Spectral Representation of Transient Signals

Tarek A. Lahlou and Alan V. Oppenheim
Digital Signal Processing Group
Massachusetts Institute of Technology

***Abstract*—Signal processing techniques exploiting natural and efficient representations of a class of signals with an underlying parametric model have been extensively studied and successfully applied across many disciplines. In this paper, we focus attention to the representation of one such class, i.e. transient structured signals. The class of transient signals in particular often results in computationally ill-conditioned problems which are further degraded by the presence of noise. We develop the Discrete Transient Transform, a biorthogonal transform to a basis parameterized by decay rate, along with algorithms for its implementation which mitigate these numerical issues and enable a spectral approach to parameter identification, estimation, and modeling for signals with transient behavior. The three algorithms developed have varying degrees of numerical robustness for generating the biorthogonal transient basis. Issues pertaining to transient spectral leakage and resolution are characterized and discussed in the context of an example related to Vandermonde system inversion.**

***Index Terms*—The discrete transient transform, nonlinear filtering, biorthogonal transform, real exponentials**

I. Introduction

Signals exhibiting exponential behavior play a fundamental role in both applied and theoretical disciplines, e.g., solutions to a broad class of differential equations, load envelopes in non-intrusive load monitoring systems, nuclear decay in quantum theory, primary electronic component voltage-current characteristics, atmospheric pressure, evanescent acoustic waves, etc. [1][2][3][4][5] An exponential signal explicitly characterizes a relationship for which a constant change in the independent variable corresponds to constant proportional change in the dependent variable. For these and additional reasons, transient signal parameters, namely decay rates and amplitude coefficients, often carry informational significance with respect to the signals origin; thus establishing a motivation to determine or estimate them in either a parameter identification or modeling scenario, respectively.

The identification and estimation of signal parameters for both exponential and more general signal models using iterative, subspace, spectral-projection, and curve-fitting techniques has been considered previously, e.g., Prony's method and related algorithms, all-pole modeling, finite rate of innovation, Cadzows method, various regression techniques, etc. [6][7][8][9] A method which specifically addresses real, continuous-time exponentials has also been considered in [10] which makes use of Jacobi polynomial expansions. However, the application of the general techniques to transient discrete-time data often results in large numerical errors due in part to the ill-conditioning of the underlying problem. To address this issue, we propose a biorthogonal transform enabling a spectral approach to representation, identification, estimation, and modeling of transient signal parameters. An algorithm for specifically generating the biorthogonal transient basis is developed which exploits the underlying structure of the transient signal model in order to better control numerical issues and increase the robustness of analysis as compared with the more general techniques.

The authors wish to thank Analog Devices, Bose Corporation, and Texas Instruments for their support of innovative research at MIT and within the Digital Signal Processing Group.

A straightforward consequence of the z-transform representation of a system whose impulse response is an infinitely long causal transient signal requires all poles of the system to lie strictly on a bounded interval of the real axis. This observation motivates the general approach underlying the methodology in this paper: to form a parameterized spectrum along this interval with respect to decay rate. Advantages in taking this spectral approach include efficient and natural representation of transient-like signals which often remain dense in other canonical representations, such as with shifted impulses, sinusoids, and wavelet bases. Interpreting the transient spectrum with regard to spectral resolution and leakage is analogous to the interpretation issues of frequency spectra produced by the Discrete Fourier Transform (DFT).

We proceed as follows: Section II defines the parametric structure of the class of transient signals considered in this paper and discusses numerical conditioning, additive noise considerations, and the shortcomings of well-known noise reduction and model-order estimation techniques. Section III defines the discrete transient transform and presents three algorithms for its implementation. Finally in Section IV, through the context of an example, issues of transient spectral leakage, resolution, and numerical stability are presented.

II. Transient Signal Structure

In defining a parametric model for transient structured signals, we limit our attention to the N-dimensional vector space $\mathbb{R}^N$ and the space of finite-energy signals ℓ_2, with primary focus on the former. For these spaces, define the inner product $\langle \cdot, \cdot \rangle : \mathcal{V} \times \mathcal{V} \to \mathbb{R}$, for any two vectors $\underline{v}, \underline{w} \in \mathcal{V}$, as

$$\langle \underline{v}, \underline{w} \rangle = \sum_n v[n]w[n] \tag{1}$$

where the limits of summation are determined by the dimensionality of the vector space $\mathcal{V}$. In the following definition we state the transient signal structure in the setting of ℓ_2 with a natural restriction via orthogonal projection onto $\mathbb{R}^N$ to follow.

Definition II-1. The Transient Signal Model.
A transient signal $x_d[n]$ is formed as a linear combination of a finite number d of real decaying exponentials with the structure

$$x_d[n] = \sum_{k=1}^{d} \alpha_k (\rho_k)^n, \;\; \rho_k \neq 0, -1 < \rho_1 < \cdots < \rho_d < 1, \tag{2}$$

for all $n \geq 0$.

We proceed utilizing $x_d[n]$ restricted to a finite interval of support, i.e., we focus our treatment from ℓ_2 to $\mathbb{R}^N$ by utilizing $x_d[n]$ only for $0 \leq n \leq N-1$. The justification of this restriction is twofold.

First, consider the relative conditioning of a Vandermonde matrix $\Phi : \mathbb{R}^d \to \mathbb{R}^N$ defined using a set of distinct real nodes $\underline{\sigma} = \{\sigma_k\}_{k=1}^d$, i.e.

$$\Phi_{i,j} = \sigma_j^{i-1}, \quad 1 \leq i \leq N, 1 \leq j \leq d. \tag{3}$$

When the nodes $\underline{\sigma}$ are chosen to be the decay rates $\underline{\rho} = \{\rho_k\}_{k=1}^d$ and are assumed fixed then Φ has an interpretation as the linear mapping between the amplitude coefficients $\underline{\alpha} = \{\alpha_k\}_{k=1}^d$ and the sequence values $\underline{x}_d = \{x_d[n]\}_{n=0}^{N-1}$. Consider a perturbation from $\underline{x}_d$, denoted $\delta\underline{x}_d$, produced at the output of the system Φ resulting from a small perturbation to the amplitude coefficients $\underline{\alpha}$, denoted $\delta\underline{\alpha}$. The corresponding relative conditioning, given by

$$\lim_{\gamma\to 0}\sup_{\|\delta\underline{\alpha}\|_d\leq\gamma}\left(\frac{\|\delta\underline{x}_d\|_N}{\|\underline{x}_d\|_N}\right)\left(\frac{\|\delta\underline{\alpha}\|_d}{\|\underline{\alpha}\|_d}\right)^{-1}, \tag{4}$$

has been shown [11] to grow exponentially in d and/or N for specific choices of amplitude coefficients and decay rates where $\|\cdot\|_d$ and $\|\cdot\|_N$ are the respective Euclidean norms on the vector spaces $\mathbb{R}^d$ and $\mathbb{R}^N$. The behavior of Eq. (4) for closely spaced real nodes illuminates the conditioning issues pertaining to the transient signal structure and underlies the instability of analysis algorithms, even for small data sets and exact arithmetic.

Second, in restricting our treatment to finite-length transient signals, Definition II-1 and the remaining presentation of the paper may be extended in a straightforward manner to include exponentially growing components, i.e., we may allow values of ρ_k for which $|\rho_k| > 1$ without introducing issues pertaining to convergence.

We next couple the transient signal model with an additive noise source to account for potential distortion resulting from sampling hardware and/or the environment, e.g., sampling jitter, thermal noise, etc. We denote this model by $\bar{x}_d[n] = x_d[n] + \eta[n]$ where $\eta[n]$ is an additive noise process with finite Fourier spectral density. The instantaneous signal-to-noise ratio (SNR) of $\bar{x}_d[n]$ tends to $-\infty$ as the interval of support N grows large. For this reason there is an inherent tradeoff in the choice of the data record length.

We next emphasize a key difficulty in using general model order identification, noise reduction, and subspace projection methods on discrete-time transient data by comparing the Singular Value Decomposition (SVD), which is utilized in a large number of these general methods, with a structurally similar transient decomposition of the data, for which a stable algorithm is unknown. Loosely speaking, the comparison highlights geometrically why SVD based methods have poor performance in capturing transient signal parameters. Consider first the SVD of the order $\hat{d} \geq d$ Toeplitz matrix $X^{(\hat{d})}$ given by

$$X^{(\hat{d})} = \begin{bmatrix} x_d\left[\hat{d}-1\right] & \cdots & x_d[0] \\ \vdots & \ddots & \vdots \\ x_d[N-1] & \cdots & x_d\left[N-\hat{d}\right] \end{bmatrix} = \sum_{k=1}^{d}\pi_k\underline{u}_k\underline{v}_k^T$$

where $\{\pi_k\}_{k=1}^d, \{\underline{u}_k\}_{k=1}^d$, and $\{\underline{v}_k\}_{k=1}^d$ are the singular values, orthonormal left-singular vectors and orthonormal right-singular vectors, and consider next the transient decomposition, given by

$$X^{(\hat{d})} = F\left(\underline{\rho}\right)D\left(\underline{\alpha}\right)H\left(\underline{\rho}\right)^T \tag{5}$$

$$= \sum_{k=1}^{d}\alpha_k\underline{f}_{\rho_k}\underline{h}_{\rho_k}^T \tag{6}$$

where

$$\underline{f}_{\rho_k} = \left[1, \rho_k, \cdots, \rho_k^{N-\hat{d}}\right]^T \tag{7}$$

$$\underline{h}_{\rho_k} = \left[\rho_k^{\hat{d}-1}, \rho_k^{\hat{d}-2}, \cdots, \rho_k, 1\right]^T \tag{8}$$

and $D(\underline{\alpha})_{ij} = \alpha_i\delta_{i-j}$. $\underline{f}_{\rho_k}$ and $\underline{h}_{\rho_k}$ respectively represent the k^{th} columns of $F(\underline{\rho})$ and $H(\underline{\rho})$. Just as the rank and number of unit-rank outer products in the SVD of $X^{(\hat{d})}$ corresponds to the number of non-zero transient components d, Eq. (6) likewise describes the number of transient components in terms of d unit-rank outer products while explicitly exposing the transient parameters. Although the decomposition in Eq. (6) has poor numerical properties as compared to the SVD, e.g., the left- and right-singular vectors form orthonormal systems while $F(\underline{\rho})$ and $H(\underline{\rho})$ are both Vandermonde and thus suffer the same conditioning issues as Φ, the singular values have little significance with respect to the amplitude coefficients. Consequently, SVD-based noise reduction techniques such as low-rank matrix approximation, alternating projections between low rank and Toeplitz spaces, and Tikhonov regularization perform poorly on transient data. Performance is additionally degraded with $\bar{X}^{(\hat{d})}$ defined analogously to $X^{(\hat{d})}$ with $\bar{x}_d[n]$ in place of $x_d[n]$ further justifying the restriction of Definition II-1 to finite data records.

III. The Discrete Transient Transform

In preparation of defining a spectral representation parameterized by decay rate, the following definition describes the parametric structure of a real exponential basis, henceforth denoted by ϕ, summarized by the decay parameter $\underline{\sigma} \in \mathbb{R}^N$.

Definition III-1. The Real Exponential Basis.
The real exponential basis is a set of N linearly independent signals $\phi = \{\underline{\phi}_k\}_{k=1}^N = \{\underline{\phi}_1, \cdots, \underline{\phi}_N\}$ with the geometric structure

$$\phi_k[n] = (\sigma_k)^n, \ 0 \leq n \leq N-1$$

or equivalently

$$\underline{\phi}_k = \left[1 \ \sigma_k \ \sigma_k^2 \ \ldots \ \sigma_k^{N-1}\right]^T$$

defined using an arbitrary set of real, non-zero decay rates $\underline{\sigma} = \{\sigma_k\}_{k=1}^N$ ordered such that $-1 < \sigma_1 < \cdots < \sigma_N < 1$.

The set transform of ϕ, i.e. generating the matrix with columns corresponding to the elements of ϕ, yields the Vandermonde matrix Φ in Eq. 3 for $d = N$. Despite the focus on real-valued decay rates in this paper, the development of the algorithms in this section avoid relying upon this restriction in such a way that it is straightforward to extend the algorithms for arbitrary distinct non-zero complex values in place of the decay rates, e.g. the the N^{th} roots of unity.

The real exponential basis ϕ, as stated in Definition III-1, is non-orthogonal under the inner product defined in Eq. (1). In order to perform a change of basis to ϕ, several approaches may be taken. For example, using the linear independence of the exponential basis elements allows an alternative inner product to be defined under which ϕ is an orthogonal system. Another option is to orthogonalize the real exponential basis using classical techniques, e.g., Gram-Schmidt or Householder orthogonalization, Givens rotations, etc. We may also generate the closest orthonormal basis to ϕ in a least squares sense using Inner Product Shaping. [12] Another alternative, which we proceed with in this paper, is to construct a dual basis which is biorthogonal under the inner product in Eq. (1).

The dual basis of ϕ consists of the set of signals $\psi = \{\underline{\psi}_k\}_{k=1}^N$ satisfying the biorthogonality conditions, i.e.,

$$\langle\underline{\phi}_i, \underline{\psi}_j\rangle = \delta_{i-j}, \quad \text{for } 1 \leq i,j \leq N. \tag{9}$$

Sections III-A-III-C present three algorithms for generating a dual exponential basis ψ satisfying Eq. (9). In particular, Section III-A and III-B present algorithms for generating the dual basis signals for a general basis while the algorithm proposed in Section III-C is

specific to generating the dual basis to ϕ. As mentioned previously, that algorithm is not limited to real-valued decay rates.

In the following definition we state the DTT and the inverse DTT as an invertible transform by which a general finite-length signal $x[n]$ may be represented using the real exponential basis ϕ. This signal representation admits filtering in the sense of decay-rate selective filtering. It is straightforward to show that this form of filtering does not correspond to the linear convolution of $x[n]$ with an impulse response, as would be the case for an orthogonal basis.

Definition III-2. The Discrete Transient Transform. [13]
The Discrete Transient Transform (DTT) $X[k]$ and Inverse DTT (IDTT) of a general length N sequence $x[n]$ are defined as

$$X[k] = \sum_{n=0}^{N-1} x[n]\,\psi_k[n], \quad 1 \le k \le N \tag{10}$$

$$x[n] = \sum_{k=1}^{N} X[k]\,\phi_k[n], \quad 0 \le n \le N-1 \tag{11}$$

where ϕ is a real exponential basis and ψ is the corresponding dual basis. Eqs. (10) *and* (11) *are the DTT analysis and synthesis equations, respectively.*

The computation of an inner product such as that appearing in Eq. (10) is often implemented by sampling the convolution of $x_d[n]$ with $\psi_k[-n]$ at $n = 0$. Any of the algorithms in the following three subsections may be used to generate the dual exponential basis ψ for use in the DTT analysis equation. The uniqueness and equivalence to linear functionals of a dual basis is well known [14], confirming that all three algorithms theoretically produce identical results though in practice the signals that result often differ due to their algorithmic organizations and numerical properties.

A. Dual Basis Generation - Direct Computation

For a general finite-dimensional basis the corresponding dual basis may always be generated by solving, either directly or iteratively, the system of simultaneous linear equations which describe the biorthogonality conditions as stated in Eq. (9). Specifically, the dual exponential basis signals under the inner product defined in Eq. (1) satisfy the system of equations

$$\Psi^T \Phi = I_N \tag{12}$$

where Φ is defined as in Eq. (3) with $d = N$ and where I_N is the identity operator on $\mathbb{R}^N$, and the k^{th} column of Ψ is $\underline{\psi}_k, 1 \le k \le N$.

We now emphasize two potential disadvantages to directly solving Eq. (12) for Ψ. First, if generic inversion routines including variations of Gaussian elimination or Krylov iterative methods are implemented then the dual basis signals must all be computed simultaneously, which is inefficient when only specific expansion coefficients are of interest. Second, the matrix Φ to be inverted is structurally Vandermonde, whose ill-conditioning was discussed following Eq. (4). The algorithms in Section III-B and III-C avoid the first disadvantage by generating the dual basis signals independently.

B. Dual Basis Generation - Orthogonal Projectors

The second algorithm presented for dual basis generation identifies the corresponding subspace of a particular dual signal $\underline{\psi}_j$, for a fixed index j, by first identifying any non-zero multiple of the dual signal of interest followed by appropriate normalization. Denote the unnormalized dual signal as $\underline{\varphi}_j$, i.e.

$$\psi_j[n] = \gamma_j \varphi_j[n], \quad \gamma_j \in \mathbb{R} \tag{13}$$

where γ_j is the normalization factor. Once the unnormalized dual signal is available, the normalization constant may be evaluated as

$$\gamma_j^{-1} = \langle \underline{\phi}_j, \underline{\varphi}_j \rangle. \tag{14}$$

Note that Eq. (9) implicitly states that the inner product in Eq. (14) equals zero if and only if $\underline{\varphi}_j = \underline{0}$, which is excluded by definition from the dual basis.

In order to generate $\underline{\varphi}_j$ for a particular index j, we decompose the vector space $\mathbb{R}^N$, spanned by the exponential basis $\{\underline{\phi}_k\}_{k=1}^N$, into the direct sum of two subspaces as

$$\mathbb{R}^N = \mathcal{H}_j \oplus \mathcal{H}_j^{\perp} \tag{15}$$

where the one-dimensional space $\mathcal{H}_j$ is the span of $\underline{\phi}_j$ and $\mathcal{H}_j^{\perp}$ is the orthogonal complement of $\mathcal{H}_j$. Note that Eq. (9) constrains the dual signal $\underline{\psi}_j$ to be orthogonal to span$(\{\underline{\phi}_k\}_{k \ne j})$. Consequently, by defining $\Phi_{\sim j}$ as any linear map from $\mathbb{R}^N$ onto $\mathcal{H}_j^{\perp}$, e.g.,

$$\Phi_{\sim j} = \begin{bmatrix} | & & | & | & & | \\ \underline{\phi}_1 & \cdots & \underline{\phi}_{j-1} & \underline{\phi}_{j+1} & \cdots & \underline{\phi}_N \\ | & & | & | & & | \end{bmatrix}, \tag{16}$$

then $\underline{\psi}_j \in \text{range}^{\perp}(\Phi_{\sim j})$. By the adjoint equivalence theorems, the dual signal $\underline{\psi}_j$ lies in the null space of the adjoint map, i.e., $\underline{\psi}_j \in \text{null}\left(\Phi_{\sim j}^*\right)$, where $\Phi_{\sim j}^*$ is the adjoint of $\Phi_{\sim j}$.

As an example of this approach, consider generating $\underline{\varphi}_j$ using an orthogonal projector. Specifically, select the signal $\underline{h} \in \mathbb{R}^N$ such that $\underline{h}$ is partially in the subspace null $\left(\Phi_{\sim j}^*\right)$, e.g., $\underline{\phi}_j$. Then subtracting the projection of $\underline{h}$ onto the space range$(\Phi_{\sim j})$ from itself yields $\underline{\varphi}_j$. In summary, for any $\underline{h} \in \mathbb{R}^N$ such that $\underline{h} \ne \mathcal{P}_{\mathcal{R}(\Phi_{\sim j})}(\underline{h})$

$$\underline{\varphi}_j = \underline{h} - \mathcal{P}_{\mathcal{R}(\Phi_{\sim j})}(\underline{h}) \tag{17}$$

where $\mathcal{P}_{\mathcal{R}\left(\Phi_{\sim j}\right)}$ is an orthogonal projector onto the space $\mathcal{R}\left(\Phi_{\sim j}\right)$. The canonic projector of this type, using the linear map in Eq. (16), is given by

$$\mathcal{P}_{\mathcal{R}\left(\Phi_{\sim j}\right)} = \Phi_{\sim j}\left(\Phi_{\sim j}^T \Phi_{\sim j}\right)^{-1} \Phi_{\sim j}^T. \tag{18}$$

C. Dual Exponential Basis Generation - Polynomial Expansion

The third and final algorithm presented in this paper for dual exponential basis generation makes use of the structural relationship between polynomials and Vandermonde systems in order to generate the unnormalized dual signal $\underline{\varphi}_j$ for a fixed index j. Explicitly writing the $N-1$ homogeneous biorthogonality constraints corresponding to $\underline{\psi}_j$, or equivalently writing $\underline{\psi}_j \in \text{null}\left(\Phi_{\sim j}^*\right)$ using the matrix $\Phi_{\sim j}$ in Eq. (16), yields the underdetermined system of equations

$$\sum_{n=0}^{N-1} \sigma_i^n \varphi_j[n] = 0, \quad i \ne j. \tag{19}$$

The structure of Eq. (19) implies that the values of $\varphi_j[n]$ must belong to the coefficients of a polynomial with zeros $\sigma_k, 1 \le k \le N, k \ne j$, to within an arbitrary non-zero scaling. Coupling this observation with the representation of the dual exponential signal as the impulse response of a system described as a factored z-transform we have a description of $\psi_j[n]$ as

$$\mathcal{Z}\{\psi_j[n]\} \propto \prod_{k=1,\, k \ne j}^{N} \left(1 - (\sigma_k z)^{-1}\right) \tag{20}$$

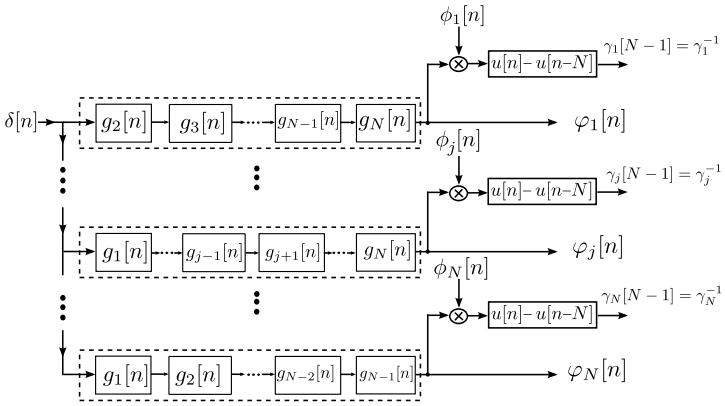

Fig. 1. The signal-flow graph representation of an algorithm which generates the dual real exponential basis signals ψ using the method described in Section III-C.

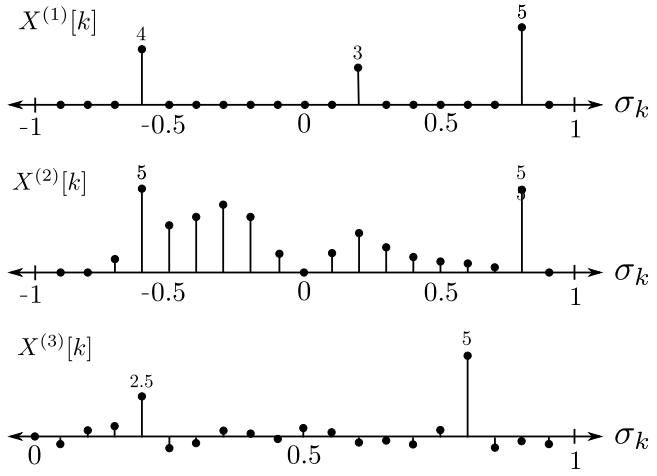

Fig. 2. The transient spectra generated by solving the linear system of equations in Eq. 23 for the three synthetic transient signals in Eqs. 24-26 using the algorithms presented in Sections III-A-III-C.

where $\mathcal{Z}\{\cdot\}$ represents the z-transform. Let the inverse z-transform of each multiplicand in Eq. (20) be denoted by $g_k[n]$, a length two temporal signal of the form

$$g_k[n] = \delta[n] - \sigma_k^{-1}\delta[n-1]. \tag{21}$$

Then it follows that the dual exponential signal is generated by $N-1$ successive linear convolutions, i.e.

$$\psi_j[n] \propto g_1[n] * \cdots * g_{j-1}[n] * g_{j+1}[n] * \cdots * g_N[n], \tag{22}$$

followed by appropriate scaling. A signal-flow graph representation for the system which generates both the unnormalized dual signals via polynomial expansion and the corresponding normalization constants is depicted in Figure 1. This algorithm, depending on the parameters $\underline{\sigma}$ and N, tends to have significant numerical advantages as compared to general inversion routines discussed in Section III-A. Since the components $g_k[n]$ are re-used in Fig. 1, computational recycling is obtainable via straightforward flowgraph manipulations or by using a recursive implementation with pole-zero cancellation.

IV. Example: Spectral Leakage and Resolution

In this section we illustrate, by means of computing the DTT or equivalently inverting a Vandermonde system with real nodes, issues of transient spectral resolution, leakage, and numerical instability. In particular, we consider solving the generally ill-conditioned linear system of equations

$$\underline{x}_d = \Phi\underline{\alpha} \tag{23}$$

using the algorithms presented in Sections III-A-III-C for three synthetic transient signals given by, for $0 \leq n \leq N-1$,

$$x_3^{(1)}[n] = 4(-0.6)^n + 3(0.2)^n + 5(0.8)^n \tag{24}$$

$$x_3^{(2)}[n] = 4(-0.62)^n + 3(0.2)^n + 5(0.8)^n \tag{25}$$

$$x_2^{(3)}[n] = 3(0.2)^n + 5(0.8)^n. \tag{26}$$

Figure 2 depicts the corresponding transient spectrum or Vandermonde expansion coefficients for each of these signals and for $N = 19$. For $X^{(1)}[k]$ and $X^{(2)}[k]$ the decay rates σ are equally spaced in the interval $[-0.9, 0.9]$ and for $X^{(3)}[k]$ the decay rates are equally spaced in the interval $[0.05, 0.95]$. The spectrum $X^{(1)}[k]$ was produced using the algorithm presented in Section III-C, the spectrum $X^{(2)}[k]$ was produced using the algorithm presented in Section III-B, and the spectrum $X^{(3)}[k]$ was produced using the algorithm presented in Section III-A. $X^{(1)}[k]$ and $X^{(3)}[k]$ depict scenarios for which $\{\rho_k\}_{k=1}^3 \subseteq \sigma$ and $\{\rho_k\}_{k=1}^2 \subseteq \sigma$, respectively, i.e., the expansion coefficients theoretically should contain no spectral leakage since the decay rates appear exactly in the nodes defining Φ. The spectrum $X^{(2)}[k]$ depicts the case for which $\{\rho_k\}_{k=1}^3 \not\subseteq \sigma$; the effect of which is manifested as a false spectral peak caused by spectral leakage appearing at $\sigma_7 = -0.3$ partially occluding the true peak at $\sigma_{12} = 0.2$. The spectral leakage in $X^{(3)}[k]$ is a consequence of the numerical instability of the algorithm used to generate this spectrum. Although not depicted in Figure 2, solving for the spectrum $X^{(3)}[k]$ using the algorithm in Section III-C on the same computational platform results in a spectrum containing no leakage, i.e. a spectrum consisting of two distinct peaks and zero-valued coefficients at all other decay rates.

References

[1] W. E. Boyce and R. C DiPrima, *Elementary Differential Equations and Boundary Value Problems*, Wiley, 5th edition, 1992.

[2] P. Markos and C. M. Soukoulis, *Wave Propagation: From Electrons to Photonic Crystals and Left-Handed Materials*, Princeton, 1st edition, 2008.

[3] J. Paris, *A framework for non-intrusive load monitoring and diagnostics*, Ph.D. thesis, MIT, 2006.

[4] R. Shankar, *Principles of Quantum Mechanics*, Springer, 2nd edition, 2008.

[5] A. Sedra and K. Smith, *Microelectronic Circuits*, Oxford, 6th edition, 2009.

[6] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, Prentice-Hall, 3rd edition, 2010.

[7] J. Makhoul, "Linear prediction: A tutorial review," *Proceedings of the IEEE*, vol. 63, no. 4, pp. 561–580, 1975.

[8] J. A. Cadzow, "Signal enhancement-a composite property mapping algorithm," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 36, no. 1, pp. 49–62, 1988.

[9] M. Vetterli P. Marziliano and T. Blu, "Sampling signals with finite rate of innovation," *Signal Processing, IEEE Transactions on*, vol. 50, no. 6, pp. 1417–1428, 2002.

[10] H. Armstrong, "On the representation of transients by series of orthogonal function," *IRE Transactions on Circuit Theory*, 1959.

[11] B. Beckermann, "The condition number of real vandermonde, krylov and positive definite hankel matrices," *Numerische Mathematik*, vol. 85, pp. 553–577, 1997.

[12] Y. C. Eldar, *Quantum Signal Processing*, Ph.D. thesis, MIT, 2001.

[13] T. A. Lahlou, "Parameter recovery for transient signals," M.S. thesis, MIT, 2013.

[14] S. Axler, *Linear Algebra Done Right*, Springer, 2nd edition, 2004.