Low-Order State-Space Models of Head-Related Transfer Function Arrays

Norman H Adams

March 16, 2007

Abstract

Low-order state-space models of collections of *head-related transfer functions* (HRTFs) are described and found to yield a substantial computational savings relative to conventional filter arrays. Several recent binaural displays model multiple HRTFs simultaneously using an array of independent FIR or IIR filters. As HRTFs are locally redundant it is possible to model collections of HRTFs with a single state-space system. In the present work, high-order state-space systems are built from measured HRTFs, and two order reduction methods, *balanced model truncation* (BMT) and *Hankel optimal approximation* (HOA), are used to design efficient approximants. Both methods are based on the Hankel operator, which is described and related to the convolution operator. Complete algorithms for both reduction methods are given. The performance of the two methods is compared and HOA is found to exhibit favorable properties for HRTF approximation, although overall performance is similar for both methods. The performance of the state-space systems is compared to that of a truncated minimum-phase FIR array of equal net cost for a variety of binaural display conditions. The state-space methods are found to outperform the FIR array in many conditions.

1 Introduction

Emerging applications in *binaural environment modeling* (BEM) present a unique opportunity for the design of efficient state-space systems. For *single-input single-output* (SISO) filter approximation, state-space systems do not

offer a computational savings over FIR or IIR systems. However, for a filter that contains either multiple inputs or outputs, a single state-space system may outperform an array of independent FIR or IIR filters of equal net cost. Contemporary approaches to BEM often employ *multiple-input multiple-output* (MIMO) systems, but are usually implemented as an array of independent filters. The transfer functions of each input-output pair of such systems exhibit similar features, and, in particular, exhibit similar pole locations. As such, a single low-cost MIMO state-space system may be suitable for BEM.

In the present work we demonstrate that a state-space system can offer a substantial computational savings for binaural environment modeling relative to conventional filter arrays. We construct high-order state-space systems from a collection of measured *head-related transfer functions* (HRTFs) and explore two methods of model reduction based on the *Hankel operator*. The Hankel-operator is not traditionally used to analyze audio filters; we relate this operator to the more common convolution operator. The two model reduction methods are closely related, although one is *ad hoc* and relatively simple to implement whereas the other is optimal in the Hankel-norm sense but is somewhat more complicated to implement. The two methods have been compared extensively for SISO systems, but relatively few comparisons have been made for the MIMO case. We are not aware of any studies that compare the two methods for systems as large as those required for binaural environment modeling. Both methods are adapted to HRTF modeling and complete algorithms are given in the appendix.

The remainder of this section gives background on BEM and reports two previous studies that employed state-space systems for this application. Section 2 develops the Hankel operator and describes the two model reduction techniques. In section 3 the results of a numerical experiment are reported in which the two techniques are directly compared to FIR arrays of equal computational cost.

1.1 Binaural Environment Modeling

An acoustic scene that consists of a single stationary sound source in the far-field of a listener in an anechoic environment can be modeled using a single pair of HRTFs [1]. The HRTF represents the acoustic filtering of a plane wave enroute to a listener's two ears due to the head, pinna and torso of the listener, and hence is unique to the listener. Such transfer functions can

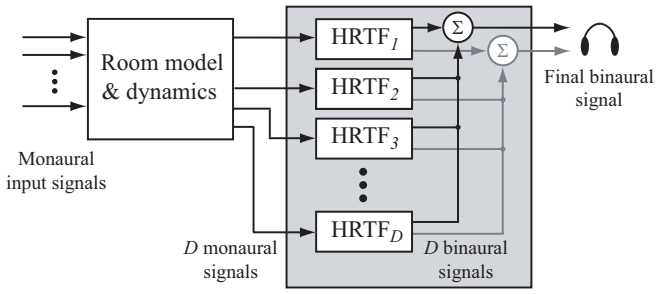

Figure 1: A conventional implementation of a multi-direction binaural display.

be implemented using appropriately measured *head related impulse responses* (HRIRs), which, empirically, require approximately 200 FIR filter coefficients at a sampling rate of 44.1 kHz. A binaural display that operates in this manner, by filtering a monaural source with two $N = 200$ FIR filters and supplying the two output signals to a pair of headphones, requires about 18 million multiplication and addition operations per second. Because this is a substantial computational load, considerable effort has been made to find low-order approximations to measured HRIRs [2].

However, it is well known that binaural displays, as described above, are perceptually unsatisfying [3]. Several recent studies propose methods for synthesizing 'realistic' auditory scenes. Binaural displays have been designed that account for reflective environments [4–9], source and listener motion [4–7,10–12], and spatially-extended sources [13]. These examples share a similar framework: monaural source signals are filtered with multiple HRTF pairs instead of a single HRTF pair. That is, a monaural source is *auralized* at D directions simultaneously, and then the D binaural signals are combined so as to model the desired auditory scene. This framework compounds the computational burden however, as the overall computational cost scales linearly with D if individual filters are employed for each HRTF. A common architecture for such binaural displays is shown in Figure 1. In the present work we consider a D-input 2-output state-space system that models the HRTFs shown in the gray box in Figure 1.

1.2 State-Space Modeling of HRTFs

HRTFs measured for different directions, and for different listeners, are locally redundant. Even if the *common transfer function* (CTF) is removed from a collection of HRTFs, leaving only the *directional transfer functions* (DTFs), the transfer functions still exhibit similar spectral features, especially for nearby directions [1]. Numerous studies have found that collections of HRTFs can be reasonably represented in low dimensional spaces. In [14] it is shown that most of the HRTF variance can be accounted for with the first five principal components of a measured dataset. However, such representations do not necessarily yield low cost filters for individual HRTFs. A system that models HRTFs at many directions simultaneously may be able to utilize the redundancy of HRTF datasets to reduce the net cost of the system. Indeed, it has been shown that HRTFs can be accurately approximated using IIR filters with common poles [15]. This implies that a collection of HRTFs can be reasonably approximated using a single MIMO state-space system, as the rational transfer functions between each input and output of the state-space system share a common denominator polynomial [16].

At first glance, modeling HRTF filters as state-space systems may not appear computationally efficient. Any order N state-space system can always be converted to an equivalent array of order N IIR filters [16]. For SISO systems an IIR implementation is necessarily lower cost than a equivalent state-space system. Nonetheless, state-space techniques can be used to design low-order IIR filters from high-order FIR filters [17]. Several previous studies use one such technique, *balanced model truncation* (BMT), to design IIR filters from measured HRIRs [11,13,18]. These studies do not consider MIMO state-space systems however, hence the computational cost still scales linearly with the number of directions D. Furthermore, converting filters from state-space form to transfer function form may yield IIR filters that are sensitive to coefficient quantization errors [16,19]. In the present work we avoid this problem by leaving the reduced-order system in state-space form and considering its net computational cost.

Two recent studies propose state-space systems that model HRTFs at multiple directions simultaneously. In [19] MISO systems are designed that model multiple HRTFs for each ear. HRTF redundancy is not fully exploited in this work however, as separate state-space systems are designed for each HRTF individually, and then combined into one large system. In contrast, [20] considers a MIMO state-space design that directly models mul-

tiple HRTFs in the horizontal plane. Both studies employed BMT to design low-order systems [21]. It was shown that, for sufficiently large system order, listeners demonstrate similar localization performance with state-space systems as with conventional high-order FIR systems [20]. However, neither study considered in detail the computational advantages of state-space implementations. One goal of the present work is to demonstrate that a substantial computational savings can be achieved for binaural environment modeling using reduced-order state-space systems.

2 Methods

Consider a stable, causal, discrete-time MIMO state-space system[1]

$$\begin{aligned} \mathbf{x}[n+1] &= \mathbf{A}\mathbf{x}[n] + \mathbf{B}\mathbf{u}[n] \\ \mathbf{y}[n] &= \mathbf{C}\mathbf{x}[n] \end{aligned} \tag{1}$$

where $\mathbf{x}[n]$ is the state vector of size N_0, $\mathbf{u}[n]$ is the input vector of size M, and $\mathbf{y}[n]$ is the output vector of size P. To simplify notation, let $\mathbf{\Sigma} = \big(\mathbf{A}, \mathbf{B}, \mathbf{C}\big)$ represent the state-space system. The matrix impulse response of $\mathbf{\Sigma}$ is

$$\begin{aligned} \mathbf{h}[n] &= \begin{bmatrix} h_{11}[n] & \dots & h_{1M}[n] \\ \vdots & \ddots & \vdots \\ h_{P1}[n] & \dots & h_{PM}[n] \end{bmatrix} \\ &= \begin{cases} \mathbf{C}\mathbf{A}^{n-1}\mathbf{B} & n > 0 \\ \mathbf{0} & n \leq 0 \end{cases} \end{aligned} \tag{2}$$

For a D-input 2-output systems that models the HRTFs shown in Figure 1 the block impulse response is

$$\mathbf{h}[n] = \begin{bmatrix} h_1^L[n] & h_2^L[n] & \dots & h_D^L[n] \\ h_1^R[n] & h_2^R[n] & \dots & h_D^R[n] \end{bmatrix} \tag{3}$$

[1]For convenience the systems considered here have no feed-through term (the $\mathbf{D}\mathbf{u}[n]$ term), similar to [19, 20]. The Hankel operator, described below, is not influenced by the $\mathbf{D}$ matrix, hence the choice of $\mathbf{D}$ is somewhat arbitrary for this class of model reduction methods. In the present work we simply set $\mathbf{D} = \mathbf{0}$. The interested reader is referred to [22] for a detailed discussion of this term.

where $h_d^L[n]$ and $h_d^R[n]$ are the HRIRs for the left and right ears for direction d. It is straightforward to design a state-space system $\mathbf{\Sigma}$ that implements the collection of $2D$ HRIRs exactly[2]. However, such a state-space system is high order and computationally prohibitive. As such, we explore model reduction techniques to design low-cost approximants $\widehat{\mathbf{\Sigma}} = \left(\widehat{\mathbf{A}}, \widehat{\mathbf{B}}, \widehat{\mathbf{C}}\right)$ with order $N < N_0$. The two reduction methods that we explore are based on the Hankel operator, which is described next.

2.1 The Hankel Operator

Relative to SISO systems, there are few methods for reducing the order of MIMO systems such that the resulting low-order approximation is optimal in some sense. One metric for which optimal solutions can be found is the *Hankel norm*. However, interpreting this metric for audio applications, such as binaural displays, is somewhat subtle. The Hankel norm is a lower bound to the *2-induced*, or $\mathcal{L}^\infty$, norm, which has a clear spectral interpretation. Furthermore, it is often observed in practice that Hankel-optimal methods often yield solutions for which the Hankel error is a fortuitously tight lower bound on the $\mathcal{L}^\infty$ error. As such, we explore Hankel-based model reduction techniques in the next section. To clarify the interpretation of these techniques, the $\mathcal{L}^\infty$ and Hankel norms are described below along with the corresponding convolution and Hankel operators. The development below is general for all MIMO systems, not only D-input 2-output systems. The interested reader is referred to [23] for an excellent review.

Consider a matrix $\mathbf{X} \in \mathbb{R}^{p \times m}$. The 2-induced norm of this matrix is defined as

$$\|\mathbf{X}\|_{2\text{-ind}} \triangleq \sup_{\mathbf{u} \neq \mathbf{0}} \frac{\|\mathbf{X}\mathbf{u}\|_2}{\|\mathbf{u}\|_2} \tag{4}$$

where $\|\cdot\|_2$ is the standard Euclidean or 2-norm of a vector. If the matrix is viewed as a linear map, $\mathbf{X} : \mathbb{R}^m \to \mathbb{R}^p$, then the 2-induced norm is the maximum gain of the map. The *singular value decomposition* (SVD) of $\mathbf{X}$ is

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^* \tag{5}$$

where $\mathbf{U}$ and $\mathbf{V}$ are square unitary matrices. $\mathbf{S}$ is a diagonal rectangular matrix with the *singular values* of $\mathbf{X}$ arranged along the diagonal, ($\sigma_1 \geq \sigma_2 \geq$

[2]For example, the controller canonical form, as described in the appendix, can be used

$\cdots \sigma_{N_0})$, where $N_0 \leq \min(p, m)$ is the rank of $\mathbf{X}$. The SVD can be viewed as a *dyadic decomposition* of $\mathbf{X}$ into a sum of rank one matrices,

$$\mathbf{X} = \sum_{k=1}^{N_0} \sigma_k u_k v_k^* \tag{6}$$

where u_k and v_k are the k^{th} columns of $\mathbf{U}$ and $\mathbf{V}$. It is simple to show that $\|\mathbf{X}\|_{2\text{·ind}} = \sigma_1$. Furthermore, for a matrix $\widehat{\mathbf{X}}$ with size $p \times m$ and rank $N < N_0$, then

$$\|\mathbf{X} - \widehat{\mathbf{X}}\|_{2\text{·ind}} = \sigma_1\big(\mathbf{X} - \widehat{\mathbf{X}}\big) \geq \sigma_{N+1}\big(\mathbf{X}\big) \tag{7}$$

where $\sigma_1\big(\mathbf{X} - \widehat{\mathbf{X}}\big)$ is the largest singular value of $\mathbf{X} - \widehat{\mathbf{X}}$ and $\sigma_{N+1}\big(\mathbf{X}\big)$ is the $(N+1)^{\text{th}}$ largest singular value of $\mathbf{X}$. This inequality is known as the Schmidt-Mirsky theorem and gives a lower bound on how well any low-rank matrix $\widehat{\mathbf{X}}$ can approximate $\mathbf{X}$ in the 2-induced norm.

Consider the state-space system $\mathbf{\Sigma}$ with M inputs and P outputs given in (1). Associate the following *convolution operator* $\mathcal{L} : \mathbf{u} \mapsto \mathbf{y}$ with $\mathbf{\Sigma}$:

$$\mathbf{y}[n] = \sum_{k=-\infty}^{n-1} \mathbf{h}[n-k]\mathbf{u}[k], \quad n \in \mathbb{Z} \tag{8}$$

where $\mathbf{h}[n]$ is the matrix impulse response defined in (2). The convolution operator can be expressed in matrix form as

$$\begin{bmatrix} \vdots \\ \mathbf{y}[-1] \\ \hline \mathbf{y}[0] \\ \mathbf{y}[1] \\ \vdots \end{bmatrix} = \underbrace{\left[\begin{array}{ccc|ccc} \ddots & & & \vdots & & \iddots \\ & \mathbf{0} & & \mathbf{0} & \mathbf{h}[1] & \\ \hline \cdots & \mathbf{0} & & \mathbf{h}[1] & \mathbf{h}[2] & \cdots \\ & \mathbf{h}[1] & & \mathbf{h}[2] & \mathbf{h}[3] & \\ \iddots & & & \vdots & & \ddots \end{array}\right]}_{\mathcal{L}} \underbrace{\begin{bmatrix} \vdots \\ \mathbf{u}[0] \\ \hline \mathbf{u}[-1] \\ \mathbf{u}[-2] \\ \vdots \end{bmatrix}}_{\mathcal{U}}$$

The 2-induced norm of the system $\mathbf{\Sigma}$ is defined as

$$\|\mathbf{\Sigma}\|_{2\text{·ind}} \triangleq \|\mathcal{L}\|_{2\text{·ind}} = \sup_{\mathcal{U} \neq \mathbf{0}} \frac{\|\mathcal{L}\mathcal{U}\|_2}{\|\mathcal{U}\|_2} \tag{9}$$

The discrete-time Fourier transform can be applied to each element of the impulse response $\mathbf{h}[n]$, yielding the $P \times M$ matrix frequency response

$\mathbf{H}(\omega)$. Due to the equivalence between time and frequency domains, it is straightforward to show that

$$\|\boldsymbol{\Sigma}\|_{2\text{-ind}} = \sup_{\omega} \sigma_1\big(\mathbf{H}(\omega)\big) \triangleq \|\boldsymbol{\Sigma}\|_{\mathcal{L}} \tag{10}$$

where $\sigma_1\big(\mathbf{H}(\omega)\big)$ is the largest singular value of $\mathbf{H}(\omega)$. Hence the 2-induced norm is also known as the $\mathcal{L}^\infty$ norm. For SISO systems, this norm is equal to the maximum spectral magnitude.

Optimal causal approximations of MIMO systems in the $\mathcal{L}^\infty$ norm are not currently known[3]. However, if the domain and range of the convolution operator are restricted, then optimal solutions are known. For this reason, we consider the *Hankel operator* of $\boldsymbol{\Sigma}$, $\mathcal{H} : \mathbf{u}_- \mapsto \mathbf{y}_+$, which maps past inputs to future outputs:

$$\mathbf{y}[n] = \sum_{k=-\infty}^{-1} \mathbf{h}[n-k]\mathbf{u}[k], \quad n \in \mathbb{Z}_+ \tag{11}$$

The Hankel operator can be expressed in matrix form as

$$\begin{bmatrix} \mathbf{y}[0] \\ \mathbf{y}[1] \\ \mathbf{y}[2] \\ \vdots \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{h}[1] & \mathbf{h}[2] & \mathbf{h}[3] & \\ \mathbf{h}[2] & \mathbf{h}[3] & \mathbf{h}[4] & \dots \\ \mathbf{h}[3] & \mathbf{h}[4] & \mathbf{h}[5] & \\ & \vdots & & \ddots \end{bmatrix}}_{\mathcal{H}} \begin{bmatrix} \mathbf{u}[-1] \\ \mathbf{u}[-2] \\ \mathbf{u}[-3] \\ \vdots \end{bmatrix}$$

The matrix $\mathcal{H}$, while potentially infinite in size, has rank not greater than N_0. The rank is exactly N_0 iff the system $\boldsymbol{\Sigma}$ is minimal. Let $(\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_{N_0})$ be the singular values of $\mathcal{H}$. The *Hankel norm* of the system $\boldsymbol{\Sigma}$ is defined as the maximum singular value of $\mathcal{H}$

$$\|\boldsymbol{\Sigma}\|_{\mathcal{H}} \triangleq \sigma_1 = \|\mathcal{H}\|_{2\text{-ind}} \tag{12}$$

It can be shown that the Hankel norm lower bounds the $\mathcal{L}^\infty$ norm. It can also be shown that twice the sum of the Hankel singular values upper bounds the $\mathcal{L}^\infty$ norm

$$\sigma_1 \leq \|\boldsymbol{\Sigma}\|_{\mathcal{L}} \leq 2(\sigma_1 + \cdots + \sigma_{N_0}) \tag{13}$$

[3]In some instances, such as one-step model reduction, $\mathcal{L}^\infty$ optimal approximations are known [23].

Let $\widehat{\mathbf{\Sigma}}$ be another state-space system with M inputs, P outputs, and order $N < N_0$. Then

$$\sigma_{N+1} \leq \|\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}\|_{\mathcal{H}} \leq \|\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}\|_{\mathcal{L}} \tag{14}$$

where the first inequality follows directly from (7) and (12) for finite Hankel operators[4], and the second inequality follows from (13). It can also be shown that the $\mathcal{L}^\infty$ error is upper bounded by both the 1-norm applied to $\mathbf{h}[n] - \widehat{\mathbf{h}}[n]$ as well as twice the sum of the Hankel singular values of the error system $\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}$ [24, 25], although we have found in practice that these bounds are typically loose.

Remarkably, it can be shown that there exists a low-order system that achieves the lower bound on the Hankel error in (14). This result was proven for Hankel operators by Adamjan, Arov and Krein and is known as the AAK theorem [26]. Later, Glover [22] extended this result to state-space systems and developed a method for computing all optimal $\widehat{\mathbf{\Sigma}}$. In the next section we describe both Glover's method, as well as a simpler suboptimal method.

2.2 Model Reduction

Two order reduction techniques are considered, *balanced model truncation* (BMT) and *Hankel-norm optimal approximation* (HOA). In the SISO case, the two methods have been directly compared for IIR filter design [27]. Both methods are briefly reviewed below, and detailed algorithms are described in the appendix.

BMT is a state-space model reduction technique that operates by discarding all but the N largest singular values of a *balanced* system [21]. Balancing a state-space system transforms the system such that the controllability and observability Gramians are identical diagonal matrices, with the singular values of $\mathcal{H}$ contained along the diagonal [28]. In this case, any state $\mathbf{x}_0$ that results in a 'small' amount of energy output (with the input set to zero) also requires a 'large' amount of energy at the input to move the system from zero to state $\mathbf{x}_0$. Such states contribute little to the input-output behavior of the system, and hence can be truncated without greatly affecting the transfer function matrix $\mathbf{H}(\omega)$.

For our application it is not necessary to explicitly balance a state-space system to perform BMT, as the original system is specified by a collections of measured HRIRs. In this case, the Hankel matrix $\mathcal{H}$ is constructed directly

[4]The inequality holds for infinite Hankel operators as well.

from the measured HRIRs. The SVD of $\mathcal{H}$ is then computed, and the $N < N_0$ largest singular values, along with the corresponding singular vectors, are used to construct a balanced order N state-space system.

While BMT is convenient, it is not optimal in any specific sense. The HOA method is also based on the theory of balanced systems. However, rather than simply truncate the $N_0 - N$ least significant states, HOA operates in a more sophisticated manner. The $N + 1^{\text{th}}$ state is removed directly by means of an *all pass dilation*, and the remaining $N_0 - N - 1$ states are transformed so as to be both antistable and anticausal, in which case they no longer influence the Hankel operator of the system. The order N stable subsystem is then extracted, yielding a Hankel-optimal approximant $\widehat{\mathbf{\Sigma}}$ [22]

$$\|\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}\|_{\mathcal{H}} = \sigma_{N+1}(\mathcal{H}) \tag{15}$$

The following bounds apply to order $N < N_0$ systems reduced using either BMT or HOA[5]

$$\|\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}\,\|_{\mathcal{H}} \leq \|\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}\,\|_{\mathcal{L}} \leq 2(\sigma_{N+1} + \cdots + \sigma_{N_0}) \tag{16}$$

In practice, the upper bound on the $\mathcal{L}^\infty$ error is often loose, whereas the Hankel error is often a relatively tight lower bound on the $\mathcal{L}^\infty$ error.

Both the Hankel error and the $\mathcal{L}^\infty$ error are reported in Section 3 for several state-space and FIR systems that approximate HRTF filter arrays. As we are primarily interested in state-space systems that are more efficient than FIR systems, we must consider a definition of computational cost that is consistent for both system types.

2.3 Computational Cost

In the next section, state-space systems are compared to FIR filter arrays of *equal computational cost*. A measure of computational cost that is consistent for both system types must be defined. We define the cost C as the number of multiplication operations required per sample period, or equivalently, the total number of non-zero coefficients in the system. This measure of computational cost is common in filter design applications when comparing FIR and IIR filters [2, 13, 29]. An FIR filter array of order N with M inputs and

[5] If the state-space system is given a feed-through path, then the upper bound on the $\mathcal{L}^\infty$ error is halved for the HOA method [22].

P outputs requires $C = PM(N+1)$ multiplies per sample period. For example, consider a $N = 255$ FIR array that models HRIRs at $D = 10$ directions. This system has $M = 10$ inputs and $P = 2$ outputs, and has cost $C = 5120$.

For state-space systems, the computational cost depends on the choice of system realization, as there are many state-space systems with the same input-output behavior. In general, a state-space system of order N with M inputs and P outputs, and no feed-through path, requires $C = N^2+(P+M)N$ multiplies per time step. However, it is possible to apply a similarity transform to the system matrices ($\mathbf{A}$,$\mathbf{B}$,$\mathbf{C}$) such that the $\mathbf{A}$ matrix becomes more sparse. For example, a *modal decomposition* can be used to diagonalize the $\mathbf{A}$ matrix [30], reducing the system cost to $C = (P+M+1)N$. However, in this case the system matrices are complex. Alternatively, the $\mathbf{A}$ can be transformed to *Jordan canonical form*, a real and nearly diagonal form. Unfortunately, the Jordan form is difficult to compute in general, and notoriously prone to quantization error.

A more practical alternative is to employ a *Schur decomposition* to triangularize the $\mathbf{A}$ matrix [30]. Because we seek system matrices that are strictly real, the new $\mathbf{A}$ matrix is only quasitriangular, with either 1×1 or 2×2 blocks down the main diagonal. An algorithm for computing the real Schur decomposition can be found in [31], and is implemented in several numerical software packages. In this case the cost of the final state-space system is not greater than $C = N^2/2 + (P+M+1)N$.

3 Performance Characterization

To characterize the performance of the state-space systems described above, a numerical experiment is conducted in which multiple HRTF systems of varying size D, but fixed cost C, are constructed. The number of directions D that are required for a binaural display depends on the application. As such, we view D as an independent variable and choose the largest system order $N \in \mathbb{Z}_+$ such that the total computational cost of the system, as defined in Section 2.3, does exceed $C_{\max}$.

An FIR filter array is used as a baseline for comparison with the state-space systems. For any given cost bound, two FIR arrays are constructed, one with cost bounded by $C_{\max}$, and the other with cost bounded by $2C_{\max}$. We include the 'double-cost' FIR array to gauge the relative improvement in approximation quality of the state-space systems. In so doing, we will

demonstrate that for some configurations, a state-space system not only outperforms an FIR array of equal cost, but also outperforms an FIR array of twice the cost. This would seem to be a large enough margin of improvement to warrant the use of state-space systems in practical binaural displays. FIR filters of order N are constructed by truncating all but the first $N+1$ samples of minimum-phase HRIRs. Hence the FIR filters are optimal FIR approximations in terms of $\mathcal{L}^2$ error [29].

Two approximation errors are reported: the Hankel error and the $\mathcal{L}^\infty$ error. Perceptually weighted error metrics are reported in a companion study [32], in which different state-space architectures are considered for binaural display applications. In addition to the two error metrics, a few specific error responses are shown to compare BMT and HOA.

3.1 HRTF Measurements

The HRTFs used to design the low-order systems were measured at the Naval Submarine Medical Research Laboratory in Groton, CT. The HRTFs of eight individuals were measured. Using a vertical-polar coordinate system, HRTFs were measured in 10° increments in azimuth around the listener, and in 18° increments in elevation from $-36°$ to $+90°$, yielding a total of 253 pairs of HRTFs for each listener. For the experiment below, systems are designed that model D of the 253 measured directions. For every system, D directions are chosen randomly subject to a constraint that the D directions be approximately uniformly distributed around the listener.

Golay codes are employed to minimize bias in the measurement and identification process. The HRTF measurement process is described in detail in [33]. At a sampling rate of 44.1 kHz the measured HRIRs have length 256, order $N = 255$. The measured HRIRs are nearly minimum-phase. To simplify the analysis, and to guarantee the performance of the truncated FIR filters, all HRIRs are reconstructed as strictly minimum-phase for this experiment[6].

3.2 Hankel and $\mathcal{L}^\infty$ Results

For the main experiment, we select a cost bound of $C_{\max} = 4000$, which is approximately the cost of eight full-order HRIR pairs. Separate systems are

[6]We also applied the state-space methods to the nearly measured HRIRs and observed that performance did not change significantly.

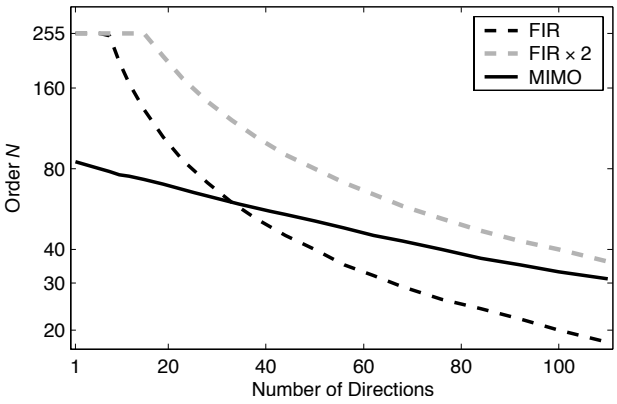

Figure 2: System order for three systems as a function of the number of directions D.

designed for each of the eight measured HRTF datasets for each number of directions. The results are then averaged across the eight individuals. We found that the approximation measures varied little across individual, so we report only mean values without standard deviation bars. For each HRTF dataset, 25 systems are approximated, with $1 \leq D \leq 110$. For the state-space approximats, low-order systems are designed using both BMT and HOA, as described in Section 2.2. Two FIR arrays are designed for each as well.

Figure 2 shows the order N for three systems: a MIMO state-space with cost $C \leq C_{\max}$, an FIR array with cost $C \leq C_{\max}$, and a second FIR array with cost $C \leq 2C_{\max}$. Note that for the FIR filter arrays, it is not necessary to truncate the HRIRs if $D \leq 8$ in order to satisfy the cost constraint. And for the 'double-cost' array no truncation in required for $D \leq 16$. Hence the approximation error for the two FIR arrays will be zero for $D \leq (8, 16)$.

To compute the Hankel error, the Hankel matrix of the error system for each low-order system is constructed and its largest singular value computed. Several efficient techniques for computing the $\mathcal{L}^\infty$ error have been published [34], although we found that a brute-force search over a finely sampled frequency grid yields accurate estimates of the $\mathcal{L}^\infty$ error with manageable computation time.

Figure 3 shows the Hankel error of state-space and FIR systems with order given by Figure 2. The performance of the FIR systems is shown with dashed lines, and the performance of the state-space systems is shown with

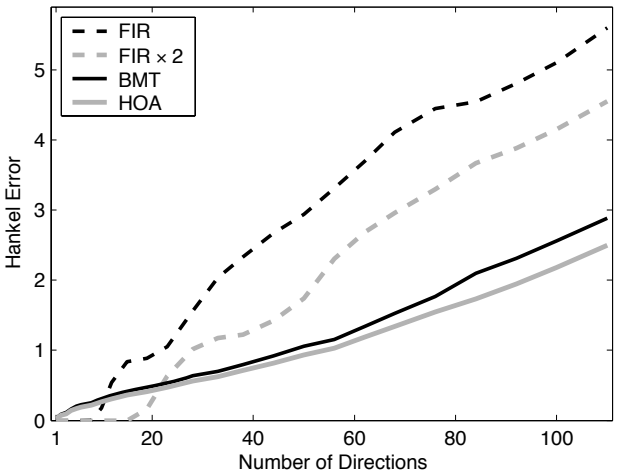

Figure 3: Hankel error as a function of D.

solid lines. The performance of the BMT state-space systems is shown in black, and that of the HOA systems is shown in gray. The performance of the 'single-cost' FIR systems is shown in black, and the 'double-cost' FIR systems in gray. Figure 4 shows the corresponding $\mathcal{L}^\infty$ error.

Both errors are non-decreasing functions of D for all systems. For small $D \leq 8$, the FIR systems exhibit no error. However, for larger D the error of the FIR systems increases rapidly relative to the error of the state-space systems. For $D > 20$ the error of the state-space systems is less than that of even the 'double-cost' FIR array. As expected, the Hankel error lower bounds the $\mathcal{L}^\infty$ error in all cases. For the FIR systems, the Hankel error is a relatively loose bound on the $\mathcal{L}^\infty$ error for $D > 40$, with the $\mathcal{L}^\infty$ error being 20-30% greater than the Hankel error. In contrast, for the state-space systems, the Hankel error is relatively tight, with the $\mathcal{L}^\infty$ error being 15-20% greater than the Hankel error.

3.3 Comparison of BMT and HOA

Overall, BMT and HOA exhibit similar performance in the design of low-cost state-space models of HRTFs. The HOA designs are found to yield slightly lower Hankel error and $\mathcal{L}^\infty$ error than the BMT designs. However, we have also found that BMT yields slightly lower $\mathcal{L}^2$ error [32,35]. This inconsistency

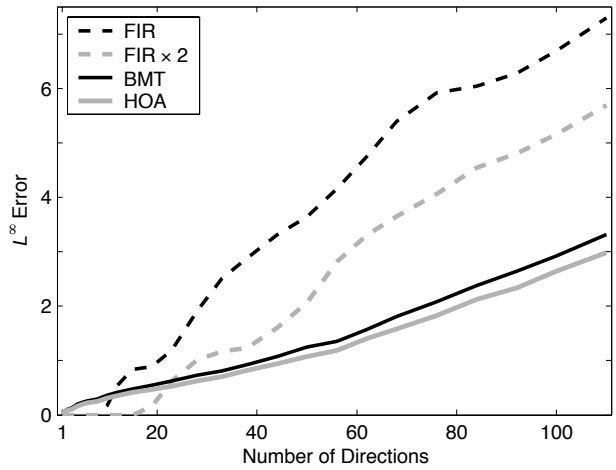

Figure 4: $\mathcal{L}^\infty$ error as a function of D.

is due to the difference between the ∞-norm and the 2-norm. The HOA method is known to 'distribute' the error approximately uniformly across frequency. In contrast, the error spectrum that results from the BMT method exhibits peaks and dips. The average (2-norm) spectral error exhibited by the BMT method is often smaller than that of the HOA method, but the peaks in the BMT error response rise above the relatively flat error response of the HOA method.

The performance trends outlined above may be observed in HRTF approximation, but are not strong. In the design of simple SISO filters, however, these trends are clear. A direct comparison of BMT and HOA for IIR filter design is given in [27]. Figure 5 shows the magnitude error for two examples. Two $N = 100$ FIR filters are constructed, one models a narrow notch filter, while the other models a wide band-pass filter. Low-order state-space systems are designed for these two FIR filters using both BMT and HOA. The HOA method exhibits relatively flat magnitude error. The BMT method exhibits lower magnitude error at most frequencies, but also exhibits substantial peaks. As has been previously observed [27], the BMT method often concentrates error near spectral notches, or near transition bands, although this trend is not universal. This property may be undesirable for binaural applications, as spectral notches are important for spatial hearing [36].

The trends outlined above for the SISO case can also be observed in simple

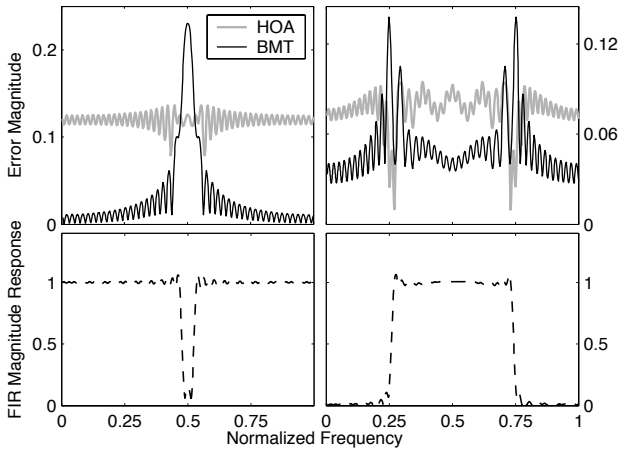

Figure 5: The bottom two panels show the magnitude response for two order 100 FIR filters; a notch filter (left) and a band-pass filter(right). The top two panels show the magnitude error of low-order state-space systems designed using either HOA or BMT. For the notch filter $N = 6$ state-space systems are designed, and for the band-pass filter $N = 10$ systems are designed.

MIMO filters. Figure 6 shows the magnitude response of four $N = 100$ FIR filters. The filters are simple in shape: band-pass, band-stop and low-pass. Two $N = 20$ state-space systems are constructed using BMT and HOA. The resulting error magnitude responses are shown in the top panel of Figure 6. The error peaks of the BMT system again occur near the transition bands of the original filters. Indeed, the highest error peak exhibited by the BMT approximant occurs at the frequency where all four transfer functions exhibit transitions between stop bands and pass bands. These performance trends are less clear when we examine measured HRTF approximations, however.

Figure 7 shows the magnitude response of one HRTF on a linear scale. The error magnitude for two $N = 6$ state-space systems is shown in the top panel, and for two $N = 30$ state-space systems is shown in the bottom panel. To facilitate comparison, the error magnitudes are shown in the same panel as the HRTF magnitude; the error magnitude values are given on the left of the figure, and the HRTF magnitude values are given on the right of the figure. The differences between the error magnitudes for the BMT and HOA systems are not as significant in this example as in the previous filter approximations. The HOA systems exhibit a flatter error magnitude than

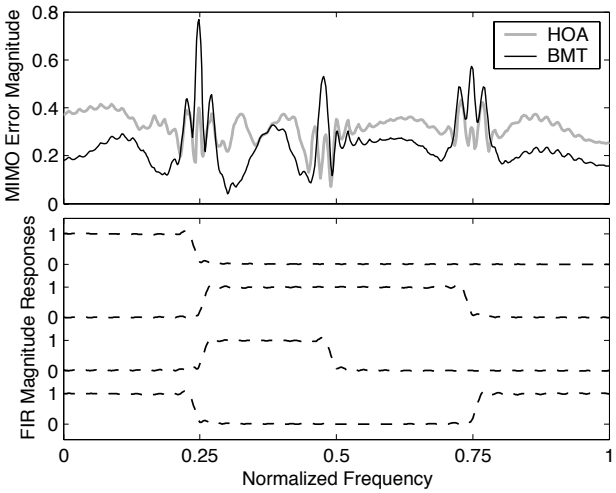

Figure 6: The bottom two panels show the magnitude responses for two-input, two-output FIR array with order $N = 100$. The top two panel shows $\sigma_1\big(\mathbf{H}(\omega) - \widehat{\mathbf{H}}(\omega)\big)$, for two $N = 20$ state-space systems designed using HOA and BMT.

the BMT systems, although at most frequencies the BMT error magnitude is lower than the HOA error. The BMT error magnitude exhibits peaks in the error at or near spectral notches in the original HRTF. For example, the errors of both the $N = 6$ and $N = 30$ BMT systems exhibit a peak near the HRTF notch at 5.5 kHz.

If MIMO HRTF systems are considered, the BMT and HOA methods yield similar performance. Figure 8 shows an example for a MIMO system that models $D = 44$ HRTF pairs. In this case the magnitudes shown are the maximum singular value of the matrix transfer functions. State-space systems with order $N = 40$ are designed with both BMT and HOA, and both systems yield error magnitudes between 1 and 1.3 for most frequencies. The BMT error magnitude again yields greater fluctuations than the HOA error magnitude. The BMT error exhibits several small peaks, some of which are located at or near notches in the HRTF response. Nonetheless, the two methods yield similar results and both are suitable for HRTF approximation.

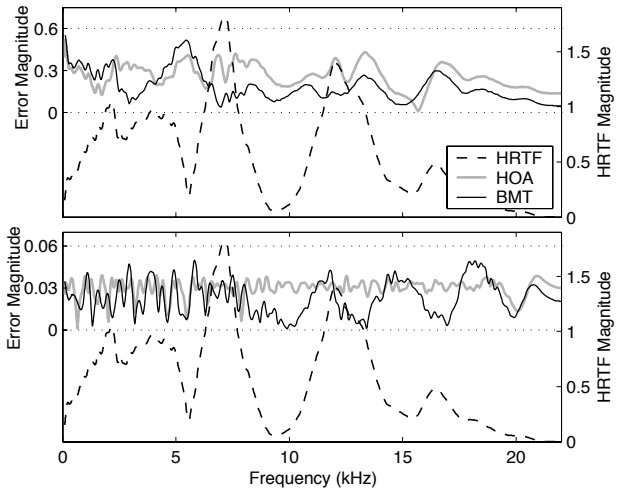

Figure 7: HRTF magnitude response for location $(\theta, \phi) = (30°, -36°)$, ipsilateral ear. Also shown are error magnitudes for four low-order approximations. The top panel shows the error responses for two $N = 6$ state-space systems designed using HOA and BMT. The bottom panel shows error responses for two $N = 30$ systems.

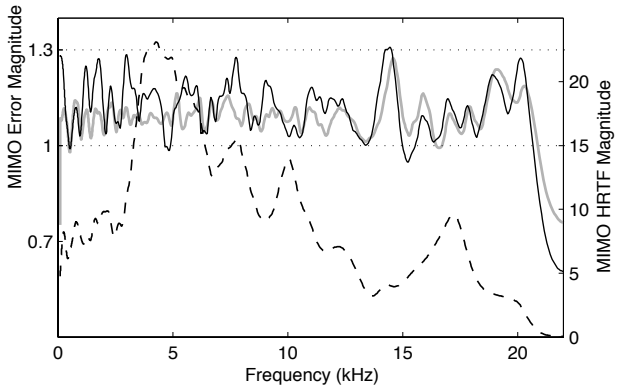

Figure 8: MIMO HRTF magnitude response for 44 directions surrounding the listener, $\sigma_1 \mathbf{H}(\omega)$, where $H(\omega)$ is a 2×44 matrix function. Also shown are MIMO error responses, $\sigma_1\big(\mathbf{H}(\omega) - \widehat{\mathbf{H}}(\omega)\big)$, for two $N = 40$ state-space systems designed using HOA and BMT.

3.4 Psychophysical Results

While the Hankel and $\mathcal{L}^\infty$ results presented above are promising, additional experiments are necessary to evaluate the perceptual efficacy of the proposed state-space methods. We have conducted experiments that employ simple models of auditory perception, based on a warped $\mathcal{L}^2$ norm [35]. Informal listening tests confirm the numerical results presented above: for a fixed C we are able to construct state-space systems that are perceptually indiscriminable from min.-phase HRIRs, whereas FIR arrays with the same cost C are discriminable from min.-phase HRIRs. A formal experiment is currently being conducted. The audible artifacts that sometimes appear in the design of low-order IIR filters for audio applications [13, 19] are not apparent in any of the state-space HRTF systems.

The results presented above do not reflect a perceptually critical feature of the HRIRs: the *interaural time delay* (ITD). We have found that in the design of low-order MIMO systems, the impulse responses are often smeared such that there is no longer a precise time-delay in the contralateral responses after model reduction. We have found that this problem can be mitigated by either employing a different state-space architecture or using a hybrid state-space/FIR system. Remarkably, we have found that the Hankel and $\mathcal{L}^\infty$ results are similarly promising with the alternative architectures. The alternative architectures and issues associated with the ITD are described in detail in [32].

4 Conclusion

The present work explores low-order state-space models of HRTFs. Many contemporary binaural environment models are implemented with an array of HRTF filters. If more than 20 directions are included in the array, we found that the array can be replaced with a state-space system of lower computational cost. Two order-reduction techniques are explored, BMT and HOA. Both methods are based on the Hankel operator and balanced systems theory, although HOA is optimal in the Hankel-norm sense whereas BMT is not optimal in any sense. In practice, the two methods are found to yield similar approximation quality for HRTFs. In terms of the Hankel error and $\mathcal{L}^\infty$ error, the HOA method yields slightly lower error. The BMT method often yields peaks in the error response near spectral notches in the original

HRTFs, although this effect is slight in the MIMO case.

Binaural displays are an ideal candidate for reduced-order state-space models, as there is a clear need to model multiple similar transfer functions simultaneously. We have demonstrated that for this application and state-space systems can achieve a significant computational savings relative to conventional filter arrays.

Appendix: Model Reduction Algorithms

Algorithms have been previously published for both Balanced Model Truncation (BMT) [20, 21] and Hankel-norm Optimal Approximation (HOA) [22, 37][7]. In the interest of completeness and reproducibility the algorithms have been adapted to the binaural display application and are described below in detail. In particular, the HOA algorithm is somewhat involved. Published HOA algorithms typically only provide a broad overview that does not describe several difficult steps. Each step of the HOA method is described below, with additional references for individual sub-algorithms as necessary.

The principal step in BMT is the computation of the SVD of Hankel matrix $\mathcal{H}$. This computation is straightforward, but $\mathcal{H}$ is often a large matrix and memory consumption may be an issue. In contrast, HOA operates entirely on the system matrices, hence memory consumption is not an issue. However, HOA may be more time consuming to compute, as solving Lyapunov equations and stable projections are computationally intensive steps for large systems.

Balanced Model Truncation

Balanced Model Truncation (BMT) is a state-space model reduction technique that operates by discarding all but the N largest singular values of a balanced system [21]. If the original system is given in state-space form, then it is *balanced* before truncation. For HRTF modeling, this is not necessary as the original system is a collection of transfer functions. In this case the Hankel matrix $\mathcal{H}$ is given by the block impulse response of the HRTFs. A balanced system is then constructed from the SVD of $\mathcal{H}$ [21].

[7]Both algorithms are included in the Robust Controls Toolbox, an extension of the Controls Toolbox, which is itself an extension of the MATLAB software package.

The BMT algorithm described below constructs an order N state-space system from a collection of measured HRIRs of order N_0. Both the originial HRIR filter array and the reduced-order state-space system have M inputs and P outputs.

1. Prepend a single zero to the beginning of each HRIR, as the state-space system in (1) has no feed-through path [20, 21]. For the MIMO architecture, the time-delay of the contralateral HRIRs must be included, but is removed for the SIMO and MISO architectures. In all other respects, the BMT algorithm is independent of system architecture.

2. Arrange the HRIRs into the desired matrix impulse response, $\mathbf{h}[n = 0, 1, \cdots N_0+1]$, where $\mathbf{h}[n]$ is $P \times M$, and construct the finite Hankel matrix $\mathcal{H}$

$$\mathbf{H} = \begin{bmatrix} \mathbf{h}[1] & \mathbf{h}[2] & \mathbf{h}[3] & \dots & \mathbf{h}[N_0+1] \\ \mathbf{h}[2] & \mathbf{h}[3] & \mathbf{h}[4] & \dots & \mathbf{0} \\ \mathbf{h}[3] & \mathbf{h}[4] & \mathbf{h}[5] & \dots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{h}[N_0+1] & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \end{bmatrix}$$

   The size of $\mathcal{H}$ is $P(N_0+1) \times M(N_0+1)$[8].

3. Compute the singular value decomposition (SVD) of $\mathcal{H}$

$$\mathcal{H} = \mathbf{U}\mathbf{S}\mathbf{V}^T \tag{17}$$

   where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices with dimensions $P(N_0+1) \times P(N_0+1)$ and $M(N_0+1) \times M(N_0+1)$, respectively, and $\mathbf{S}$ is a diagonal matrix of size $P(N_0+1) \times M(N_0+1)$ with the singular values, $(\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_R)$, arranged along the main diagonal, where $R = \min\big(P(N_0+1), M(N_0+1)\big)$.

[8] $\mathcal{H}$ may be very large depending upon (P, M, N_0), and computing its singular value decomposition may require a prohibitive amount of memory. However, almost half of $\mathcal{H}$ is zero, and given that the HRIRs are minimum-phase, most of the energy in $\mathcal{H}$ is contained in the upper-left corner. We found that the results are not affected substantially if only the upper-left quarter of $\mathcal{H}$ is used for BMT. All results given in the present study are computed using the full Hankel matrix.

4. Weight the singular vectors, given by the rows of $\mathbf{U}$ and $\mathbf{V}$, with the square root of the singular values[9]: $\widetilde{\mathbf{U}} = \mathbf{U}\mathbf{S}^{1/2}$ and $\widetilde{\mathbf{V}}^T = \mathbf{S}^{1/2}\mathbf{V}^T$. $\widetilde{\mathbf{U}}$ and $\widetilde{\mathbf{V}}$ both have dimensions $P(N_0+1) \times M(N_0+1)$. Hence $\mathcal{H} = \widetilde{\mathbf{U}}\widetilde{\mathbf{V}}^T$, where $\widetilde{\mathbf{U}}$ and $\widetilde{\mathbf{V}}$ are orthogonal but not orthnomornal.

5. Partition the SVD so as to retain only the N largest singular values[10]

$$\widetilde{\mathbf{U}} = \begin{bmatrix} \widetilde{\mathbf{U}}_1 & \dots \\ \widetilde{\mathbf{U}}_2 & \dots \\ \widetilde{\mathbf{U}}_3 & \dots \end{bmatrix}, \quad \widetilde{\mathbf{V}} = \begin{bmatrix} \widetilde{\mathbf{V}}_1 & \dots \\ \vdots & \ddots \end{bmatrix} \tag{18}$$

   where $\widetilde{\mathbf{U}}_1$ and $\widetilde{\mathbf{U}}_3$ have dimension $P \times N$, $\widetilde{\mathbf{U}}_2$ has dimension $P(N_0 - 1) \times N$, and $\widetilde{\mathbf{V}}_1$ has dimension $M \times N$.

6. Construct an order N state-space system from $\widetilde{\mathbf{U}}$ and $\widetilde{\mathbf{V}}$. The system matrices $\left(\widehat{\mathbf{A}}, \widehat{\mathbf{B}}, \widehat{\mathbf{C}}\right)$ are given by

$$\begin{aligned} \widehat{\mathbf{A}} &= \left(\begin{bmatrix} \widetilde{\mathbf{U}}_1 \\ \widetilde{\mathbf{U}}_2 \end{bmatrix}^T \begin{bmatrix} \widetilde{\mathbf{U}}_1 \\ \widetilde{\mathbf{U}}_2 \end{bmatrix} \right)^{-1} \begin{bmatrix} \widetilde{\mathbf{U}}_1 \\ \widetilde{\mathbf{U}}_2 \end{bmatrix}^T \begin{bmatrix} \widetilde{\mathbf{U}}_2 \\ \widetilde{\mathbf{U}}_3 \end{bmatrix} \\ \widehat{\mathbf{B}} &= \widetilde{\mathbf{V}}_1^T, \qquad \widehat{\mathbf{C}} = \widetilde{\mathbf{U}}_1 \end{aligned} \tag{19}$$

   where $\widehat{\mathbf{A}}$ is $N \times N$, $\widehat{\mathbf{B}}$ is $N \times M$, and $\widehat{\mathbf{C}}$ is $P \times N$.

Hankel-Norm Optimal Approximation

The Hankel-norm Optimal Approximation (HOA) algorithm also designs a low-order system by discarding all but the N largest singular values of the original system, albeit in such a way as to minimize the Hankel error of the resulting system. For the SISO case, closed-form HOA algorithms have been published [27, 38]. These studies use HOA to construct low-order IIR filters from FIR filters. For the MIMO case the algorithm is considerably

[9]The square-root operation is 'element-wise,' as $\mathbf{S}$ is not necessarily square.

[10]The SVD is the most computation-intensive part of the BMT algorithm. Because only the N largest singular values are used to construct the final state-space system, there is no need to compute the entire SVD. Only the N largest singular values, along with their singular vectors, need to be computed.

more dense. The HOA algorithm is further complicated if the original system contains repeated singular values. This complication is neglected in the algorithm below, as measured HRIRs invariably contain sufficient white observation noise to preclude any repeated singular values [23]. The observation noise also guarantees that the system is minimal.

The HOA algorithm described below constructs an order N state-space system from a collection of measured HRIRs of order N_0. Both the original HRIR filter array and the reduced-order state-space system have M inputs and P outputs.

1. Prepend a single zero to the beginning of each HRIR, as in the BMT algorithm. In addition to removing the need for a feed-through path in the state-space system, prepending a zero also transforms the transfer functions so as to be strictly proper, a requirement for HOA [22]. For the MIMO architecture, the time-delay in the contralateral HRIRs must be included, but is removed for the SIMO and MISO architectures.

2. Construct a high-order state-space system that implements the measured HRIRs exactly. The HOA algorithm operates directly on the system matrices ($\mathbf{A}$,$\mathbf{B}$,$\mathbf{C}$), hence it is necessary to realize the HRIR filter array as a state-space system prior to performing order reduction. This is readily accomplished with the controller canonical form [39]. Without loss of generality, consider a system with more outputs than inputs[11], $P \geq M$. The controller canonical realization of this filter array is an order $M(N_0+1)$ state-space system. The $\mathbf{A}_0$ matrix is $M(N_0+1) \times M(N_0+1)$, the $\mathbf{B}_0$ matrix is $M(N_0+1) \times M$, and both matrices are block diagonal

$$\mathbf{A}_0 = \begin{bmatrix} \mathbf{I}' & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{I}' & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{I}' \end{bmatrix} \quad \mathbf{B}_0 = \begin{bmatrix} \underline{1} & \underline{0} & \dots & \underline{0} \\ \underline{0} & \underline{1} & \dots & \underline{0} \\ \vdots & \vdots & \ddots & \vdots \\ \underline{0} & \underline{0} & \dots & \underline{1} \end{bmatrix}$$

[11] If $M > P$, swap inputs with outputs, perform HOA, and swap back when done: $\widehat{\mathbf{A}} = \mathbf{A}^T$, $\widehat{\mathbf{B}} = \mathbf{C}^T$, $\widehat{\mathbf{C}} = \mathbf{B}^T$.

where

$$\mathbf{I}' = \begin{bmatrix} 0 & 0 & 0 & 0 & & \cdot \\ 1 & 0 & 0 & & \cdot & \\ 0 & 1 & & \cdot & & 0 \\ 0 & & \cdot & & 0 & 0 \\ & \cdot & & 1 & 0 & 0 \\ \cdot & & 0 & 0 & 1 & 0 \end{bmatrix} \quad \underline{1} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{20}$$

$\mathbf{I}'$ has dimension $(N_0+1)\times(N_0+1)$, and $\underline{1}$ has dimension $(N_0+1)\times 1$. The $\mathbf{C}_0$ matrix is $P \times M(N_0+1)$ and is constructed from the measured HRIRs

$$\mathbf{C}_0 = \begin{bmatrix} h_{11}[1] & h_{11}[2] & \dots & h_{11}[N_0+1] & \\ h_{21}[1] & h_{21}[2] & \dots & h_{21}[N_0+1] & \\ \vdots & \vdots & \ddots & \vdots & \cdots \\ h_{P1}[1] & h_{P1}[2] & \dots & h_{P1}[N_0+1] & \end{bmatrix}$$
$$\begin{bmatrix} h_{12}[1] & h_{12}[2] & \dots & h_{1M}[N_0+1] \\ h_{22}[1] & h_{22}[2] & \dots & h_{2M}[N_0+1] \\ \vdots & \vdots & \ddots & \vdots \\ h_{P2}[1] & h_{P2}[2] & \dots & h_{PM}[N_0+1] \end{bmatrix}$$

where $h_{pm}[n]$ is the impulse response between input m and output p.

3. Convert the discrete-time system above to continuous-time. The HOA algorithm is simpler in continuous-time, and it is common even when designing low-order discrete-time systems to convert the original system to continuous-time using a bilinear transform, and then convert back to discrete-time after performing HOA [22, 23, 40]. A discrete-time HOA algorithm is given in [37], although it is more dense than the algorithm below. The bilinear transform to continuous-time is

$$\begin{aligned} \mathbf{A}_c &= \left(\mathbf{I}+\mathbf{A}_0\right)^{-1}\left(\mathbf{A}_0-\mathbf{I}\right) \\ \mathbf{B}_c &= \sqrt{2}\left(\mathbf{I}+\mathbf{A}_0\right)^{-1}\mathbf{B}_0 \\ \mathbf{C}_c &= \sqrt{2}\,\mathbf{C}_0\left(\mathbf{I}+\mathbf{A}_0\right)^{-1} \end{aligned} \tag{21}$$

4. The controllability and observability Gramians, $\mathcal{P}$ and $\mathcal{Q}$, of the system

$(\mathbf{A_c}, \mathbf{B_c}, \mathbf{C_c})$ are defined as

$$\begin{aligned} \mathcal{P} &\triangleq \int_0^\infty e^{\mathbf{A}_c t}\, \mathbf{B}_c\, \mathbf{B}_c^*\; e^{\mathbf{A}_c^* t}\, dt \\ \mathcal{Q} &\triangleq \int_0^\infty e^{\mathbf{A}_c^* t}\, \mathbf{C}_c^*\, \mathbf{C}_c\; e^{\mathbf{A}_c t}\, dt \end{aligned} \tag{22}$$

Both Gramians are size $M(N_0{+}1) \times M(N_0{+}1)$. The system $(\mathbf{A}_c,\mathbf{B}_c,\mathbf{C}_c)$ is stable, and the eigenvalues of $\mathbf{A}_c$ are strictly in the left half of the complex plane. Hence the integrals above converge. However, this is a numerically prohibitive integral to evaluate. The Gramians are typically computed by considering the corresponding matrix differential equations, which yield the following linear equations, known as the Lyapunov equations

$$\begin{aligned} \mathbf{A}_c\mathcal{P} + \mathcal{P}\mathbf{A}_c^* + \mathbf{B}_c\mathbf{B}_c^* &= 0 \\ \mathbf{A}_c^*\mathcal{Q} + \mathcal{Q}\mathbf{A}_c + \mathbf{C}_c^*\mathbf{C}_c &= 0 \end{aligned} \tag{23}$$

An efficient algorithm for solving matrix equations of this form is given in [41]. Because the system $(\mathbf{A}_0, \mathbf{B}_0, \mathbf{C}_0)$ is in controller canonical form, the controllability Gramian is simply the identity matrix, and this simplifies the computation of a balancing transform somewhat [22]. The observability Gramian $\mathcal{Q}$ must be computed by solving the Lyapunov equation.

5. Find a balancing transform $\mathbf{T}$ from the observability Gramian $\mathcal{Q}$. There are many balancing transforms, however care must be taken in choosing a transform, as they may be ill-conditioned [17]. The transform below is well-conditioned for all of the HRTF data used in the present study [28].

   Compute the SVD of $\mathcal{Q}$

   $$\mathcal{Q} = \mathbf{V}\mathbf{S}\mathbf{V}^T \tag{24}$$

   where $\mathbf{V}$ is a unitary matrix and the singular values of $\mathcal{Q}$, $(\sigma_1 > \sigma_2 > \cdots > \sigma_{M(N_0+1)})$, are arranged along the diagonal of $\mathbf{S}$. The symmetry of the SVD in this case is due to the symmetry of the $\mathcal{Q}$. Let

   $$\mathbf{U} = \mathbf{V}\mathbf{S}^{1/4} \tag{25}$$

   The matrix $\mathbf{U}^T$ is itself a balancing transform. However, we seek to isolate the state that corresponds to the $(N{+}1)^{\text{th}}$ largest singular value.

We move this state to the end of the state-vector by permuting the columns of the transform matrix

$$\widetilde{\mathbf{U}} = \big[\underline{u}_1 \ \cdots \ \underline{u}_N, \ \underline{u}_{N+2} \ \cdots \ \underline{u}_{M(N_0+1)}, \ \underline{u}_{N+1}\big]$$

where $\underline{u}_k$ is the k^{th} column of $\mathbf{U}$. The balancing transform is then given by $\mathbf{T} = \widetilde{\mathbf{U}}^T$.

6. Balance the system using similarity transform $\mathbf{T}$

$$\begin{aligned} \mathbf{A}_b &= \mathbf{T}\,\mathbf{A}_c\,\mathbf{T}^{-1} \\ \mathbf{B}_b &= \mathbf{T}\,\mathbf{B}_c \\ \mathbf{C}_b &= \mathbf{C}_c\,\mathbf{T}^{-1} \end{aligned}$$

7. Partition the matrices so as to isolate the state corresponding to the $(N+1)^{\text{th}}$ singular value

$$\mathbf{A}_b = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \quad \mathbf{B}_b = \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \end{bmatrix} \quad \mathbf{C}_b^T = \begin{bmatrix} \mathbf{C}_1^T \\ \mathbf{C}_2^T \end{bmatrix}$$

where $\mathbf{A}_{11}$ has dimensions $\big(M(N_0+1)-1\big) \times \big(M(N_0+1)-1\big)$, $\mathbf{A}_{12}$ and $\mathbf{A}_{21}$ are vectors, and $\mathbf{A}_{22}$ is a scalar. Matrix $\mathbf{B}_1$ has dimensions $\big(M(N_0+1)-1\big) \times M$, matrix $\mathbf{C}_1$ has dimensions $P \times \big(M(N_0+1)-1\big)$, and $\mathbf{B}_2$ and $\mathbf{C}_2$ are vectors.

8. Let $\mathbf{\Gamma} = \mathbf{\Sigma}_1 - \sigma_{N+1}\mathbf{I}$, where

$$\mathbf{\Sigma}_1 = \text{diag}\big(\sigma_1 \ \cdots \ \sigma_N, \ \sigma_{N+2} \ \cdots \ \sigma_{P(N_0+1)}\big)$$

and σ_k is the k^{th} singular value of $\mathcal{Q}$. Also let $\mathbf{W} = (\underline{C}_2^T)^\dagger \underline{B}_2$, where † represents the Moore-Penrose pseudoinverse[12].

9. Construct the following order $M(N_0+1)-1$ system

$$\begin{aligned} \widetilde{\mathbf{A}} &= \mathbf{\Gamma}^{-1}\big(\mathbf{A}_{11}\sigma_{N+1} + \mathbf{\Sigma}_1\mathbf{A}_{11}\mathbf{\Sigma}_1 + \sigma_{N+1}\mathbf{C}_1^T\mathbf{W}\mathbf{B}_1^T\big) \\ \widetilde{\mathbf{B}} &= \mathbf{\Gamma}^{-1}\big(\mathbf{\Sigma}_1\mathbf{B}_1 - \sqrt{\sigma_{N+1}}\mathbf{C}_1^T\mathbf{W}\big) \\ \widetilde{\mathbf{C}} &= \mathbf{C}_1\mathbf{\Sigma}_1 - \sqrt{\sigma_{N+1}}\mathbf{W}\mathbf{B}_1^T \end{aligned} \tag{26}$$

[12]This step is commonly stated as: find a unitary matrix $\mathbf{W}$ that satisfies $\mathbf{C}_2^T\mathbf{W} = \mathbf{B}_2$. The pseudoinverse provides one possible solution [22, 42].

which is referred to as an *all pass dilation* of $\left(\mathbf{A}_b, \mathbf{B}_b, \mathbf{C}_b\right)$, as the spectral error between the two systems is constant

$$\sigma_{\max}\left(\mathbf{T}_b(j\omega) - \widetilde{\mathbf{T}}(j\omega)\right) = \sqrt{\sigma_{N+1}} \tag{27}$$

where $\mathbf{T}_b(j\omega)$ and $\widetilde{\mathbf{T}}(j\omega)$ are the matrix transfer functions for the two systems.

10. The system $\left(\widetilde{\mathbf{A}}, \widetilde{\mathbf{B}}, \widetilde{\mathbf{C}}\right)$ is not stable. Of the $M(N_0+1)-1$ eigenvalues in this system, exactly N of them are stable. It is necessary to extract the stable order N subsystem from $\left(\widetilde{\mathbf{A}}, \widetilde{\mathbf{B}}, \widetilde{\mathbf{C}}\right)$. There are several methods for accomplishing this. The most direct method is to compute the modal decomposition of $\widetilde{\mathbf{A}}$ [30], and constructing a diagonal $\widehat{\mathbf{A}}_c$ from the N eigenvalues with negative real part. The N corresponding eigenvectors similarly transform $\widetilde{\mathbf{B}}$ and $\widetilde{\mathbf{C}}$. A more involved method that yields real system matrices is described below [22, 42].

    (a) Compute the complex Schur decomposition of $\widetilde{\mathbf{A}}$. Find a unitary matrix $\mathbf{U}$ such that $\mathbf{U}^*\widetilde{\mathbf{A}}\mathbf{U} = \widetilde{\mathbf{A}}_t$ is a triangular matrix with diagonal elements given by the eigenvalues of $\widetilde{\mathbf{A}}$.

    (b) It is necessary to transform $\widetilde{\mathbf{A}}_t$ such that the N stable eigenvalues appear as the first N diagonal elements. This can be accomplished by applying a sequence of *Givens rotations* to the $\widetilde{\mathbf{A}}_t$ and $\mathbf{U}$ matrices [31], yielding a new unitary transform matrix $\widetilde{\mathbf{U}}$

$$\widetilde{\mathbf{U}}^*\widetilde{\mathbf{A}}\widetilde{\mathbf{U}} = \widetilde{\mathbf{A}}_p = \begin{bmatrix} \widetilde{\mathbf{A}}_{11} & \widetilde{\mathbf{A}}_{12} \\ \mathbf{0} & \widetilde{\mathbf{A}}_{22} \end{bmatrix} \tag{28}$$

    where $\widetilde{\mathbf{A}}_p$ is real and $\widetilde{\mathbf{A}}_{11}$ is $N \times N$.

    (c) Find a matrix $\mathbf{X}$ that satisfies

$$\widetilde{\mathbf{A}}_{11}\mathbf{X} - \mathbf{X}\widetilde{\mathbf{A}}_{22} + \widetilde{\mathbf{A}}_{12} = \mathbf{0} \tag{29}$$

    This is similar to solving the continuous-time Lyapunov equations, and the algorithm in [41] can be used.

(d) The system matrices of an order N HOA continuous-time system are given by

$$\begin{aligned} \widehat{\mathbf{A}}_c &= \widetilde{\mathbf{A}}_{11} \\ \widehat{\mathbf{B}}_c &= \left[\mathbf{I},\ -\mathbf{X}\right]\widetilde{\mathbf{U}}^*\,\widetilde{\mathbf{B}} \\ \widehat{\mathbf{C}}_c &= \widetilde{\mathbf{C}}\,\widetilde{\mathbf{U}}\begin{bmatrix}\mathbf{I}\\ \mathbf{0}\end{bmatrix} \end{aligned} \tag{30}$$

where $\mathbf{I}$ is $N \times N$.

11. The system matrices of the final order N discrete-time system are given by the bilinear transformation of the continuous-time solution

$$\begin{aligned} \widehat{\mathbf{A}} &= \left(\mathbf{I} + \widehat{\mathbf{A}}_c\right)\left(\mathbf{I} - \widehat{\mathbf{A}}_c\right)^{-1} \\ \widehat{\mathbf{B}} &= \sqrt{2}\left(\mathbf{I} - \widehat{\mathbf{A}}_c\right)^{-1}\widehat{\mathbf{B}}_c \\ \widehat{\mathbf{C}} &= \sqrt{2}\,\widehat{\mathbf{C}}_c\left(\mathbf{I} - \widehat{\mathbf{A}}_c\right)^{-1} \end{aligned} \tag{31}$$

References

[1] C. Cheng and G. Wakefield, "Introduction to Head-Related Transfer Functions (HRTFs): Representations of HRTFs in Time, Frequency, and Space," *J. Audio Eng. Soc.*, vol. 9, no. 4, pp. 231–249, 2001.

[2] J. Huopaniemi, N. Zacharov, and M. Karjalainen, "Objective and Subjective Evaluation of Head-Related Transfer Function Filter Design," *J. Audio Eng. Soc.*, vol. 47, no. 4, pp. 218–239, 1999.

[3] B. Shinn-Cunningham and A. Kulkarni, "Recent Developments in Virtual Auditory Space," in *Virtual Auditory Space: Generation and Applications*, S. Carlile, Ed. Berlin, Germany: Springer-Verlag, 1996.

[4] D. Begault and E. Wenzel, "Direct Comparison of the Impact of Head Tracking, Reverberation, and Individualized Head-Related Transfer Functions on Spatial Perception of a Virtual Speech Source," *J. Audio Eng. Soc.*, vol. 49, no. 10, pp. 904–916, Oct. 2001.

[5] D. Zotkin, R. Duraiswami, and L. Davis, "Rendering localized spatial audio in a virtual auditory space," *IEEE Trans. Multimedia*, vol. 6, no. 4, pp. 553–564, Aug. 2004.

[6] L. Savioja, J. Huopaniemi, T. Lokki, and R. Väänänen, "Creating Interactive Virtual Acoustic Environments," *J. Audio Eng. Soc.*, vol. 47, no. 9, pp. 675–705, Sept. 1999.

[7] S. Takane, Y. Suzuki, T. Miyajime, and T. Sone, "ADISE: A new method for high definition virtual acoustic display," in *Proc. Int. Conf. on Auditory Display*, 2002, Kyoto, Japan.

[8] H. Hacihabiboglu, "A fixed-cost variable-length auralization filter model utilizing the precedence effect," in *Proc. IEEE Workshop of App. of Signal Processing to Audio and Acoust.*, 2003, New Paltz, NY.

[9] R. Heinz, "Binaural Room Simulation Based on an Image Source Model with Addition of Statistical Methods to Include the Diffuse Sound Scattering of Walls and to Predict the Reverberant Tail," *Applied Acoustics*, vol. 38, pp. 145–159, 1993.

[10] V. Algazi, R. Duda, and D. Thompson, "Motion Tracked Binaural Sound," *J. Audio Eng. Soc.*, vol. 52, no. 11, pp. 1142–1156, Nov. 2004.

[11] F. Freeland, L. Biscainho, and P. Diniz, "Interpositional Transfer Function for 3D-Sound Generation," *J. Audio Eng. Soc.*, vol. 52, no. 9, pp. 915–930, Sept. 2004.

[12] C. Cheng and G. Wakefield, "Moving Sound Source Synthesis for Binaural Electroacoustic Music Using Interpolated Head-Related Transfer Functions (HRTFs)," *Computer Music Journal*, vol. 25, no. 4, pp. 57–80, 2001.

[13] N. Adams and G. Wakefield, "The binaural display of clouds of point sources," *Proc. IEEE Workshop on App. of Signal Processing to Audio and Acoust.*, October 2005, New Paltz, NY.

[14] D. Kistler and F. Wightman, "A model of head-related transfer functions based on principal components analysis and minimum-phase reconstruction," *J. Acoust. Soc. Am.*, vol. 91, no. 3, pp. 1637–1647, Mar. 1992.

[15] Y. Haneda, S. Makino, Y. kaneda, and N. Kitawaki, "Common-Acoustical-Pole and Zero Modeling of Head-Related Transfer Functions," *IEEE Trans. Speech and Audio Processing*, vol. 7, no. 2, pp. 188–196, Mar. 1999.

[16] J. Bay, *Fundamentals of linear state space systems.* Boston, MA: McGraw-Hill, 1999.

[17] B. Beliczynski, I. Kale, and G. Cain, "Approximation of FIR by IIR Digital Filters: An Algorithm Based on Balanced Model Truncation," *IEEE Trans. Signal Processing*, vol. 40, no. 3, pp. 532–542, Mar. 1992.

[18] J. Mackenzie, J. Huopaniemi, V. Välimäki, and I. Kale, "Low-Order Modeling of Head-Related Transfer Functions Using Balanced Model Truncation," *IEEE Signal Processing Letters*, vol. 4, no. 2, pp. 39–41, Feb. 1997.

[19] P. Georgiou and C. Kyriakakis, "Modeling of Head Related Transfer Functions for Immersive Audio Using a State-Space Approach," in *Proc. IEEE Asilomar Conf. on Signals, Systems and Computers*, vol. 1, 1999, pp. 720–724.

[20] D. Grantham, J. Willhite, K. Frampton, and D. Ashmead, "Reduced order modeling of head related impulse responses for virtual acoustic displays," *J. Acoust. Soc. Am.*, vol. 117, no. 5, pp. 3116–3125, May 2005.

[21] S. Kung, "A new identification and model reduction algorithm via singular value decompositions," in *Proc. IEEE Asilomar Conf. on Signals, Systems and Computers*, 1978, pp. 705–714.

[22] K. Glover, "All optimal Hankel-norm approximations of linear multivariable systems and their L^∞-error bounds," *Int'l. J. Control*, vol. 39, pp. 1115–1193, 1984.

[23] A. Antoulas, "Approximation of linear dynamical systems," in *Wiley Encyclopedia of Electrical and Electronics Engineering*, J. Webster, Ed. New York: John Wiley and Sons, Inc., 1999, vol. 11, pp. 403–422.

[24] T. Hartley, R. Veillette, J. Abreu-Garcia, A. Chicatelli, and R. Hartmann, "To Err is Normable: The Computation of Frequency-Domain Error Bounds From Time-Domain Data," *NASA Contractor Report NASA/CR-1998-208516*, 1998, available online: http://gltrs.grc.nasa.gov/reports/1998/CR-1998-208516.pdf.

[25] K. Glover, R. Curtain, and J. Partington, "Realization and approximation of linear finite-dimensional systems with error bounds," *SIAM J. Control and Optimization*, vol. 26, no. 4, pp. 863–898, July 1988.

[26] V. Adamjan, D. Arov, and M. Krein, "Analytic properties of Schmidt pairs for a Hankel operator and the generalized Schur-Takagi problem," *Math. USSR Sbornik*, vol. 15, pp. 31–73, 1971.

[27] B. Beliczynski, J. Gryka, and I. Kale, "Critical comparison of hankel-norm optimal approximation and balanced model truncation algorithms as vehicles for fir-to-iir filter order reduction," in *Proc. IEEE Conf. on Acoust., Speech, and Signal Processing (ICASSP)*, 1994.

[28] B. Moore, "Principal Component Analysis in Linear Systems: Controllability, Observability, and Model Reduction," *IEEE Trans. Automatic Control*, vol. AC-26, no. 1, pp. 17–32, Feb. 1981.

[29] A. Oppenheim and R. Schafer, *Discrete-Time Signal Processing*. Englewood Cliffds, NJ: Prentice Hall, 1989.

[30] R. Horn and C. Johnson, *Matrix Analysis*. Cambridge, UK: Cambridge University Press, 1990.

[31] G. Golub and C. V. Loan, *Matrix Computations*. Baltimore, MD: John Hopkins University Press, 1996.

[32] N. Adams, "State-space architectures for binaural environment modeling," 2007, Technical Report available online at http://www.eecs.umich.edu/systems/TechReportList.html.

[33] C. Cheng, "Visualiation, Measurement, and Interpolation of Head-Related Transfer Functions (HRTF'S) with Applications in Electro-Acoustic Music," Ph.D. dissertation, Univ. of Michigan, Ann Arbor, 2001.

[34] G. Robel, "On computing the infinite norm," *IEEE Trans. on Auto. Control*, vol. 34, no. 8, pp. 882–884, August 1989.

[35] N. Adams and G. Wakefield, "Efficient binaural display using MIMO state-space systems," *Proc. IEEE Int. Conf. on Acout. Speech and Sig. Proc.*, April 2007, Honolulu, HI.

[36] J. Blauert, *Spatial Hearing: The Psychophysics of Human Sound Localization.* Cambridge, MA: MIT Press, 1997.

[37] C. K. Chui and G. Chen, *Discrete H^∞ Optimization.* Berlin: Springer-Verlag, 1997.

[38] B. Chen, S. Peng, and B. Chiou, "IIR filter design via optimal Hankel-norm approximation," *IEE Proc. G., Circuits, Devices and Systems*, vol. 139, no. 5, pp. 586–590, Oct. 1992.

[39] D. Luenberger, "Canonical Forms for Linear Multivariable Systems," *IEEE Trans. Auto. Control*, pp. 290–293, June 1967.

[40] J. Huang and G. Gu, "A Direct Approach to the Design of QMF Banks via Frequency Domain Optimization," *IEEE Trans. on Signal Processing*, vol. 46, no. 8, pp. 2131–2138, August 1998.

[41] R. Bartels and G. Stewart, "Solution of the Matrix Equation AX + XB = C," *Comm. of the ACM*, vol. 15, no. 9, pp. 820–826, Sept. 1972.

[42] L. Fortuna, G. Nunnari, and A. Gallo, *Model Order Reduction Techniques with Applications in Electrical Engineering.* London: Springer-Verlag, 1992.