sum_{n=1}^{m/2} (r(n) - 0.5) \quad (2)$$

$$p(n) = \pi \sum_{n=1}^{m/2} |r(n) - 0.5| \quad (3)$$

Figures 2.2-2.4 show example distributions of phase values for each investigated PDF. A phase weighting constant has been applied to limit phase values to ±0.94π (approximately ±2.95) [4].

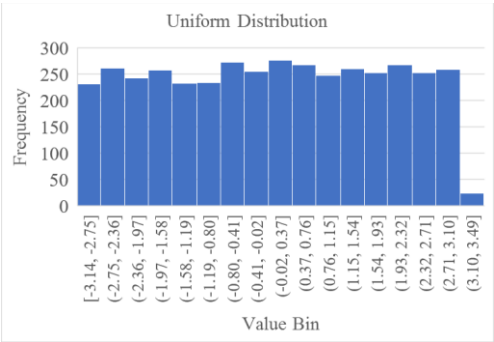

Fig. 2.2 Uniform PDF distribution

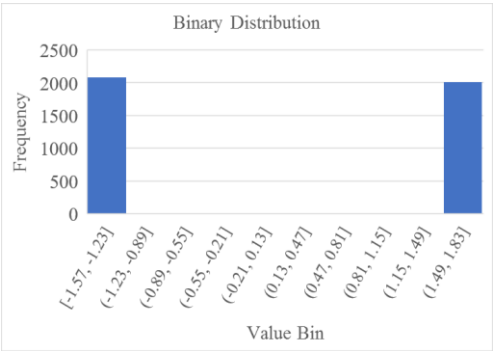

Fig. 2.3 Binary random PDF distribution

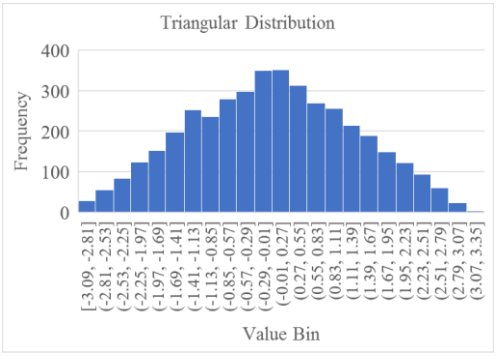

Fig. 2.4 Triangular PDF distribution

2.3 Windowing Methods

Five TDI windowing methods are examined here, defined as: all-pass, windowed noise, logarithmic scan, linear scan and variable decay constant.

The linear scan, logarithmic scan and variable decay constant methods all make use of frequency dependent exponential windowing, and are of most interest here. A full description of the all-pass and windowed noise methods is given in [4], and these methods have been included in the testing for completeness.

Each noise generation method described in Section 2.2 produces one vector of random phase values for each speaker in an array. The TDI is then constructed using the code in Figure 2.5 (adapted from [4]).

```
% Create initial TDI matrix (holding all TDIs)
% nS = number of sources, m = TDI length)
TDIs = zeros(nS, m);

% generate cosine waves with appropriate decay and phase for
% each frequency up to Nyquist and sum to form composite TDIs.
for z = 1 : nS
    for p = 1 : m/2
        cosDec = cos(phaseMat(z,p) + 2*p*pi*mm/m);
        TDIs(z, 1 : m) = TDIs(z, 1 : m) + cosDec/std(cosDec);
    end
end
```

Fig. 2.5 MATLAB code for TDI generation (without frequency-dependent windowing)

A matrix of *nS* rows (to match the number of sources) and *m* columns (the FFT resolution of the TDI) is initialized. For each iteration of a *for* loop, up to *m/2* (Nyquist frequency), a cosine wave is generated of

frequency corresponding to index *p*. The phase of a given cosine wave is set by the corresponding value held in the *phaseMat* matrix, for the given source (index *z* in Fig 2.5) and given *p* value. *mm* is a vector from 0 to $m-1$. The cosines are summed within each iteration of the loop up to Nyquist, including a division by the cosine's standard deviation to normalize.

To apply decay, each cosine is multiplied by an exponential decay function, which updates for every iteration of the for loop, thus enabling a frequency-dependent decay. The rate of change of the decay constant is defined either linearly or logarithmically between DC and Nyquist frequency.

The user defines the minimum and maximum rates of decay (at DC and Nyquist, respectively) in milliseconds (t_{max}, t_{min}). Equations 4 and 5 describe the conversion to upper and lower decay constant where *m* is TDI FFT resolution and *Fs* is sample rate.

$$x_{\max} = 1000\frac{m}{t_{\max}Fs} \tag{4}$$

$$x_{\min} = 1000\frac{m}{t_{\min}Fs} \tag{5}$$

Equation 6 describes how the slope of the rate of change is calculated for the logarithmic method, and Eq. 7 describes how it is calculated for the linear method.

$$k = \left(\frac{1}{x_{\max} - x_{\min}}\right)\log_{10}\left(\frac{m}{2}\right) \tag{6}$$

$$b = \left(\frac{x_{\max} - x_{\min}}{0.5m - 1}\right) \tag{7}$$

The initial rate of change at DC is calculated using Eq. 8 for the logarithmic method, and Eq. 9 for the linear method.

$$aa = e^{-k*x_{\min}} \tag{8}$$

$$a = x_{\min} - b \tag{9}$$

The frequency-dependent decay constant is applied during the cosine generation process, as previously described, for both the logarithmic and linear scan methods as shown in Fig. 2.6 (adapted from [4]).

```
% Logarithmic method
for p = 1 : m/2
    cosDec = cos(phi(p) + 2*p*pi*mm/m).*...
        exp(-(log(p/aa)/k)*mm/m);

    TDIs(1 : m) = TDIs(1 : m) +...
        cosDec/std(cosDec);
end

% Linear method
for p = 1 : m/2
    cosDec = cos(phi(p) + 2*p*pi*mm/m).*...
        exp(-(a + b*p)*mm/m);

    TDIs(1 : m) = TDIs(1 : m) +...
        cosDec/std(cosDec);
end
```

Fig. 2.6 MATLAB code for TDI generation (with frequency-dependent windowing)

Figures 2.7 and 2.8 show logarithmic and linear scan TDIs after normalization. The same decay constant limits of 100 ms to 2 ms (t_{max} & t_{min}) were used for both.

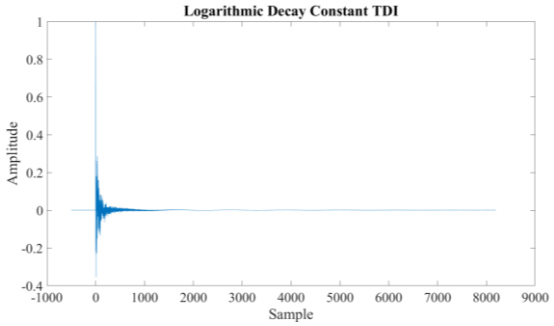

Fig. 2.7 Example TDI (logarithmic decay constant method)

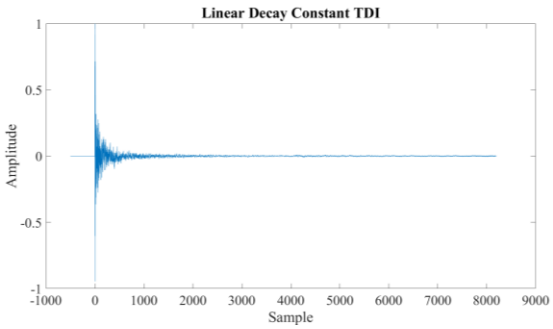

Fig. 2.8 Example TDI
(linear decay constant method)

Figures 2.9 and 2.10 show the magnitude response of the TDIs shown in Figs. 2.7 and 2.8. Here the random spectral notches created by the random phase generation can be seen. It is important to note the increased density of notches in the linear scan TDI at low frequencies. This is a result of longer low frequency decays with this method. The logarithmic method leads to low frequency components being attenuated more aggressively, thus having lower density of notches at low frequencies.

Regardless of the TDI generation method utilized, minimum phase equalization is applied to give a uniform (flat) magnitude response (Fig. 2.11), using the method described in [4].

In this work, a low frequency time constant limit of 100 ms, and high frequency time constant limit of 2 ms were chosen for the linear and logarithmic methods. These values were chosen to give the required notch density, with minimal perceptual effects.

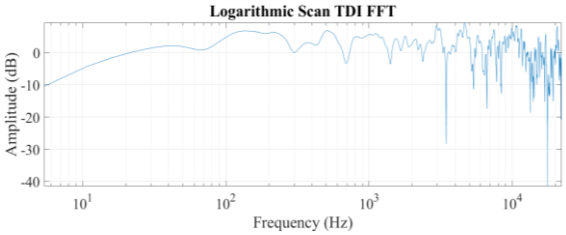

Fig. 2.9 Example TDI magnitude response
(logarithmic decay constant method)

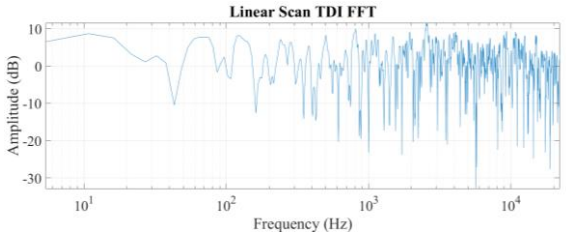

Fig. 2.10 Example TDI magnitude response
(linear decay constant method)

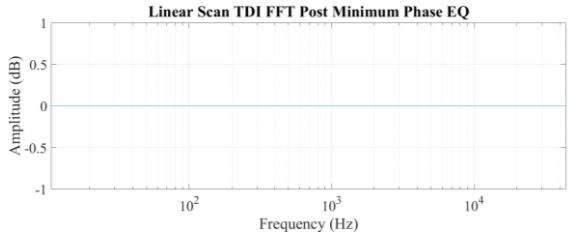

Fig. 2.11 Linear decay TDI post minimum phase EQ

2.4 Variable Decay Windowing Method

The proposed variable decay windowing method is a variation on the linear scan method. As previously described, both the linear and logarithmic scan methods only allow a user to define time constants for the lowest and highest frequencies.

In this method, the spectrum is divided into several bins. The user can define the time constant at each bin boundary, with intermediate values defined by Eq. 7. A GUI was developed to facilitate this. After TDI generation for *n* sources, a pre-equalization FFT of each individual TDI can be analyzed for average Q per frequency bin and overall TDI coherence.

Adjusting the time constants at each bin boundary allows for the frequency persistence to be increased or decreased. An increase in the persistence of components will increase depth and density of spectral notches, leading increased diffuse behavior in that part of the spectrum.

3 Simulation Method

To assess the TDI generation methods in terms of sound energy homogeneity over an audience area, two models were utilized: one anechoic and one utilizing an image source model (IS), similar to that

described in [9]. In both models, multiple point sources are placed at any position in a user-defined space.

In this investigation, a room size of 10 m x 10 m was chosen to simulate a small-sized venue. Appropriate materials were chosen to give a reverb time of 0.7 s in the IS model. The chosen materials are summarized in Table 2. It should be noted that the IS model at the time of testing did not incorporate frequency dependent absorption coefficients; the coefficients listed are average values.

For this testing, two point sources were utilized, as it was felt this would provide a "worst case scenario" due to only two degrees of freedom being available for DiSP. The sources were positioned 1 m away from the front wall, 1.6 m off the floor and 5 m apart from one another (centered along the median of the space). The audience area was divided into a 4 x 4 grid, positioned centrally within the room, thus providing 16 measurement positions.

Surface	**Ceiling**	**Floor**	**Back Wall**
Material	*Ceiling Tiles*	*Carpet*	*Plaster*
Coefficient	0.03	0.37	0.05
Surface	**Front Wall**	**Left Wall**	**Right Wall**
Material	*Heavy Fabric*	*Plaster*	*Plaster*
Coefficient	0.42	0.05	0.05

Table 2 Materials used in the image source model

3.1 Spatial Variance

Spatial variance has been used previously [10-13] to quantify sound pressure level variation across a listening area, and is calculated using Eq. 7.

$$SV = \frac{1}{N-1}\sum_{i=1}^{N}\left|A_i - \overline{A_i}\right|^2 \quad (7)$$

Where A_i is the measured SPL (dB) of a given frequency bin at each measurement position i, $\overline{A_i}$ is the mean SPL at all listening positions for that frequency bin. N represents the number of listening positions measured. A range of frequency bins can be assessed by repeating Eq. 7 for each bin in question. *SV* then becomes a vector of length corresponding to the number of frequency bins in the range. A single *SV* value for a range of bins can then be obtained by taking the mean of the *SV* vector. In this work, *SV* for frequency bins of the ranges 20-200 Hz, 200-4000 Hz and 4-15 kHz were calculated.

It should be noted that the grid spacing should ideally be no more than half the wavelength of the highest frequency of interest to avoid spatial aliasing [11]. However, in this case the result of interest is the frequency response before and after processing rather than an accurate mapping of the total sound field, therefore, fewer measurement points were used to aid computational efficiency.

It is expected that when strongly correlated sources are used in a multiple speaker array, spatial amplitude nulls will be created by summations of peak rarefaction and compression across the audience area. This should lead to a high spatial variance. The objective of this work is to reduce spatial variance in the measured SPL across the audience area by using DiSP to decorrelate otherwise coherent sources. This decorrelation should result in a more diffuse sound field where energy is distributed more evenly. If the spatial variance value obtained for a given piece of audio material, transduced by a given speaker array is reduced when DiSP is utilised, the objective is met. To quantify this, the percentage change in spatial variance with and without DiSP is analysed.

Note that in this work, spatial variance is only analysed for steady state responses. An assessment of the transient performance of DiSP is required in future investigations.

3.2 Spectral Smoothing

To bring the measurements in line with human perception each frequency response was subjected to $1/9^{th}$ octave smoothing before spatial variance was calculated. This level of smoothing was derived from the subjective test described in Section 2.1, by taking the mean Q value for notch detection obtained in those results (found to be 14.6).

3.3 Test material

Five audio tracks were selected for testing: classical

music, electronic dance music, rock music, pink noise and an impulse. Tables 3 and 4 show initial spatial variance for each track in the anechoic and IS models, with the loudspeaker array and measurement grid as described in Section 3. No processing has been applied at this point.

Sample	20-200 Hz	200-4000 Hz	4-15 kHz
Classical	23.79	10.44	1.80
EDM	38.29	5.82	1.05
Rock	25.08	7.42	1.97
Pink Noise	39.86	5.41	0.02
Impulse	52.14	6.44	0.01

Table 3 Unprocessed spatial variance (dB) for the anechoic model

Sample	20-200 Hz	200-4000 Hz	4-15 kHz
Classical	22.09	8.56	2.38
EDM	17.83	2.54	0.64
Rock	18.47	3.11	1.43
Pink Noise	15.61	2.18	0.21
Impulse	15.07	2.28	0.22

Table 4 Unprocessed spatial variance (dB) for the image source model

Both tables show the greatest spatial variance in the 20-200 Hz range, decreasing in the medium to high frequency ranges. As noted in Section 3.1, the grid spacing used meant that mid to high frequencies were not accurately mapped due to spatial aliasing. A finer grid size would be necessary for accurate analysis of mid to high frequencies, but the low frequency response is of central interest for this current work.

The reverberant nature and dimensions of the IS model lead to a critical distance, where direct and reverberant sound level are equal, of around 4 meters and a Schroeder frequency of 97 Hz. Given that several measurement positions included in the average spatial variance calculation were beyond 4 meters distance from the sound sources, and frequencies up to 200 Hz were also included in this average, it is expected that the measured unprocessed sound field will already be quite diffuse before any processing, and this is reflected in the lower initial spatial variance results for this model.

3.4 Test Methodology

The three previously described PDF functions (uniform, binary random and triangular) were used with the previously described TDI generation methods (logarithmic decay constant, linear decay constant, variable decay constant) to generate TDIs for testing. The all-pass and windowed noise methods described in [4] were also included for completeness, giving 11 TDI generation methodologies in total. Each methodology was tested in both acoustic models, with the goal of reducing sound pressure level spatial variance in the 20-200 Hz range.

4 Results

Table 5 shows the average change in spatial variance after processing due to each PDF method. This data spans all test configurations.

PDF Function	Triangular	Uniform	Binary Random
% change	-1.62	-16.47	-10.92

Table 5 Spatial variance % change by PDF method after DiSP. Negative results indicate average spatial variance has been reduced due to DiSP.

The triangular PDF function for generating random phase values was found to be ineffective. This is because this distribution results in a greater density of phase values towards the centre of the scale. Extreme values are required to give large phase shifts for properly uncorrelated TDIs. The following results analysis will therefore focus on methodologies including uniform and binary random PDF functions.

Figures 4.1 and 4.2 show results for each TDI generation methodology in the anechoic and IS models, respectively. All audio tracks were investigated, with the average taken for the results.

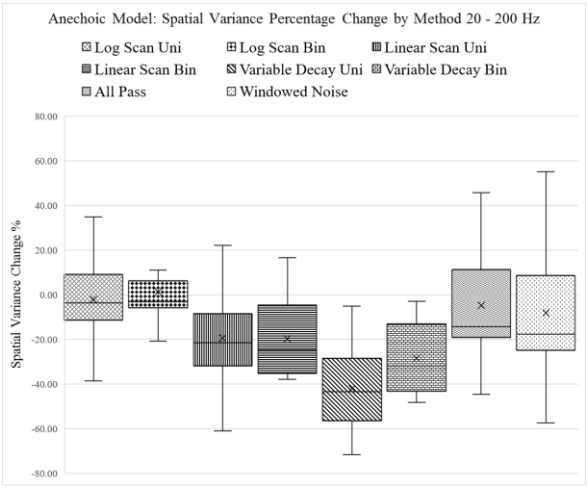

Fig 4.1 Percentage of spatial variance change by TDI generation method in the anechoic model

The anechoic results show that the most effective methodology for TDI generation was the variable decay constant method with the uniform PDF function. This method achieved a maximum spatial variance reduction of 71.6% with the impulse sound source. Figure 4.2 shows generally poor DiSP performance, and in some cases an increase in the calculated spatial variance. This is a result of the already diffuse sound field of the IS model. The value of DiSP in such environments is questionable, and the processing is of more value where the properties of the acoustic space do not already facilitate a diffuse sound field.

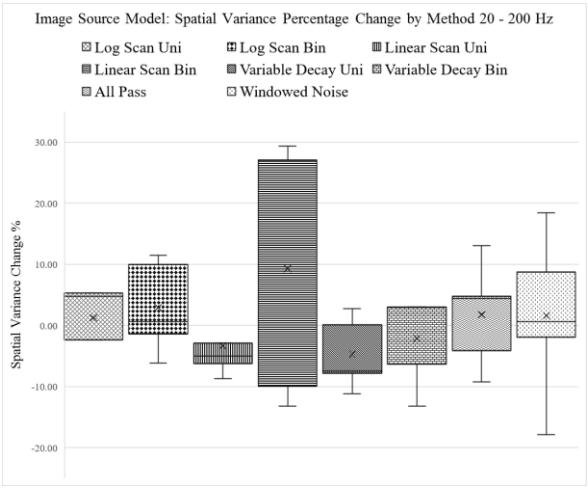

Fig 4.2 Percentage of spatial variance change by TDI generation method in the image source model

4.1 Multiple Speaker Systems

A test was conducted to examine the effect of adding further sources to the system. Additional sources were placed in between the initial two sources, with even spacing maintained. The array was positioned 1 m from the front wall, equidistantly between the two side walls and 1.6 m off the floor. Table 8 summarizes the results from the anechoic model.

No. Speakers	SV (no DiSP)	SV (w/DiSP)	% Change
2	39.86	29.01	-27.22
4	19.68	16.61	-15.59
8	16.21	16.45	1.47
16	16.57	16.30	-1.65

Table 8 Mean SV due to number of speakers using variable windowing method in the anechoic model.

As more speakers are added to the system, the measured unprocessed spatial variance decreases, leading to diminishing returns from DiSP processing. This is because sources placed within one-half wavelength at a given frequency will couple and act as a single source, thus reducing spatial variance in the anechoic model.

5 Conclusion

The results obtained by the anechoic model indicate that the most effective TDI generation methodology for the reduction of spatial variance below 200 Hz, is the variable decay constant method when used with the uniform PDF function. The results indicate that using this method consistently reduces low frequency spatial variance by between 20-40%, regardless of source material.

This method allows for the frequency-dependent filter persistence to be accurately defined. When low frequency diffuse behavior is desired, the persistence of low frequency components in the TDI can be increased whilst holding medium/high frequency components (where the perceptual impact is most noticeable) to minimal persistence in the filter.

Whilst the logarithmic and linear decay constant methods allow some flexibility in defining the persistence of elements by frequency, in these methods the decay constants are only set at the highest and lowest frequencies, meaning that mid-range frequencies be subject to too long decay times when trying to optimize TDI generation for low frequency diffuse behavior.

It should be noted that the large variances in TDI performance, as shown by Figures 4.1 and 4.2 were due to different behavior of the TDIs with different source material. This is likely due to source material characteristics, indicating that source phase may need to be considered in the random phase generation process to increase consistency across different source material. Future work will investigate this.

It is important to note that reflections from surfaces will maintain coherence, even when static DiSP is applied. This leads to the logical conclusion that to achieve optimum performance, dynamic DiSP must be investigated, whereby the TDIs for each individual source should vary with time to avoid correlated reflections within a closed acoustic space. This is the subject of current investigations by the authors.

References

[1] Kendall, G. “The Decorrelation of Audio Signals and its Impact on Spatial Imagery,” *Computer Music Journal,* vol. 19, pp. 71-87, 1995.

[2] Boueri, M.; C. Kyriakakis. “Audio Signal Decorrelation Based on a Critical Band Approach,” presented at the *117th Convention of the Audio Engineering Society,* Convention Paper 6291, 2004.

[3] Molle, B.D.; J. Pinkl; M. Blewett. “Decorrelated Audio Imaging in Radial Virtual Reality Environments,” presented at the *139th Convention of the Audio Engineering Society,* Convention e-Brief 208, 2015.

[4] Hawksford, M.O.J.; N. Harris. "Diffuse signal processing and acoustic source characterization for applications in synthetic loudspeaker arrays," presented at the *112th Convention of the Audio Engineering Society* (2002, May.), convention paper 5612.

[5] Hawksford, M.O.J. “Digital Signal Processing Tools for Loudspeaker Evaluation and Discrete-Time Crossover Design”, *Journal of Audio Engineering Society.* Volume 45 Issue 1/2 pp. 37-62, 1997.

[6] Bucklein, R. “The Audibility of Frequency Response Irregularities,” *J. Audio Eng. Soc.* vol. 29, no.3*,* pp. 126 – 131, 1981.

[7] Moore, B.C.; S. R. Oldfield; G.J. Dooley. “Detection and discrimination of spectral peaks and notches at 1 and 8 kHz,” *The Journal of the Acoustical Society of America,* 85.2, pp. 820-836, 1989.

[8] Olive, S.E.; P. L. Shuck; J. G. Ryan; S. L. Sally; M. E. Bonneville. “The Detection Thresholds of Resonances at Low Frequencies,” *J. Audio Eng. Soc.* vol. 45, no. 3, pp – 116-128, 1997.

[9] Allen, J.B.; D. A. Berkley. “Image Method for Efficiently Simulating Small-Room Acoustics,” *The Journal of the Acoustical Society of America,* vol. 65, no. 4, 1979.

[10] Hill, A.J.; Hawksford, M.O.J.; Newell, P. “Enhanced Wide-Area Low-Frequency Sound Reproduction in Cinemas: Effective and Practical Alternatives to Current Calibration Strategies.” *Journal of the AES*, vol. 64, no. 5, pp. 280-298. May, 2016.

[11] Hill, A.J. “Analysis, modeling and wide-area spatiotemporal control of low-frequency sound reproduction,” *Ph. D Thesis, University of Essex,* (2012).

[12] Celestinos, A.; Nielsen, S. B. “Optimizing placement and equalization of multiple low frequency loudspeakers in rooms,” presented at the *119th Convention of the Audio Engineering Society,* Convention Paper 6545, 2005

[13] Welti, T.; Devantier, A. “Low-frequency optimization using multiple sub-woofers,” *Journal of the Audio Engineering Society,* vol. 54, no. 5, pp. 347-364. May, 2006.