Parameter Recovery for Transient Signals

by

Tarek A. Lahlou

B.S., Electrical Engineering, George Mason University (2011)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

Author . . . .
Department of Electrical Engineering and Computer Science
May 1, 2013

Certified by. . . . .
Alan V. Oppenheim
Ford Professor of Engineering
Thesis Supervisor

Accepted by . . . .
Leslie A. Kolodziejski
Chair, Department Committee on Graduate Students

Parameter Recovery for Transient Signals

by

Tarek A. Lahlou

Submitted to the Department of Electrical Engineering and Computer Science
on May 1, 2013, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

Transient signals naturally arise in numerous disciplines for which the decay rates and amplitudes carry some informational significance. Even when the decay rates are known, solving for the amplitudes results in an ill-conditioned formulation. Transient signals in the presence of noise are further complicated as the signal-to-noise ratio asymptotically decreases in time.

In this thesis the Discrete-Time Transient Transform and the Discrete Transient Transform are defined in order to represent a general signal using a linear combination of decaying exponential signals. A common approach to computing a change of basis is to make use of the dual basis. Two algorithms are proposed for generating a dual basis: the first algorithm is specific to a general exponential basis, e.g., real exponential or harmonically related complex exponential bases are special cases of the general exponential basis, while the second algorithm is usable for any general basis. Several properties of a transient domain representation are discussed. Algorithms for computing numerically stable approximate transient spectra are additionally proposed.

The inherent infinite bandwidth of a continuous-time transient signal motivates in part the development of a framework for recovering the decay rates and amplitudes of a discrete-time lowpass filtered transient signal. This framework takes advantage of existing parameter modeling, identification, and recovery techniques to determine the decay rates while an alternating projection method utilizing the Discrete Transient Transform determines the amplitudes.

Thesis Supervisor: Alan V. Oppenheim
Title: Ford Professor of Engineering

Acknowledgments

This thesis greatly benefited from the stimulating discussions, insightful contributions, unconventional creativity, and extraordinary patience of my thesis supervisor, Al Oppenheim. Al presented me with a rare, non-traditional opportunity to take an unconstrained intellectual journey, on which I would have, without a doubt, ended up lost without his mentorship. Al, over the past two years I know that I have gained more from working with you than I have yet to realize, both academically and personally; I sincerely look forward to continuing this development.

On the road to completing this thesis I have been fortunate enough to cross paths with many members of the DSPG family, each of whom have had a positive influence on my thesis. Specifically, I would like thank the following people for helping me to connect the dots when I was unable to untangle myself from the details, and for providing creative ideas and suggestions throughout this process: Tom Baran, Sefa Demirtas, Martin McCormick, Milutin Pajovic, and Guolong Su. Additionally, I'd like to thank Laura von Bosau, Petros Boufounos, Dan Dudgeon, Xue Feng, Yuantao Gu, Steve Isabelle, Charlie Rohrs, and Charles Sestok for their insightful contributions during group brainstorming sessions and otherwise.

I will forever owe a great debt to my parents, Jaye and Mohammed, for having provided me with the best opportunities they could, while never asking for anything in return. Jaye and Mohammed, you've instilled in me a belief in myself that has been, and forever will be, invaluable. You've always told me you'd support me in whatever I do, and more importantly, I know you will. As I continue to grow up I further appreciate all the time, effort, and money you selflessly spent in order to make me who I am today. Unfortunately, my gratitude cannot fit within these pages.

Growing up, and still to this day, I've had a first-class role model in my sister, Rita. Rita, when I think of either honesty, empathy, integrity, or compassion, I immediately think of you for the extraordinary example you've set for me and the world. One of my proudest accomplishments is being your little brother, a title I can only strive to live up to.

I have been blessed to have an extended family that has never allowed the distance between us to alter our friendships. Bassam Dourassi, Erik Flading, Phillip Mai, Ryan Marks, and Brandt Williams, thank you for always making my visits home memorable and keeping my up to date on your lives, and occasionally mine. Each time I made it home, you picked up our friendship where we left it off, as if I had never left.

In the past two years a few people have made MIT feel much like a home to me, and to them I owe a great debt of gratitude. They've each made contributions to this thesis in their own ways. Thanks to Danica Chang, Itthi Chatnuntawech, Farrah Sabouni, Erin (Dongying) Shen, Yi Sun, and Filiz Yetisir. Special thanks to Itthi Chatnuntawech for countless hours of studying, brainstorming, and just surviving the day-to-day. Special thanks to Yi Sun, for the always entertaining discussions, both academic and otherwise. Special thanks to Farrah Sabouni, for enduring, juicing, and always being up for an adventure.

There have been many educators who have both inspired and assisted me to take the path I am currently on; I'd like to specifically thank Gerald Cook, Joseph Hibey, Clark Lindgren, and Jill

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Exponential signals are often used to model relationships for which a constant change in the independent variable corresponds to a constant proportional change in the dependent variable. This property, and the fact that the derivative of an exponential signal is itself an exponential signal, are just two of the many reasons that this class of signals plays an important role in a broad range of disciplines. For example, exponential signals play a fundamental role in: the solutions to differential equations, the study of radioactivity in nuclear physics, the behavior of various electronic components, compounding interest in finance, tunneling in quantum physics, the behavior of atmospheric pressure in geophysics, and in a number of other places.

In the field of signal processing, exponential signals play a fundamental role in characterizing the behavior of signals and systems, e.g., Fourier and Laplace transforms, eigenfunctions of Linear and Time-Invariant (LTI) systems, etc. With respect to exponential-based transforms, exponential signals are used as a basis for representing a set of data, irrespective of the quantities represented by the data. Consistent with this methodology, an important contribution in this thesis is the derivation of a representation for a useful class of discrete-time signals based upon linear combinations of real, decaying exponential signals. After establishing this representation, important key properties and consequences of such a representation are identified and discussed.

Digital signal processing has enjoyed widespread use on both standard computers and special purpose hardware, e.g., application-specific integrated circuits, field-programmable gate arrays, etc. In order to use digital processing techniques on a continuous-time data source, the signal must be appropriately sampled. However, the process of sampling a transient signal inherently includes the effects of aliasing due to the infinite bandwidth of each exponential component. More commonly, it is often only the low frequency content of any signal that can be measured due to the lowpass characteristic of numerous physical sampling systems. In general, a signal is not guaranteed to be uniquely recovered from the output of a non-invertible system, e.g., an ideal lowpass filter (LPF), without further knowledge of the input signal prior to the system. In other words, the heavily attenuated portion of the frequency spectrum cannot be recovered in any meaningful sense. As a consequence, the process of sampling a transient signal cannot completely avoid aliasing effects. Therefore it is of particular interest in many fields that rely upon transient signals to convey infor-

mation to account for this effect. For this reason, an important aspect of this thesis deals with the recovery of a transient signal after it has been processed by an LPF. As we will show, by exploiting prior knowledge on the structure of transient signals, it becomes possible to recover transient signal parameters under fairly broad conditions relating the passband of the lowpass filter and the fastest decaying exponential component of the filtered transient signal.

The algorithms developed in this thesis are easily extendable to a more general structure of signals, specifically a linear combination of damped complex exponentials, all of the same frequency. Given a signal of this structure, the algorithms in this thesis are directly applicable by simply implementing a pre-processing stage where the signal to be analyzed is frequency-modulated to baseband, therefore becoming a real, transient signal.

1.1 Organization of Thesis

This thesis progressively discusses three stages of parameter recovery for transient signals in which each stage builds upon the previous. The remainder of this section outlines this progression as well as the organization of results.

Chapter 2 defines the structure of transient signals and motivates the need for the algorithms and techniques proposed in this thesis. In addition, detailed definitions of successful parameter recovery are presented for the three stages discussed below. A transient signal, in what follows, is completely characterized by its decay rates and amplitude coefficients, where each amplitude coefficient corresponds to an exponential component with a unique decay rate.

The first stage of this thesis considers determining the amplitude coefficients of a transient signal given both the samples of the signal and the decay rates present in the signal. Chapter 3 proposes two algorithms for determining the amplitude coefficients for this case, one of which avoids solving an ill-conditioned formulation using matrix operations, but instead exploits the structure inherent in exponential signals. Although the linear system of equations derived for the amplitude coefficients is uniquely invertible, the solutions found using the proposed algorithms result in fewer numerical errors.

The second stage considers determining the amplitude coefficients and decay rates of a transient signal given only the samples of the signal. (Definitions 1 and 3 in Table 2.2) Two approaches used to solve for the decay rates are considered in this thesis. The first approach is discussed in Chapter 3 where the transient transform is defined in order to use spectral analysis techniques, similar to Fourier spectral analysis, in order to determine which decay rates are present in the transient signal. Due to the numerical instability of implementing this transform for large data sets, several approximate transient spectral analysis techniques are proposed in Chapter 4. The second approach considered is to make use of well known parameter recovery algorithms, a survey of which is presented in Appendix B. The result of either of these approaches is the set of decay rates present in the transient signal. The algorithms presented for the first stage may then be used to determine the amplitude coefficients.

The third and final stage of this thesis considers determining the amplitude coefficients and decay

rates of a transient signal given samples of the transient signal after lowpass filtering. (Definitions 2 and 4 in Table 2.2) Chapter 5 presents a framework for determining the decay rates in this scenario based upon overdetermined parameter recovery. The recovery of the decay rates is shown to be possible under a broad range of cutoff frequencies of the LPF. This framework requires the number of exponential components present in the signal to be known *a priori*. To address this, an algorithm for determining the model order when it is unknown is proposed. Further, an alternating projection-based algorithm is also proposed which utilizes both the transient transform and the Fourier transform, i.e., representations using both real exponential and complex exponential bases, in order to determine the amplitude coefficients.

Chapter 6 provides a summary of the key results within this thesis as well as highlighting several potential directions for future research.

Chapter 2

Transient Structure, Spectra, and Parameter Recovery

This chapter defines the transient structure of signals as well as identifies the need for transient parameter recovery through example applications. The concept of a transient domain representation of a signal is also motivated by an example. Further, the full definitions of successful parameter recovery under various scenarios are defined.

2.1 Multi-Component Decaying Exponential Signal Structure

The structure of signals under consideration in this thesis is that of a linear combination of decaying, real exponentials with distinct decay rates. Specifically, the d^{th}-order continuous-time model is defined by

$$x_d(t) = \sum_{k=1}^{d} \alpha_k e^{-\lambda_k t}, \quad t \geq 0, \tag{2.1.1}$$

where $\lambda_k \neq \lambda_j$, for $j \neq k$, for a finite order d. The parameters α_k and λ_k are real, non-zero amplitude coefficients and positive decay rates, respectively, for $1 \leq k \leq d$. A time-series is defined as a sequence of uniformly spaced samples of a continuous-time model. For the remainder of this thesis, unless otherwise stated, the signal in Eq. (2.1.1) is assumed to be uniformly sampled, producing a time-series of the form

$$x_d[n] = \sum_{k=1}^{d} \alpha_k (\sigma_k)^n, \quad n \geq 0. \tag{2.1.2}$$

The shorthand notation $\{\alpha_k\}_{1:d}$ and $\{\sigma_k\}_{1:d}$ will be used to denote the set of amplitude coefficients, α_k, and the set of decay rates, σ_k, for $1 \leq k \leq d$, respectively. The decaying restriction of each exponential component implies that the value of each σ_k must lie within the open interval of $(0, 1)$. Eq. (2.1.2) is henceforth referred to as the transient signal structure.

Consider N samples of the transient signal $x_d[n]$ over a finite and possibly non-consecutive

interval of support, e.g., for $n = n_1, ..., n_N$, $n_i < n_j$ for $1 \leq i < j \leq N$. This observed set of transient data, using the notation in Eq. (2.1.2), is completely characterized using the structure of a generalized Vandermonde matrix and is given by

$$
\begin{bmatrix} x_d[n_1] \\ x_d[n_2] \\ \vdots \\ x_d[n_N] \end{bmatrix} = \begin{bmatrix} \sigma_1^{n_1} & \sigma_2^{n_1} & \cdots & \sigma_d^{n_1} \\ \sigma_1^{n_2} & \sigma_2^{n_2} & \cdots & \sigma_d^{n_2} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_1^{n_N} & \sigma_2^{n_N} & \cdots & \sigma_d^{n_N} \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_d \end{bmatrix}. \tag{2.1.3}
$$

In this formulation, the decay rates, $\{\sigma_k\}_{1:d}$, are seen to have a non-linear relationship with the observed sample values, while the amplitude coefficients, $\{\alpha_k\}_{1:d}$, have a linear relationship to the samples− given that the decay rates are fixed. An important special case of Eq. (2.1.3) is considered next.

A convenient and commonly appropriate interval of support is given by

$$
n_i = i - 1, \quad \text{for } 1 \leq i \leq N \tag{2.1.4}
$$

in which case Eq. (2.1.3) becomes

$$
\begin{bmatrix} x_d[0] \\ x_d[1] \\ \vdots \\ x_d[N-1] \end{bmatrix} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \sigma_1 & \sigma_2 & \cdots & \sigma_d \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_1^{N-1} & \sigma_2^{N-1} & \cdots & \sigma_d^{N-1} \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_d \end{bmatrix} \tag{2.1.5}
$$

or expressed in matrix notation as $x_d = V\left(\sigma^T\right)\alpha$, where $\sigma^T = [\sigma_1, \cdots, \sigma_d]$ is defined as the pole vector, $V\left(\sigma^T\right)$ is a Vandermonde structured matrix defined by $\left[V\left(\sigma^T\right)\right]_{ij} = \sigma_j^{i-1}$, $\alpha = [\alpha_1, \cdots, \alpha_d]^T$ is the amplitude vector, and $x_d = [x_d[0], \cdots, x_d[N-1]]^T$ is defined as the observation vector. For this special case, the structure of the decay rates is completely characterized by the geometric progression in each column of the Vandermonde matrix. Unless indicated otherwise, the interval of support in Eq. (2.1.4) is assumed for the remainder of this thesis.

When $N = d$ the determinant of the Vandermonde matrix in Eq. (2.1.5) is given by [25]

$$
\prod_{1 \leq i < j \leq d} (\sigma_j - \sigma_i), \tag{2.1.6}
$$

which can be seen by induction on d. Since the decay rates are distinct, the determinant is guaranteed to be non-zero and strictly less than one, which is equivalent to stating that the columns of the Vandermonde matrix are linearly independent. These facts will be exploited in various places throughout this thesis.

In order to represent the physical process of sampling a continuous-time signal, we consider a model which includes the presence of additive Gaussian noise to account for possible effects of the sampling hardware, e.g., jitter or thermal noise. The model of a sampled continuous-time d^{th}-order

multi-component real decaying exponential signal with additive noise is given by

$$\bar{x}_d[n] = x_d[n] + \eta[n], \tag{2.1.7}$$

where $\eta[n]$ is an additive white Gaussian noise process with constant finite spectral density.

2.2 Motivation and Need

The values taken by the amplitude coefficients and decay rates in a transient signal typically have some physical significance with respect to the origin of the signal. In these applications it is often of greater importance to accurately identify these parameters rather than, for example, model the observed data samples. This is especially true when the observed samples are known to contain additive noise and/or aliasing effects. In the following two examples, various disciplines, both scientific and otherwise, use a physical model or a mathematical description best represented by a sum of real, decaying exponentials. In the third example, the need for a parametric spectrum based upon decay rate is motivated.

2.2.1 The Smart Grid: Non-Intrusive Load Monitoring

A Non-Intrusive Load Monitoring (NILM) system provides the capability of recording and/or reacting to different electrical components being added or removed from an electrical grid. One approach for identifying which component is added to the system is to classify the initial transient signature of each component, resulting from the sudden change in impedance seen by the grid when a component is added or removed. This signature dies off in steady-state. Therefore proper identification of this transient signal is important to a NILM system. In the United States, the electric grid operates at 60 Hz, and therefore the recorded signal must be appropriately modulated to baseband in order to abstract the transient envelope from the data. The NILM systems described in [19] use this approach followed by a least-squares comparison to each template signal in order to identify which component was added or removed from the grid. An alternative to this procedure, using several algorithms proposed within this thesis, would be to store the exponential decay rates corresponding to each electrical component in place of the template signals. Then a transient signal would be analyzed to identify if each possible decay rate was present. The resulting amplitude coefficients would then signify which electrical component was added or removed from the electrical grid accordingly.

2.2.2 Economic Recession Intervention

A common measure of the economic growth of a country is the percentage change in that countries real Gross Domestic Product (GDP), i.e., a measure of the goods and services produced by that country within a fixed time period using a fixed price for the good or service. For example, the United States real GDP per capita has been approximately growing by two percent per year for over 200 years, implying an exponential rate of growth. Various methods for predicting a countries GDP during a fiscal quarter are often used to indicate whether market regulation is needed to prevent

major economic recession. This process requires the estimation of an exponential parameter in a noisy signal, motivating the recovery of decay rates when the signal is time-reversed.

2.2.3   Dual Tone Multi-Frequency

In what follows, as an illustration of an advantage to using a parametric spectrum based upon decay rate, we consider a comparison to the commonly used frequency spectrum using an example of when each spectrum is best utilized. In order to do so, we define a periodic signal and a transient signal, over the finite-duration interval $0 \leq n \leq N - 1$, by

$$v_1\,[n] = \alpha_a \cos\left(\omega_a n\right) + \alpha_b \cos\left(\omega_b n\right) \tag{2.2.1}$$

$$v_2\,[n] = \alpha_a \left(\sigma_a\right)^n + \alpha_b \left(\sigma_b\right)^n. \tag{2.2.2}$$

In many contexts, a received signal may convey information through identifying which signal, if any, out of a set of template signals was received. This may be accomplished through comparing distinguishable characteristics of the received signal to the characteristics of each of the possible template signals. For example, $v_2\,[n]$ has the structure of a template signal of a specific load in a NILM system for fixed values of σ_a, $\sigma_b \in (0, 1)$. In this scenario it is useful to define a parameter spectrum in which the calculated coefficients describe the amount of each of the specific decay rates present in the received signal. Henceforth the parameter spectrum based upon decay rate will be termed the transient spectrum, an example of which is shown in Figure 2.2.1(b). As a second example, $v_1\,[n]$ has the form of the template signals used in Dual-Tone Multi-Frequency (DTMF) signaling, the current industry standard for landline telecommunication service. [6, 22] This communication scheme uses the presence of a pair of frequencies to distinguish one template signal from another.

DTMF is a multi-frequency tone dialing system by which the push button keypad used for dialing in a landline telephone call, or in response to an in-call menu, conveys the number or keys dialed by the caller. With DTMF, each key pressed transmits the linear combination of two sinusoids comprised of predetermined frequencies; one from a high frequency group and one from a low frequency group. The high frequency group ranges in frequency from $1209 - 1633$ Hz and the low frequency group ranges from $697 - 941$ Hz. Two frequencies are used per transmitted signal to ensure that a human voice cannot imitate one of the template signals. The sinusoidal component from the high frequency group is slightly louder than the sinusoidal component from the the low frequency group to compensate for the high-frequency roll off of most communication channels. Table 2.1 shows the pairs of frequencies used in each of the template signals corresponding to the transmittable symbols.

	1209 Hz	1336 Hz	1477 Hz	1633 Hz
697 Hz	1	2	3	A
770 Hz	4	5	6	B
852 Hz	7	8	9	C
941 Hz	*	0	#	D

Table 2.1: Dual-tone multi-frequency signaling codebook

For DTMF, one method for identifying which symbol was transmitted is to check the received signal for the presence of each of the eight possible frequency components by computing the Discrete Time Fourier Transform (DTFT) coefficient at each frequency. Then, each of these coefficients is compared to a threshold to determine which pair of frequencies is present, or if no key was pressed. A plot of the DTFT coefficients with respect to these frequencies is termed the parameter spectrum for DTMF. An example is shown in Figure 2.2.1(a).

Many techniques are available to compute the DTFT coefficients needed to produce the parameter spectrum corresponding to DTMF. Among these techniques, with respect to efficiency in terms of the number of multiplications and additions, the Goertzel algorithm is often used. The Goertzel algorithm simultaneously computes the real and imaginary portions of a sample of the DTFT of a finite length signal using a second order recursion followed by a correction term. [10] Note that the Goertzel algorithm is often more efficient than using a Fast Fourier Transform (FFT) algorithm to compute the Discrete Fourier Transform (DFT) followed by discarding the non-relevant frequency bins. The Goertzel algorithm also has the advantage of not having to compute DTFT coefficients at evenly spaced points on the unit circle in the z-plane, but can be used instead for any harmonic. Finally, an understanding of additive noise in a communication channel is well understood in the frequency domain− the parameter spectrum based on frequency. This allows for symbol decoding to take the effects of noise into account.

Figure 2.2.1 displays the resulting parameter spectra of these two signals; the spectra in Figure 2.2.1(a) are for the sinusoidal signal $v_1[n]$ while the spectra in Figure 2.2.1(b) are for the transient signal $v_2[n]$. Each spectrum has N possible components present. Although both signals are completely represented using either spectrum, we see that in Figure 2.2.1(a) it is straightforward to identify which frequencies are present in $v_1[n]$, however, identifying the two decay rates in $v_2[n]$ is not as straightforward. Likewise, in Figure 2.2.1(b) it is straightforward to identify which decay rates are present in $v_2[n]$ while it is not straightforward to identify which frequencies are present in $v_1[n]$. Therefore, for selecting which decay rates are present in a template signal, e.g., $v_2[n]$, a transient spectrum is advantageous.

(a) Parametric spectra of $v_1[n]$

(b) Parametric spectra of $v_1[n]$
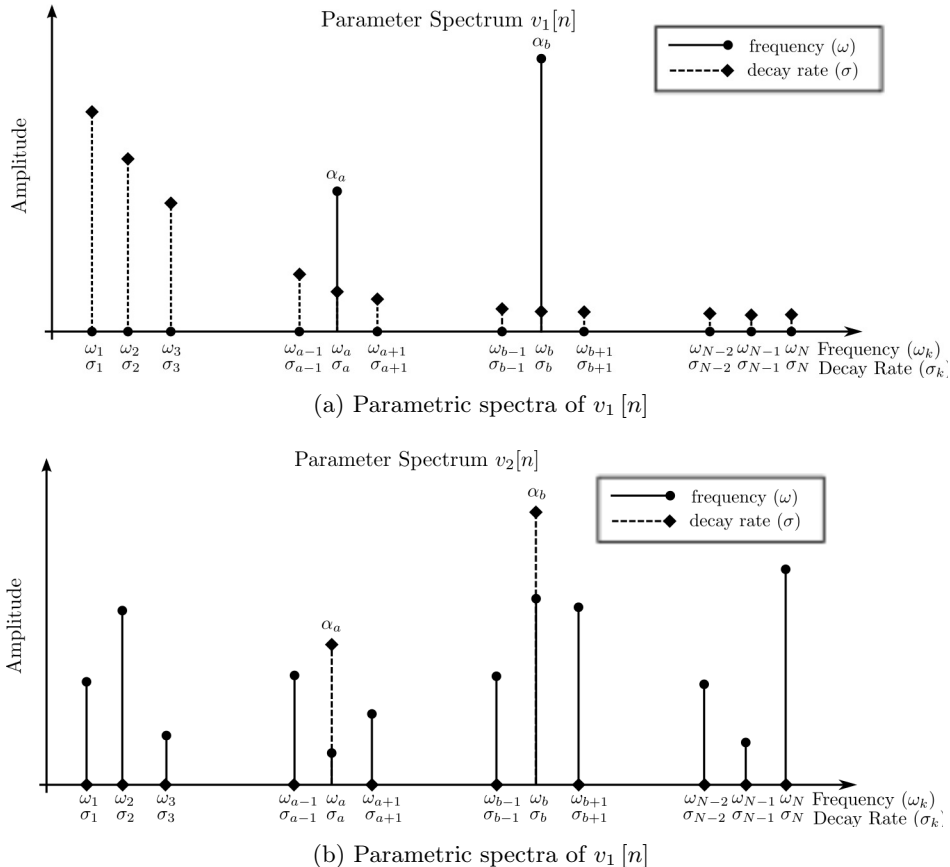
Figure 2.2.1: A comparison of parametric spectra based upon (a) frequency and (b) decay rate

2.3 Transient Spectra

Consider the scenario in which computing the amplitude coefficients is of interest when both samples from the signal in Eq. (2.1.1) and the decay rates, $\{\sigma_k\}_{1:d}$, are known. This is equivalent to computing the transient spectrum for the signal $x_d[n]$ as discussed in Section 2.2.3. Procedurally, this involves solving the system of equations described by Eq. (2.1.5) for α. In Section 3.7 we show that solving for α using a matrix inverse is a poorly conditioned approach, resulting in considerable error in the amplitude vector found. In Section 3.4, two algorithms are proposed that exploit the structure of these equations to produce the solution without requiring the direct inversion of the matrix $V\left(\sigma^T\right)$.

An inefficiency arises when a matrix inverse is used to to solve for α in Eq. (2.1.5) and only a subset of the amplitude coefficients corresponding to a subset of the decay rates is desired, suggesting that a technique which allows each amplitude coefficient to be found independently would be advantageous, rather than solving for the entire set and discarding the undesired coefficients. To address this inefficiency, the two algorithms proposed are each capable of solving for an amplitude coefficient corresponding to a specific decay rate independent of the other coefficients.

We next consider the scenario in which computing the amplitude coefficients is of interest when only samples from the signal in Eq. (2.1.2) are known. In this case, the transient spectrum may be

computed using any set of values for the decay rates. Peaks in the resulting transient spectrum provide an indication of which decay rates are present in the transient structured signal. Interpreting the results of a transient spectrum, similar to spectral analysis in the frequency domain, is discussed in Section 3.8.

2.4 Parameter Recovery

Distinguishing between a parameter identification algorithm and a parameter estimation algorithm is often quite subtle, depending on the context. In the context of this thesis, parameter estimation is defined to be parameter identification in the presence of noise, where a parameter identification algorithm processes data in order to identify the parameter values of a model which the data is known to fit. Therefore, a parameter estimation algorithm attempts to approximate the values of model parameters based upon measured or empirical data containing a random noise component, where the data is known to fit the model if the noise component was not present. Under these definitions, a specific algorithm is not classified as strictly an identification or estimation algorithm, but can be either depending on the relationship between the data and the model. In other cases, a set of samples, with or without a random component, does not accurately fit the model selected, but the model is chosen for other reasons. This thesis does not consider the determination of model parameters for this scenario.

The term "parameter recovery" is used in this thesis as an umbrella for which both parameter identification and parameter estimation, as defined above, fall beneath. Regardless of the algorithm or data used, the objective of a parameter recovery algorithm is to determine transient signal parameters by processing a set of observed samples. Appendix B provides a summary of several well known parameter recovery algorithms, which are later used in Chapter 5.

2.5 Parameter Identification

The successful recovery of the decay rates present in a transient signal is defined next when only the samples from a transient structured signal in Eq. (2.1.2) and the model order d are known. Because the observed samples are guaranteed to satisfy Eq. (2.1.2), this scenario is classified as parameter identification, as described in Section 2.4.

Definition 1. Successful Recovery in the Identification Case: *Given a set of observed samples that exactly fit the model in Eq. (2.1.2), successful recovery corresponds to accurately determining the full set of decay rates, $\{\sigma_k\}_{1:d}$, that were used to generate the observed samples, $\{x_d[n]\}_{0:N-1}$. Equivalently, accurate recovery of the pole vector σ.*

Several existing methods, summarized in Appendix B, achieve successful recovery of the decay

rates as defined above.

The next scenario considered is an extension of Definition 1 where the samples available are from the result of lowpass filtering the transient structured signal. In order to simplify the upcoming definition, define the signal $g_d[n]$ as the output of an LPF when the input is the signal $x_d[n]$, and let $\{g_d[n]\}_{0:N-1}$ be the set of N observed samples for $n = 0, ..., N-1$.

Definition 2. Successful Recovery in the Identification Case Post-Lowpass Filtering:
Given the cutoff frequency of the LPF and the set of observed samples, $\{g_d[n]\}_{0:N-1}$, corresponding to the output of the LPF when the input exactly fits the model in Eq (2.1.2), successful recovery corresponds to accurately determining the full set of decay rates, $\{\sigma_k\}_{1:d}$, that were used to generate the samples prior to the LPF. Equivalently, accurate recovery of the pole vector σ.

One approach for achieving successful recovery, as stated in Definition 2, is to perform a two step procedure of signal recovery followed by parameter identification. The first step is to recover the signal $x_d[n]$ prior to the LPF. The second step is then exactly the parameter recovery formulation in Definition 1. This approach inherently requires the extrapolation of frequency content from a non-bandlimited signal and is considered in Section 5.5.

Another approach for achieving successful recovery is to use the observed samples, $\{g_d[n]\}_{0:N-1}$, to directly solve for the desired parameters. By defining a parametric model $\hat{x}_d[n]$, e.g.,

$$\hat{x}_d[n] = \sum_{k=1}^{d} \hat{\alpha}_k(\hat{\sigma}_k)^n, \quad 0 \leq n \leq N-1, \tag{2.5.1}$$

where the parameters $\{\hat{\alpha}_k\}_{1:d}$ and $\{\hat{\sigma}_k\}_{1:d}$ are to be determined, we can formulate the second and more direct approach as an optimization problem. Doing so yields

$$(\hat{\alpha}, \hat{\sigma}) \in \arg\min_{\alpha,\sigma} \|g_d - \hat{x}_d * f_{lp}\|_2^2 \tag{2.5.2}$$

where $*$ denotes linear convolution and $f_{lp}[n]$ represents an ideal LPF. This formulation determines parameters that produce a residual signal with minimum ℓ_2-norm, where the residual signal is defined to be

$$r[n] = g_d[n] - \hat{x}_d[n] * f_{lp}[n]. \tag{2.5.3}$$

The residual signal, depending on the cutoff frequency of the LPF and the form of the parametric model $\hat{x}_d[n]$, can be either the sequence of modeling error or the sequence of linear prediction error. [15, 18] These two interpretations and their relation to one another are further discussed in Appendix B. The following two subsections formulate direct parameter recovery in two ways: first

in the time domain and then in the frequency domain. From the frequency domain formulation we conclude that the parameters that minimize the ℓ_2-norm of the residual signal are unique. A proof of this uniqueness is presented in Appendix A.

2.5.1 Time Domain Formulation

Denote the cutoff frequency of the ideal lowpass filter, $f_{lp}[n]$, as ω_c, $0 < \omega_c < \pi$. Directly solving for the parameters of a transient structured signal, as in Eq. (2.1.2), yields the non-linear set of equations

$$
\begin{aligned}
g_d[n] &= f_{lp}[n] * x_d[n] & (2.5.4) \\
&= \frac{\sin(\omega_c n)}{\pi n} * \sum_{k=1}^{d} \alpha_k (\sigma_k)^n u[n] & (2.5.5) \\
&= \frac{1}{\pi}\left[\alpha_1 \sigma_1^n \left(\sum_{l=-\infty}^{n} \frac{\sin(\omega_c l)}{l}\sigma_1^{-l}\right) + \cdots + \alpha_d \sigma_d^n \left(\sum_{l=-\infty}^{n} \frac{\sin(\omega_c l)}{l}\sigma_d^{-l}\right)\right]. & (2.5.6)
\end{aligned}
$$

For each value of n, $g_d[n]$ is an equation of the $2d$ parameters to be determined. Note that for a fixed set of decay rates, the amplitude coefficients still maintain a linear relationship to the lowpass filtered sample values.

2.5.2 Frequency Domain Formulation

An equivalent formulation of this problem is to consider the DTFT of each signal involved. Define the DTFTs of $x_d[n]$, $f_{lp}[n]$, and $g_d[n]$ as $X_d(e^{j\omega})$, $F_{lp}(e^{j\omega})$, and $G_d(e^{j\omega})$, respectively. This optimization problem is then formulated by selecting the parameters of $X_d(e^{j\omega})$ that minimize the energy in the DTFT of the residual, $G_d\left(e^{j\omega}\right) - X_d\left(e^{j\omega}\right) F_{lp}\left(e^{j\omega}\right)$. The parameters of the signal $x_d[n]$ are related to $X_d(e^{j\omega})$ by,

$$
X_d(e^{j\omega}) = \frac{\displaystyle\sum_{r=1}^{d}\alpha_r \prod_{l=1, l\neq r}^{d}\left(1 - \sigma_l e^{-j\omega}\right)}{\displaystyle\prod_{k=1}^{d}\left(1 - \sigma_k e^{-j\omega}\right)}. \tag{2.5.7}
$$

Writing this as an explicit optimization problem yields

$$
(\hat{\alpha},\ \hat{\sigma}) \in \arg\min_{\alpha,\sigma}\left\|G_d\left(e^{j\omega}\right) - X_d\left(e^{j\omega}\right) F_{lp}\left(e^{j\omega}\right)\right\|_2^2. \tag{2.5.8}
$$

2.6 Parameter Estimation

The successful recovery of the decay rates present in a transient signal is defined next when only

noisy samples of the transient structured signal in Eq. (2.1.2) and the model order d are known. Because the observed samples are guaranteed to satisfy Eq. (2.1.7), this scenario is classified as parameter estimation, as described in Section 2.4.

Definition 3. Successful Recovery in the Estimation Case: *Given the set of observed samples* $\{\bar{x}_d [n]\}_{0:N-1}$, *corresponding to samples which fit the model in Eq. (2.1.7), successful recovery corresponds to determining the full set of decay rates,* $\{\sigma_k\}_{1:d}$, *that produce a recovered pole vector,* $\hat{\sigma}$, *that is closest to the true pole vector,* σ, *for a given distance metric.*

This parameter estimation formulation motivates the discussion of robustness for each of the algorithms presented in Appendix B.

The final scenario considered is an extension of Definition 3 where noise is added to the result of a transient structured signal processed by an LPF. Denote $\bar{g}_d [n]$ as the output of the lowpass filter with the inclusion of additive white Gaussian noise, i.e., $\bar{g}_d [n] = g_d [n] + \eta [n]$ and $\{\bar{g}_d [n]\}_{0:N-1}$ as the set of N observed samples for $n = 0, ..., N - 1$.

Definition 4. Successful Recovery in the Estimation Case Post-Lowpass Filtering: *Given the cutoff frequency of the LPF and the set of observed samples* $\{\bar{g}_d [n]\}_{0:N-1}$, *corresponding to the sum of additive Gaussian noise and the output of the LPF when the input exactly fits the model in Eq (2.1.2), successful recovery corresponds to determining the full set of decay rates,* $\{\sigma_k\}_{1:d}$, *that produce a recovered pole vector,* $\hat{\sigma}$, *that is closest to the true pole vector,* σ, *for a given distance metric.*

Both Definitions 3 and 4 use a distance metric in order to measure the distance from a recovered pole vector to the true pole vector. Two examples of possible distance or error metrics include the Chebyshev distance and the ℓ_p-norm.

Table 2.2 summarizes the four definitions pertaining to successful recovery of the decay rates, $\{\sigma_k\}_{1:d}$. For each definition, the structure of the data assumed is stated for the interval $0 \leq n \leq N - 1$. Figure 2.6.1 depicts the generation of the available data for the different cases where the signals are also observed over the same interval of support.

Definition	Data
Definition 1: Successful Recovery in the Identification Case	$x_d[n]$
Definition 2: Successful Recovery in the Identification Case Post-Lowpass Filtering	$g_d[n] = f_{lp}[n] * x_d[n]$
Definition 3: Successful Recovery in the Estimation Case	$\bar{x}_d[n] = x_d[n] + \eta[n]$
Definition 4: Successful Recovery in the Estimation Case Post-Lowpass Filtering	$\bar{g}_d[n] = g_d[n] + \eta[n]$

Table 2.2: Summary of the definitions for successful transient parameter recovery

Figure 2.6.1: Summary of the formulations of transient parameter recovery

Chapter 3

Transient Spectral Analysis

This chapter formally defines the transient transform in order to represent a signal using a basis comprised of real exponential signals. For finite-length signals, several algorithms for converting to this representation are derived. Further, several properties of the algorithms for generating the transient spectrum as well as properties of the transient spectrum itself are discussed. In order to establish notation, a brief review of orthogonal and non-orthogonal change of bases is first presented.

3.1 Background on Basis Expansions

Consider representing a general discrete-time signal using a basis expansion of the form

$$x[n] = \sum_k \alpha_k \phi_k[n].$$

(3.1.1)

The set of signals $\{\phi_k[n]\}_{-\infty:\infty}$ is a countable set of pre-specified basis signals while the scalar coefficients $\{\alpha_k\}_{-\infty:\infty}$ are the expansion coefficients with respect to the corresponding basis signals. The representation of a signal by its expansion coefficients with respect to any basis will be referred to as a spectrum. For example, consider the discrete sifting equation, i.e.,

$$x[n] = \sum_k x[k] \delta[n-k].$$

(3.1.2)

Here the standard sampling basis is used, i.e., $\phi_k[n] = \delta[n-k]$, and the resulting expansion coefficients are the observed data values. In this example, note that the expansion coefficients, with respect to the standard sampling basis, provide a complete characterization of the signal $x[n]$. As a second example, the representation of a signal by its DTFT completely characterizes a signal with respect to a complex exponential basis.

The following two subsections describe a change of basis for the cases of an orthogonal and non-orthogonal basis.

3.1.1 Orthogonal Basis Expansions

For any two signals $u[n]$ and $v[n]$ in a vector space $\mathcal{V}$ defined over the field $\mathcal{F}$, the standard

inner product, $\langle \cdot, \cdot \rangle : \mathcal{V} \times \mathcal{V} \to \mathcal{F}$, is defined by

$$\langle u, v \rangle = \sum_n u[n] \, v^*[n] = v^H u, \tag{3.1.3}$$

where v^* and v^H denote the complex conjugate and complex conjugate transpose of v, respectively. Note that the standard inner product is a sample of the cross-correlation between the two signals, which is often implemented using linear convolution. When the basis signals $\{\phi_k[n]\}_{-\infty:\infty}$ are chosen such that they are orthogonal, i.e., $\langle \phi_k, \phi_l \rangle = 0$ for $k \neq l$, the resulting transformation is an orthogonal transform. When the basis signals are additionally normalized such that the basis signals also satisfy the set of constraints

$$\langle \phi_k, \, \phi_l \rangle = \delta_{kl} = \begin{cases} 1, & \text{if } k = l \\ 0, & \text{otherwise} \end{cases}, \tag{3.1.4}$$

then the basis expansion is an orthonormal basis expansion. For either the orthogonal or the orthonormal case, the expansion coefficients may be directly computed using the standard inner product, as defined in Eq. (3.1.3). This allows the signal $x[n]$ to be decomposed by

$$x[n] = \sum_k \frac{\langle x, \, \phi_k \rangle}{\langle \phi_k, \, \phi_k \rangle} \phi_k[n]. \tag{3.1.5}$$

The second example of an orthogonal basis expansion given above was the DTFT where the basis signals are complex exponentials. In Section 3.9 the algorithms proposed in this chapter, for finite-length signals, will be seen to hold for a complex exponential basis. Additional examples of orthogonal basis expansions include the Discrete Cosine Transform and the Hadamard Transform. [14]

3.1.2   Biorthogonal Basis Expansions

When the basis signals selected are not orthogonal under the standard inner product, then Eq. (3.1.5) cannot be used to represent $x[n]$. Several alternative approaches are discussed in the following section. One approach, which is taken in this thesis, is to generate a different set of signals with a desirable inner product structure with respect to the original set of basis signals in order to determine the desired expansion coefficients. That is, given a non-orthogonal basis $\{\phi_k[n]\}_{-\infty:\infty}$, define the set of signals $\{\psi_k[n]\}_{-\infty:\infty}$ such that they satisfy the set of constraints

$$\langle \phi_k, \, \psi_l \rangle = \sum_n \phi_k[n] \, \psi_l^*[n] = \delta_{kl}. \tag{3.1.6}$$

These constraints will be collectively referred to as the biorthogonality constraints. The set of signals $\{\psi_k[n]\}_{-\infty:\infty}$ is the dual basis of $\{\phi_k[n]\}_{-\infty:\infty}$ and each signal is a dual signal. Given $\{\psi_k[n]\}_{-\infty:\infty}$, the expansion coefficient α_l, for a fixed index l, is determined by performing the standard inner product of $x[n]$ with $\psi_l[n]$, resulting in the desired coefficient, i.e.,

$$\langle x, \, \psi_l \rangle = \left\langle \sum_k \alpha_k \phi_k, \, \psi_l \right\rangle \tag{3.1.7}$$

$$= \cdots + \alpha_{l-1}\langle \phi_{l-1}, \, \psi_l \rangle + \alpha_l \langle \phi_l, \, \psi_l \rangle + \alpha_{l+1}\langle \phi_{l+1}, \, \psi_l \rangle + \cdots \tag{3.1.8}$$

$$= \alpha_l. \tag{3.1.9}$$

The second equality comes from the linearity of an inner product in the first argument, and the third equality comes from the biorthogonality constraints. Therefore, the signal $x\,[n]$ may be represented by

$$x\,[n] = \sum_k \langle x, \, \psi_k \rangle \, \phi_k\,[n]. \tag{3.1.10}$$

The dual signal corresponding to an orthogonal basis signal is easily obtained from the orthogonal basis signal, e.g., $\psi_k \propto \phi_k^*$. This result implies that if the dual basis can be obtained, then a non-orthogonal basis expansion may be computed using a similar procedure to the orthonormal case.

If we choose the set of linearly independent, non-orthogonal signals $\{\phi_k\,[n]\}_{-\infty:\infty}$ as the basis for a vector space $\mathcal{V}$, then the vector space $\mathcal{V}$ may be decomposed as

$$\mathcal{V} = \mathcal{V}_l \oplus \mathcal{V}_{\sim l}$$

where $\oplus$ is a direct sum, $\mathcal{V}_l = \text{span}\,\{\phi_l\,[n]\}$ and $\mathcal{V}_{\sim l} = \text{span}\,\big\{\{\phi_k\,[n]\}_{-\infty:\infty} \setminus \{\phi_l\,[n]\}\big\}$ where $\setminus$ denotes set subtraction. From the definition of the biorthogonality constraints we have that the dual signal $\psi_l\,[n]$, for a fixed index l, is orthogonal to the space $\mathcal{V}_{\sim l}$, i.e., if we denote $\Phi_{\sim l}$ to be any onto linear map from $\mathcal{V}$ to $\mathcal{V}_{\sim l}$, then $\psi_l\,[n] \in \mathcal{R}^\perp\,(\Phi_{\sim l})$, where $\mathcal{R}\,(\cdot)$ is the range of a linear map. Using the adjoint equivalence identities, the dual signal $\psi_l\,[n]$ equivalently lies in the nullspace of the adjoint of $\Phi_{\sim l}$, i.e., $\psi_l\,[n] \in \mathcal{N}\,(\Phi_{\sim l}^*)$, where $\mathcal{N}\,(\cdot)$ is the nullspace of a linear map and $\Phi_{\sim l}^*$ is the adjoint of $\Phi_{\sim l}$. [24]

3.2 The Exponential Bases

This section defines the general exponential basis, using which we define two special cases: the complex exponential basis and the real exponential basis. The orthogonality of each of these two special cases of the general exponential basis is then discussed.

Definition. The General Exponential Basis. *The general exponential basis is defined as the set of signals denoted by*

$$\{\phi\,(\sigma, \, n]\}_{\sigma \in \mathcal{I}}, \tag{3.2.1}$$

where $\mathcal{I}$ is either a continuous contour in the complex plane or a finite set of distinct complex numbers, and, for each value of σ, $\phi(\sigma, n]$ is a general exponential signal of the form

$$\phi(\sigma, n] = (\sigma)^n, \quad n \geq 0, \quad \sigma \in \mathbb{C}. \tag{3.2.2}$$

Consider the general exponential basis $\{\phi(\sigma, n]\}_{\sigma \in \mathcal{I}}$ when $\mathcal{I}$ is a continuous contour in the complex plane. For any finite value of N, the first N non-zero samples of any N basis signals with distinct values of $\sigma \in \mathcal{I}$ form a valid basis for $\mathbb{C}^N$. This means that a finite general exponential basis may be constructed using values of σ_k that are samples along a contour in the complex plane. Specifically, define the set of signals $\{\phi_k[n]\}_{1:N}$ such that

$$\{\phi_k[n]\}_{1:N} = \{\phi_1[n], \ldots, \phi_N[n]\}, \quad \phi_k[n] = \begin{bmatrix} 1 \\ \sigma_k \\ \sigma_k^2 \\ \vdots \\ \sigma_k^{N-1} \end{bmatrix}, \quad 1 \leq k \leq N, \quad \sigma_k \neq \sigma_j \text{ for } k \neq j. \tag{3.2.3}$$

Define the matrix Φ, where each column corresponds to a general exponential signal, to have the structure

$$\Phi = \begin{bmatrix} | & & | \\ \phi_1[n] & \cdots & \phi_N[n] \\ | & & | \end{bmatrix}. \tag{3.2.4}$$

The matrix Φ takes the form of a Vandermonde matrix with $[\Phi]_{ij} = (\sigma_j)^{i-1}$, hence the columns of Φ are linearly independent. This implies that the general exponential signals $\{\phi_k[n]\}_{1:N}$ are also linearly independent, and since $\mathbb{C}^N$ is finite dimensional, this is sufficient to show that $\{\phi_k[n]\}_{1:N}$ forms a basis of $\mathbb{C}^N$.

Unless otherwise stated, the notation ϕ is used to denote an exponential signal for the remainder of this thesis. The first special case of the general exponential basis considered is the complex exponential basis, which is the basis used in Fourier analysis.

Definition. The Complex Exponential Basis. *The complex exponential basis is defined to be structurally identical to the general exponential basis where the interval $\mathcal{I}$ is chosen to be*

$$\mathcal{I}_{inf} = \left\{ \sigma \mid Re^2(\sigma) + Im^2(\sigma) = 1 \right\} \tag{3.2.5}$$

for the infinite complex exponential basis, or

$$\mathcal{I}_{fin} = \left\{ \sigma_k \mid \sigma_k = e^{j\frac{2\pi}{N}k}, \ k = 0, \cdots, N - 1 \right\} \qquad (3.2.6)$$

for the finite complex exponential basis.

Under the definition of the standard inner product in Eq. (3.1.3), it is straightforward to verify that both the finite and infinite complex exponential bases are orthogonal.

The second special case of the general exponential basis considered is the real exponential basis, which is the basis used in the transient domain representation of a signal. For brevity, this basis will often be referred to as simply the exponential basis.

Definition. The (Real) Exponential Basis. *The (real) exponential basis is defined to be structurally identical to the general exponential basis where the interval $\mathcal{I}$ is chosen to be*

$$\mathcal{I}_{inf} = \{ \sigma \mid \sigma \in (0, 1) \} \qquad (3.2.7)$$

for the infinite (real) exponential basis, or

$$\mathcal{I}_{fin} = \{ \sigma_k \mid \sigma_k \in (0, 1), \ k = 0, \cdots, N - 1, \ \sigma_k < \sigma_j \ for \ k < j \} \qquad (3.2.8)$$

for the finite (real) exponential basis.

We denote the transient spectrum of a signal by $\mathcal{A}(\sigma)$ or $\mathcal{A}[k]$ depending on whether the basis $\{\phi(\sigma, n]\}_{\sigma \in (0,1)}$ or $\{\phi_k[n]\}_{1:N}$ is used, respectively. Consider the standard inner product $\langle \cdot, \cdot \rangle$ given in Eq. (3.1.3), for which it immediately follows that

$$\langle \phi(\sigma_k), \phi(\sigma_l) \rangle \neq \delta_{\sigma_k \sigma_l}, \quad \text{for any } \phi(\sigma_k), \phi(\sigma_l) \in \{\phi(\sigma, n]\}_{\sigma \in (0,1)}, \qquad (3.2.9)$$

and

$$\langle \phi_k, \phi_l \rangle \neq \delta_{kl}, \quad \text{for any } \phi_k, \phi_l \in \{\phi_k[n]\}_{1:N}. \qquad (3.2.10)$$

Consequently, in order to compute the transient spectrum, i.e., $\mathcal{A}(\sigma)$ or $\mathcal{A}[k]$, several alternative approaches may then be taken. For example, the exponential basis may be orthogonalized through the Gram-Schmidt procedure, or we may generate the closest orthogonal basis to the exponential basis in a least squares sense using Inner Product Shaping. [7] These two approaches are considered in Section 4.6. Alternatively, we may define a new inner product under which the basis is orthogonal. The existence of such an inner product is guaranteed because for any set of linearly independent signals an inner product always exists under which the signals are orthogonal. This thesis takes the

route of deriving the dual basis signals under the standard inner product. However, the possibility of a finite exponential basis being orthogonal under a weighted standard inner product is ruled out next.

For any two signals $u[n]$ and $v[n]$ in a vector space $\mathcal{V}$ defined over the field $\mathcal{F}$, the weighted standard inner product, $\langle \cdot, \cdot \rangle_w : \mathcal{V} \times \mathcal{V} \to \mathcal{F}$, is defined by

$$\langle u, v \rangle_w = \sum_n u[n] v^*[n] w[n], \tag{3.2.11}$$

where $w[n]$ is a non-negative weighting function. For $\mathcal{V} = \mathbb{R}^N$, an equivalent formulation of the weighted inner product is given by, for $u[n], v[n] \in \mathbb{R}^N$,

$$\langle u, v \rangle_w = v^H W u \qquad \text{where} \qquad W = \begin{bmatrix} w[0] & 0 & \cdots & 0 \\ 0 & w[1] & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & w[N-1] \end{bmatrix}. \tag{3.2.12}$$

We are interested in determining if a non-negative weighting function $w[n]$, or a positive semi-definite, diagonal matrix W, exists such that $\phi_k[n]$ and $\phi_l[n]$ are orthogonal for $k \neq l$.

We denote G_w as the matrix of weighted inner product constraints, meaning $[G_w]_{ij} = \langle \phi_i, \phi_j \rangle_w$. The desired orthogonality constraints are given in matrix notation as

$$\Phi^T W \Phi = G_w \tag{3.2.13}$$

where G_w is a diagonal matrix with strictly positive elements on the principal diagonal. The following proposition shows that under this definition of a weighted inner product, there exists no non-negative weighting function $w[n]$ such that the elements of $\{\phi_k[n]\}_{1:N}$ are orthogonal. This is equivalent to saying there is no matrix W which is both diagonal and satisfies Eq. (3.2.13).

Proposition. *Given a real exponential basis* $\{\phi_k[n]\}_{1:N}$, *then there exists no non-negative weighting function* $w[n]$ *such that*

$$\langle \phi_k, \phi_l \rangle_w = \begin{cases} g_k, & k = l \\ 0, & k \neq l \end{cases} \tag{3.2.14}$$

for any strictly positive scalars $\{g_k\}_{1:N}$.

Proof. The weighted inner product for two elements of an exponential basis reduces to

$$\langle \phi_k, \phi_l \rangle_w = \sum_{n=0}^{N-1} (\sigma_k \sigma_l)^n w[n]. \tag{3.2.15}$$

The terms in the summand resulting from the exponential signals are strictly positive. This implies that the only non-negative weighting function that can be used in order to sum strictly positive terms and result in 0 is $w[n] = 0$.

$\square$

Therefore this proposition shows that there exists no weighted standard inner product, including the standard inner product, for which the exponential basis signals are orthogonal.

3.3 Transient Transforms

In this section we formally define the Discrete-Time Transient Transform (DTTT)−a reversible mapping through which a causal, transient signal is represented by a linear combination of exponential signals.

Definition. The Discrete-Time Transient Transform.

$$\mathcal{A}(\sigma) = \sum_n x_d[n]\,\psi(\sigma, n], \quad 0 < \sigma < 1, \tag{3.3.1}$$

$$x_d[n] = \int_0^1 \mathcal{A}(\sigma)\,\phi(\sigma, n]\,d\sigma, \quad n \geq 0, \tag{3.3.2}$$

where $\{\phi(\sigma, n]\}_{\sigma \in (0,1)}$ is the real exponential basis and $\{\psi(\sigma, n]\}_{\sigma \in (0,1)}$ is the corresponding dual basis. Furthermore, we refer to Eq. (3.3.1) as the DTTT analysis equation and Eq. (3.3.2) as the DTTT synthesis equation.

No known method for generating a dual exponential signal $\psi(\sigma, n]$ exists for any value of $\sigma \in (0, 1)$, and it appears that without access to the dual basis the DTTT cannot be used for analysis. However, for any finite-duration signal the dual exponential basis may be found. Thus we define the Discrete Transient Transform (DTT)−a reversible mapping through which a general finite-length signal is represented by a linear combination of exponential signals. Algorithms for generating the dual finite exponential basis are derived in Section 3.4. The DTT is a set of samples of the DTTT where the spacing of these spectral samples is determined by the decay rates selected in the finite exponential basis used, i.e., $\{\sigma_k\}_{1:N}$. The invertibility of the DTT follows by substituting the analysis equation into the synthesis equation and simplifying.

Definition. The Discrete Transient Transform.

$$\mathcal{A}[k] = \sum_{n=0}^{N-1} x[n]\,\psi_k[n]\,, \quad 1 \leq k \leq N, \tag{3.3.3}$$

$$x[n] = \sum_{k=1}^{N} \mathcal{A}[k]\,\phi_k[n]\,, \quad 0 \leq n \leq N-1, \tag{3.3.4}$$

where $\{\phi_k[n]\}_{1:N}$ is the real exponential basis and $\{\psi_k[n]\}_{1:N}$ is the corresponding dual basis. Furthermore, we refer to Eq. (3.3.3) as the DTT analysis equation and Eq. (3.3.4) as the DTT synthesis equation.

The DTT, as defined, possesses both homogeneity and additivity with respect to the input. To see this, denote the DTT pairs $v_1[n] \leftrightarrow \mathcal{V}_1[k]$ and $v_2[n] \leftrightarrow \mathcal{V}_2[k]$, corresponding to signals of length N; if $x[n] = a \cdot v_1[n] + b \cdot v_2[n]$ then $\mathcal{A}[k] = a \cdot \mathcal{V}_1[k] + b \cdot \mathcal{V}_2[k]$, for any real scalars a and b. Therefore this transformation is linear.

Given the dual signals $\{\psi_k[n]\}_{1:N}$, a transient analysis filter bank may be constructed to compute the transient spectrum $\mathcal{A}[k]$ according to Eq. (3.3.3). This process is identical to the filter bank implementation of the DFT where the dual signals correspond to the finite complex exponential basis. Computationally the standard inner product is a single sample of the correlation sequence, which can be implemented using convolution. For example, to compute the standard inner product of a transient signal $x_d[n]$ with $\psi_k[n]$, for a fixed index k, a linear convolution is performed with the time-reversed dual basis signal $\psi_k[-n]$. Denoting the output of this linear convolution as $s_k[n]$, the output is given by

$$s_k[n] = x_d[n] * \psi_k[-n] = \sum_{l=0}^{N-1} x_d[l]\,\psi_k[l-n]\,. \tag{3.3.5}$$

By sampling the output $s_k[n]$ at $n = 0$, the expansion coefficient α_k is computed, i.e., $s_k[0] = \langle x_d, \psi_k \rangle$. A signal flow graph representing this procedure for transient spectral analysis is depicted in Figure 3.3.1.
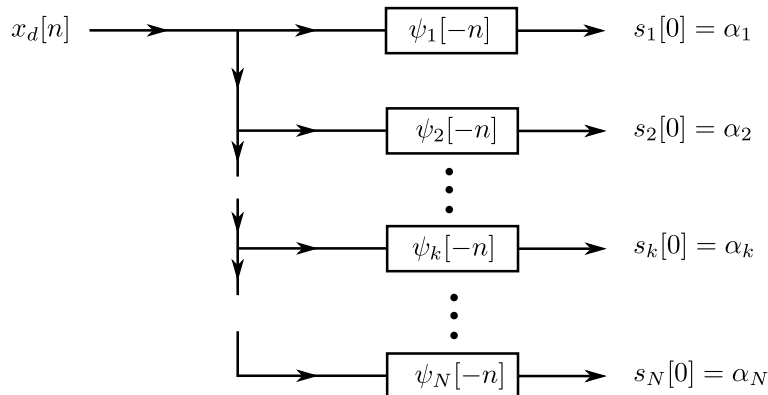
Figure 3.3.1: The DTT analysis filter bank used to generate the transient spectral coefficients $\mathcal{A}[k]$

The definition of another mapping, referred to as transient filtering, is established by the following procedure. First, compute the DTT of the signal $x[n]$, i.e. $\mathcal{A}[k] = \mathcal{T}\{x[n]\}$ where $\mathcal{T}\{\cdot\}$ is the transient analysis mapping in Eq. (3.3.3). Next, modify the transient spectrum $\mathcal{A}[k]$ by multiplying it element-wise by a function $\mathcal{H}[k]$, for each value of k. Finally, the output is defined by $y[n]$ where $y[n] = \mathcal{T}^{-1}\{\mathcal{A}[k]\mathcal{H}[k]\}$ and $\mathcal{T}^{-1}\{\cdot\}$ is the transient synthesis mapping in Eq. (3.3.4). The input-output relationship of this mapping depends on the structure of $\mathcal{H}[k]$, which may be designed freely. This mapping describes the procedure for a novel signal processing algorithm. Note that the mapping is linear but not circularly time-invariant, and therefore cannot, in general, be represented as the circular convolution of an input signal with any impulse response. As a trivial example, $\mathcal{H}[k] = 1$, for $1 \leq k \leq N$, results in an identity mapping.

As a more interesting example, the design of a high decay pass filter is given by

$$\mathcal{H}_{HP}[k] = \begin{cases} 1, & \sigma_1 \leq \sigma_k < \sigma_c \\ 0, & \sigma_c \leq \sigma_k \leq \sigma_N \end{cases}$$

for some $0 < \sigma_c < 1$. This choice of a transient domain scaling sequence passes all transient components in a signal which decay faster than the cutoff decay rate σ_c, and rejects all transient components decaying slower than σ_c. The design of the complementary low decay pass filter is then defined by $\mathcal{H}_{LP}[k] = 1 - \mathcal{H}_{HP}[k]$ for the same value of σ_c. Using these two filters to separate a signal into two transient channels allows for different processing of the slowly decaying and rapidly decaying components. In a similar fashion to this example, numerous different regions of decay rate may be selected, emphasized, or attenuated depending on the desired processing of the input signal. In Chapter 5, an adaptive transient filtering algorithm is proposed.

3.4 Algorithms for Generating Dual Exponential Bases

We have shown that in order to compute the expansion coefficients with respect to a general

exponential basis, the orthogonality of the basis cannot be assumed. Therefore, the standard inner product of the corresponding dual basis signals may be taken with the signal to be analyzed in order to produce the desired expansion coefficients. In practice, this requires the dual signals to be available for computation. This section derives the dual basis signals using three algorithms: the first generates the dual signals for any basis while the remaining two are specialized to generating the dual basis of a general exponential basis specifically. Before presenting these algorithms, the uniqueness of a dual basis is presented in order to confirm that the three algorithms all theoretically result in identical bases, though in practice the resulting bases are often different due to implementation errors.

3.4.1 Uniqueness of the Dual Basis

We begin by defining a linear map ζ from a vector space to the field of scalars over which the vector space is defined, i.e., $\zeta(\cdot) : \mathcal{V} \to \mathcal{F}$. The uniqueness of a dual basis is established by showing that the biorthogonality constraints define a unique linear map ζ for each dual signal. In order to understand how a dual signal paired with the standard inner product acts as the linear map ζ, consider Figure 3.3.1. This figure depicts the transient spectral coefficient $\mathcal{A}[k]$ as the result of sampling the correlation of two signals at a specific time, i.e.,

$$\mathcal{A}[k] = s_k[0] = \sum_{l=0}^{N-1} x_d[l]\,\psi_k[l] = \langle x_d, \psi_k \rangle. \tag{3.4.1}$$

In this example, the combination of the dual signal $\psi_k \in \mathcal{V}$ and the standard inner product is then interpreted as taking a linear combination of the input signal $x_d[n]$ to produce a scalar $\mathcal{A}[k] \in \mathcal{F}$, i.e., $\zeta(\cdot) = \langle \cdot, \psi_k \rangle$. We next show that there is a unique signal $v \in \mathcal{V}$ such that

$$\xi(u) = \langle u, v \rangle \tag{3.4.2}$$

for every $u \in \mathcal{V}$. To see this we first show that there exists a vector $v \in \mathcal{V}$ such that $\xi(v) = \langle u, v \rangle$ for every $u \in \mathcal{V}$, and then we show that only one vector $v \in \mathcal{V}$ has this desired behavior.

Let $\{\phi_k\}_{-\infty:\infty}$ be an orthonormal basis of $\mathcal{V}$. Then decompose the transformation as

$$\xi(u) = \xi\left(\sum_k \langle u, \phi_k \rangle \phi_k\right) \tag{3.4.3}$$

$$= \sum_k \langle u, \phi_k \rangle \xi(\phi_k) \tag{3.4.4}$$

$$= \left\langle u, \sum_k \xi^*(\phi_k)\phi_k \right\rangle. \tag{3.4.5}$$

By setting $v = \sum_k \xi^*(\phi_k)\phi_k$ we have $\xi(u) = \langle u, v \rangle$ for every $u \in \mathcal{V}$, as desired.

To see that only one vector $v \in \mathcal{V}$ has this behavior we define $v_1, v_2 \in \mathcal{V}$ such that

$$\xi(u) = \langle u, v_1 \rangle = \langle u, v_2 \rangle \tag{3.4.6}$$

for every $u \in \mathcal{V}$. Then

$$0 = \langle u, v_1 \rangle - \langle u, v_2 \rangle = \langle u, v_1 - v_2 \rangle \tag{3.4.7}$$

for every $u \in \mathcal{V}$. Taking $u = v_1 - v_2$ shows the uniqueness. In other words, $v_1 = v_2$. Therefore we may conclude that the dual signals are the unique set of signals as defined by the biorthogonality constraints.

3.4.2   Generating Dual Bases by Matrix Inversion

The algorithm for generating the dual basis by matrix inversion may be used for any basis, however we specialize the discussion below to a general exponential basis $\{\phi_k[n]\}_{1:N}$ in order to establish useful notation. Define the set of dual signals $\{\psi_k[n]\}_{1:N}$ to satisfy the biorthogonality constraints given by

$$\langle \phi_k, \psi_l \rangle = \delta_{kl}, \tag{3.4.8}$$

for $1 \leq l, k \leq N$, i.e., there are no inner product constraints between the original basis signals or the dual signals themselves. The only inner product constraints are pairwise between the exponential basis signals and the dual basis signals. This is distinct from enforcing orthogonality between elements of the dual basis, as will be explored in Section 4.6 on Inner Product Shaping. In matrix notation the biorthogonality constraints in Eq. (3.4.8) are

$$\Phi \Psi^H = I_N \tag{3.4.9}$$

where the matrix Ψ is defined by

$$\Psi = \begin{bmatrix} | & & | \\ \psi_1[n] & \cdots & \psi_N[n] \\ | & & | \end{bmatrix} \tag{3.4.10}$$

and I_N is the identity operator on $\mathcal{V}$. Therefore the dual signals $\{\psi_k[n]\}_{1:N}$ may be determined by taking the columns of the matrix Ψ given by

$$\Psi = \left(\Phi^H \right)^{-1}. \tag{3.4.11}$$

The matrix Φ^H is guaranteed to be invertible as its columns form a basis, i.e., are linearly independent. The uniqueness of a matrix inverse, when a matrix is full rank, is consistent with the uniqueness of the dual basis previously discussed. Solving this system of equations using a matrix inverse provides the entire set of dual signals. Each of the expansion coefficients are found

simultaneously by taking the inner product of $x_d[n]$ with each dual signal, or equivalently

$$\alpha = \Psi^H x_d. \tag{3.4.12}$$

Note that when the decay rates of a transient structured signal are known, we may choose $\Phi = V\left(\sigma^T\right)$. In this case α is the solution for the amplitude coefficients by directly solving Eq. (2.1.5) using matrix operations.

The canonical procedure for finding the dual basis signals, given in Eq. (3.4.11), has two major disadvantages. First, the dual basis signals must all be computed simultaneously, which is inefficient if only specific expansion coefficients are of interest. Second, for the real exponential basis, the matrix to be inverted, Φ^H, has a Vandermonde structure, for which it is well known that Vandermonde structured matrices are poorly conditioned when they are constructed with real distinct roots, as will be discussed in Section 3.7. [9] As a consequence, the solutions found using matrix inversion may contain large amounts of computational error. Motivated largely by these these two factors, two alternative algorithms for generating the dual basis signals for a general exponential basis is proposed in the following two subsections.

3.4.3   Generating Dual General Exponential Bases by Modulation

The algorithm proposed in this subsection eliminates both of the disadvantages discussed when the dual exponential basis is generated by matrix inversion. Specifically, it avoids computing the inverse of a potentially ill-conditioned matrix and allows for the dual signals to be found independently. To begin developing this method, define the mapping $\kappa : \mathcal{V} \times \mathcal{V} \to \mathcal{F}$ as

$$\kappa\left(x, y\right) = \sum_n x\left[n\right] y^*\left[n\right] \mu\left(y^*\left[n\right]\right) \tag{3.4.13}$$

for any $x, y \in \mathcal{V}$ and where the function $\mu\left(\cdot\right) : \mathcal{V} \to \mathcal{V}$ will be defined shortly. Note that κ is linear in its first argument and non-linear in its second. Also, note that κ is structurally identical to the standard inner product given by $\langle x, y \circ \mu\left(y\right)\rangle$, where $\circ$ denotes element by element vector multiplication, or equivalently, signal modulation.

Given a finite general exponential basis $\{\phi_k\left[n\right]\}_{1:N}$, our interest lies in determining a set of modulating signals $\{\mu\left(\phi_k\left[n\right]\right)\}_{1:N}$, such that the biorthogonality constraints hold, i.e., $\kappa(\phi_k, \phi_l) = \delta_{kl}$ for $1 \leq l, k \leq N$ where $\psi_l\left[n\right] = \phi_l\left[n\right] \circ \mu\left(\phi_l\left[n\right]\right)$. Thus the modulating signals, $\{\mu\left(\phi_l\left[n\right]\right)\}_{1:N}$, are named primarily because of their interpretation as the set of signals used to modulate the basis signals in order to generate the dual signals. Note that this algorithm relies upon the basis signals being non-zero for all sample values, which is guaranteed for the general exponential basis. For

$1 \leq l,\, k \leq N$, the biorthogonality constraints are then given by

$$
\kappa\left(\phi_k, \phi_l\right) = \phi_l^H \mathcal{K}\left(\phi_l\right) \phi_k = \begin{cases} 1, & k = l \\ 0, & k \neq l \end{cases}, \quad \text{where } \mathcal{K}\left(\phi_l\right) = \begin{bmatrix} \mu\left(\phi_l^*[0]\right) & 0 & \cdots & 0 \\ 0 & \mu\left(\phi_l^*[1]\right) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mu\left(\phi_l^*[N-1]\right) \end{bmatrix}.
$$

(3.4.14)

The following proposition states that for a general exponential basis, the set of modulating signals, $\{\mu\left(\phi_k[n]\right)\}_{1:N}$, which satisfy the desired biorthogonality constraints exist and are unique. However, for simplicity in the following proposition, the constraint of $\kappa(\phi_l, \phi_l)$ being unity is relaxed such that any non-zero scalar is accepted. In this way an appropriate normalization of the resulting modulating signal can be applied later to enforce the biorthogonality constraints as defined in Eq. (3.4.14). These unnormalized modulating signals are denoted by $\{\nu\left(\phi_k[n]\right)\}_{1:N}$, i.e., $\nu\left(\phi_k[n]\right) \propto \mu\left(\phi_k[n]\right)$.

Proposition. *Given a general exponential basis* $\{\phi_k[n]\}_{1:N}$*, then a set of unnormalized modulating signals* $\{\nu\left(\phi_k[n]\right)\}_{1:N}$ *exists such that*

$$
\kappa\left(\phi_k, \phi_l\right) = \sum_{n=0}^{N-1} \phi_k[n]\,\phi_l^*[n]\,\nu\left(\phi_l^*[n]\right) = \begin{cases} \gamma_l^{-1}, & k = l \\ 0, & k \neq l \end{cases}
\tag{3.4.15}
$$

for $1 \leq l, k \leq N$ *and for some non-zero constants* $\{\gamma_k\}_{1:N}$*. Furthermore, the unnormalized modulating signals are unique to within a scaling.*

Proof. We have previously shown that the dual signal $\psi_l[n]$, for a fixed index l, lies in the orthogonal complement of any linear map $\Phi_{\sim l} : \mathcal{V} \to \mathcal{V}_{\sim l}$, where the vector space $\mathcal{V}$ is decomposed as

$$
\mathcal{V} = \mathcal{V}_l \oplus \mathcal{V}_{\sim l}
$$

where $\mathcal{V}_l = \mathrm{span}\{\phi_l[n]\}$ and $\mathcal{V}_{\sim l} = \mathrm{span}\{\{\phi_k[n]\}_{1:N} \setminus \{\phi_l[n]\}\}$. For each l, $1 \leq l \leq N$, repeat the following argument. We choose $\Phi_{\sim l}$ to be the linear transformation

$$
\Phi_{\sim l} = \begin{bmatrix} | & & | & | & & | \\ \phi_1 & \cdots & \phi_{l-1} & \phi_{l+1} & \cdots & \phi_N \\ | & & | & | & & | \end{bmatrix}
$$

By definition we have that $\psi_l[n] \propto \phi_l[n] \circ \nu\left(\phi_l[n]\right)$, therefore the modulating signal $\nu(\phi_l[n])$ must be in the orthogonal complement to the space spanned by $\{\phi_1 \circ \phi_l, \cdots, \phi_{l-1} \circ \phi_l, \phi_{l+1} \circ \phi_l, \cdots, \phi_N \circ \phi_l\}$,

i.e.,

$$\nu\left(\phi_l\left[n\right]\right) \in \mathcal{R}^\perp\left(\begin{bmatrix} | & & | & | & & | \\ \phi_1 & \cdots & \phi_{l-1} & \phi_{l+1} & \cdots & \phi_N \\ | & & | & | & & | \end{bmatrix} \begin{bmatrix} \phi_l\left[0\right] & & 0 \\ & \ddots & \\ 0 & & \phi_l\left[N-1\right] \end{bmatrix}\right).$$

Written in matrix notation this is $\nu(\phi_l) \in \mathcal{R}^\perp\left(\Phi_{\sim l}\Phi_l\right)$. The product $\Phi_{\sim l}\Phi_l$ is guaranteed to span an $N-1$ dimensional space because each of the $N-1$ basis signals are modulated by the same non-zero signal, so no two basis signals end up being collinear after the modulation. Therefore $\dim\left(\mathcal{R}^\perp\left(\Phi_{\sim l}\Phi_l\right)\right) = 1$ by the Rank Plus Nullity Theorem, meaning the modulating signals are unique to within a scaling.

$\square$

The proof of the above proposition not only shows the existence of the unnormalized modulating signals, but also suggests an algorithm to determine them one at a time. In order to compute a nullspace we define the full Singular Value Decomposition (SVD) next. However, any technique to determine the nullspace of a linear map may be used in the following algorithm.

Definition. The Singular Value Decomposition (SVD). *If A is an $N \times N$, rank r matrix, then there exists an $N \times r$ matrix U and an $N \times r$ matrix V such that $U^H U = V^H V = I_r$ and $A = U\Pi V^H$, where Π_r is an $r \times r$ diagonal matrix whose diagonal entries, called singular values, satisfy*

$$\pi_1 \geq \pi_2 \geq \cdots \geq \pi_r > 0. \tag{3.4.16}$$

The extended or full SVD can be written as

$$A = \begin{bmatrix} u_1 & \cdots & u_r & \bar{u}_1 \cdots \bar{u}_{N-r} \end{bmatrix} \begin{bmatrix} \Pi_r & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_1^H \\ \vdots \\ v_r^H \\ \bar{v}_1^H \\ \vdots \\ \bar{v}_{N-r}^H \end{bmatrix} = \sum_{k=1}^{r} \pi_k u_k v_k^H + \sum_{k=1}^{N-r} 0\bar{u}_k\bar{v}_k^H. \tag{3.4.17}$$

The vectors u_k and v_k are the left singular vectors and right singular vectors of A, respectively, and comprise the columns of the matrices U and V. The vectors $\bar{u}_k$ and $\bar{v}_k$ span the null space of A^H and A, respectively.

By definition, all N values of σ_k are distinct. Therefore, $\sigma_1\sigma_l \neq \cdots \neq \sigma_{l-1}\sigma_l \neq \sigma_{l+1}\sigma_l \neq \cdots \neq \sigma_N\sigma_l$, so all $N-1$ of these values are also distinct. Define the function $\zeta_j\left[n\right] = \left(\sigma_l\sigma_j\right)^n$ for $0 \leq n \leq N-1$ and $l \neq j$. The standard inner product of $\nu\left(\phi_l\left[n\right]\right)$ with $\zeta_j\left[n\right]$, for each j, must be

0 for the desired constraints to be met. Writing these constraints in matrix notation, we have that $\nu\left(\phi_l\left[n\right]\right)$ must satisfy

$$
\begin{bmatrix}
\zeta_1\left[0\right] & \cdots & \zeta_1\left[N-1\right] \\
\vdots & \ddots & \vdots \\
\zeta_{l-1}\left[0\right] & \cdots & \zeta_{l-1}\left[N-1\right] \\
\zeta_{l+1}\left[0\right] & \cdots & \zeta_{l+1}\left[N-1\right] \\
\vdots & \ddots & \vdots \\
\zeta_N\left[0\right] & \cdots & \zeta_N\left[N-1\right]
\end{bmatrix}
\begin{bmatrix}
\nu\left(\phi_l\left[0\right]\right) \\
\vdots \\
\nu\left(\phi_l\left[N-1\right]\right)
\end{bmatrix}
=
\begin{bmatrix}
0 \\
\vdots \\
0 \\
0 \\
\vdots \\
0
\end{bmatrix}.
\tag{3.4.18}
$$

The matrix above is Vandermonde with $N-1$ linearly independent rows. Therefore the null space of this matrix contains exactly one non-trivial vector corresponding to $\nu(\phi_l[n])$.

Denote the matrix in Eq. (3.4.18) as Q_l. The vector $\bar{v}_N$ in the full SVD representation of Q_l is a basis for the null space and consequently is a scalar multiple of $\mu\left(\phi_l\left[n\right]\right)$, i.e., $\bar{v}_N = \nu\left(\phi_l\right)$. Therefore, the modulating signal, for a fixed index l, is found to be

$$
\mu\left(\phi_l\left[n\right]\right) = \gamma_l \cdot \nu\left(\phi_l\left[n\right]\right)
\tag{3.4.19}
$$

for some scalar γ_l. In order to use Eq. (3.1.10) to produce the transient spectrum we enforce that $\langle\phi_l, \phi_l \circ \mu\left(\phi_l\right)\rangle = 1$. Therefore we find that, for each l,

$$
\gamma_l = \frac{1}{\langle\phi_l,\, \nu\left(\phi_l\right) \circ \phi_l\rangle}.
\tag{3.4.20}
$$

If the entire transient spectrum is desired, then this process only needs to be repeated for each l, $1 \le l \le N$. The normalization by γ_l is the appropriate scaling for the set of biorthogonality constraints to be met.

This algorithm has the advantage of enabling the separate computation of the modulating signal corresponding to each exponential basis signal. The ability to produce the dual basis signals one at a time is given by

$$
\psi_k\left[n\right] = \gamma_k \cdot \phi_k\left[n\right] \circ \nu\left(\phi_k\left[n\right]\right),
\tag{3.4.21}
$$

for $1 \le k \le N$. This resolves the first of the two limitations discussed when the dual basis is constructed by matrix inversion by allowing the dual signals to be computed independently. The second issue of computational noise can be mitigated by using either algorithm based upon whether computing the nullspace of the matrix Q_l or inverting the matrix Φ^H results in lower computational error.

3.4.4 Generating Dual General Exponential Bases by Polynomial Expansion

The derivation of the algorithm for generating the dual general exponential basis by modulation

took advantage of the structure of a general exponential basis such that a matrix inverse was avoided. However, any technique for finding the null space of a matrix may be used to find the dual signals. A third algorithm, also specific to a general exponential basis, is now proposed for generating the dual exponential signals directly by exploiting the structure in the unnormalized modulating signals $\{\nu\left(\phi_k\left[n\right]\right)\}_{1:N}$. This results in a formulation where the dual basis signals are determined using polynomial expansion and knowledge of the parameters $\{\sigma_k\}_{1:N}$ used to define the general exponential basis.

An important observation about Vandermonde structured matrices and their relation to zeros of polynomials, which will be useful in the ensuing discussion, is highlighted first. This approach is based on the fact that a polynomial P of degree $N < \infty$ with scalar coefficients of the form

$$P\left(z\right) = \sum_{k=0}^{N} b_k z^k \tag{3.4.22}$$

$$= \prod_{k=1}^{N} b_N \left(z - \beta_k\right) \tag{3.4.23}$$

has N, not necessarily distinct or real, zeros $\{\beta_k\}_{1:N}$. P evaluated at each zero, β_k, must equal 0, i.e., $P\left(\beta_k\right) = 0$, for $1 \leq k \leq N$. Consequently, the polynomial coefficients must also satisfy the matrix equation

$$\begin{bmatrix} 1 & \beta_1 & \beta_1^2 & \cdots & \beta_1^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \beta_N & \beta_N^2 & \cdots & \beta_N^N \end{bmatrix} \begin{bmatrix} b_0 \\ \vdots \\ b_N \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} \tag{3.4.24}$$

where the matrix above has a Vandermonde structure. Note that in Eq. (3.4.24), the N roots of P appear explicitly in the second column of the $N \times N + 1$ matrix. Further, the vector of coefficients $b = [b_0, \cdots, b_N]^T$ lies in the nullspace of the Vandermonde structured matrix. When the roots of the polynomial are unique, the vector b is a basis for the nullspace as well.

We use the relationship between a polynomial and a Vandermonde matrix to generate a dual basis signal by exploiting the matrix Q_l to find the unnormalized modulating signal $\nu\left(\phi_l\left[n\right]\right)$, for a fixed index l, using polynomial expansion. The matrix Q_l in Section 3.4.3 is Vandermonde with $[Q_l]_{ij} = \left(\beta_i^{(l)}\right)^{j-1}$, where

$$\beta_i^{(l)} = \begin{cases} \sigma_l \sigma_i, & 1 \leq i < l \\ \sigma_l \sigma_{i+1}, & l \leq i < N \end{cases}. \tag{3.4.25}$$

Using the relationship between a polynomial and a Vandermonde matrix established in Eq. (3.4.24), we construct a polynomial in z^{-1} given by

$$\mathcal{V}^{(l)}(z) \;=\; \prod_{j=1}^{N-1} \left(1 - \left(\beta_j^{(l)}\right)^{-1} z^{-1}\right) \tag{3.4.26}$$

$$=\; \prod_{j=1,\,j\neq l}^{N} \left(1 - (\sigma_j\sigma_l)^{-1} z^{-1}\right) \tag{3.4.27}$$

$$=\; \sum_{n=0}^{N-1} \nu\left(\phi_l\left[n\right]\right) z^{-n} \tag{3.4.28}$$

where for each l, $1 \leq l \leq N$, Eq. (3.4.27) defines an expression for the z-transform of $\nu\left(\phi_l\left[n\right]\right)$, denoted by $\mathcal{V}^{(l)}(z)$. Therefore the z-transform of the modulating signal $\mu\left(\phi_l\left[n\right]\right)$, denoted by $\mathcal{M}^{(l)}(z)$, is defined by $\mathcal{M}^{(l)}(z) = \gamma_l\mathcal{V}^{(l)}(z)$. Using Eq. (3.4.27) is an improvement upon the previous two algorithms for dual basis construction in the sense that neither a matrix factorization nor a matrix inverse need to be computed. Instead, any computational error arising in this calculation is due only to the multiplication of the inverse parameters $\{\sigma_k\}_{1:N}$ used to define the general exponential basis signals.

Continuing further with the interpretation of Eq. (3.4.27) as the z-transform of $\nu\left(\phi_l\left[n\right]\right)$, for some fixed index l, the product of $N-1$ binomial terms in the z-domain is equivalent to $N-1$ convolutions of length-two signals in the temporal domain. Each length two signal has the form

$$h_k^{(l)}\left[n\right] = \mathcal{Z}^{-1}\left\{1 - (\sigma_k\sigma_l)^{-1} z^{-1}\right\} \tag{3.4.29}$$

$$= \delta\left[n\right] - (\sigma_k\sigma_l)^{-1} \delta\left[n-1\right] \tag{3.4.30}$$

for $1 \leq k \leq N$, $k \neq l$. The convolutional procedure for generating an unnormalized modulating signal is given by

$$\nu\left(\phi_l\left[n\right]\right) = h_1^{(l)}\left[n\right] * \cdots * h_{l-1}^{(l)}\left[n\right] * h_{l+1}^{(l)}\left[n\right] * \cdots * h_N^{(l)}\left[n\right] \tag{3.4.31}$$

for $1 \leq l \leq N$. Together Eqs. (3.4.30) and (3.4.31) yield an expression for determining the unnormalized modulating signals $\{\nu\left(\phi_k\left[n\right]\right)\}_{1:N}$ using only knowledge of the parameters $\{\sigma_k\}_{1:N}$.

Figure 3.4.1 shows a signal flow graph depiction of the algorithm for dual exponential basis generation by polynomial expansion. The dual basis signals are given by $\psi_k\left[n\right] = \gamma_k\tilde{\psi}_k\left[n\right]$, for $1 \leq k \leq N$, where the unnormalized dual signals $\tilde{\psi}_k\left[n\right]$ are generated independently by constructing scalar multiples of the unnormalized modulating signals, $\nu\left(\phi_k\left[n\right]\right)$, and then modulating each by the appropriate exponential basis signal. Also depicted is the construction of the scaling coefficients in Eq. (3.4.20), where the summation needed for the inner-product is performed by sampling a rectangular window filter at the appropriate sample value.
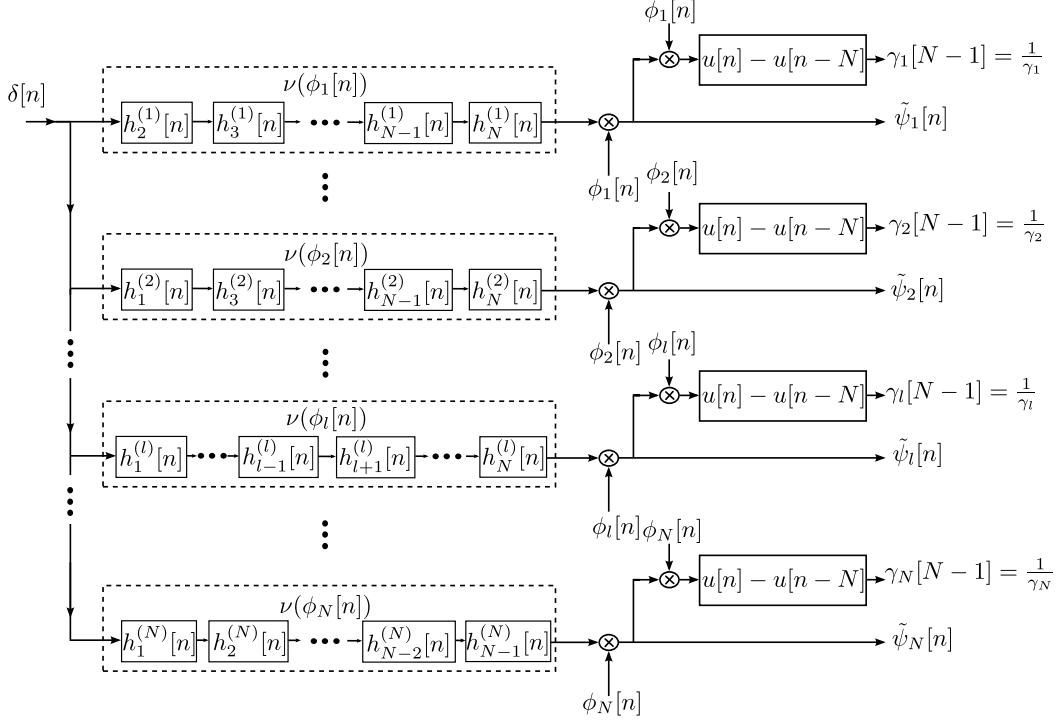
Figure 3.4.1: The signal flow graph for generating the unnormalized dual exponential basis signals $\left\{\tilde{\psi}_k[n]\right\}_{1:N}$ and the normalization constants $\{\gamma_k\}_{1:N}$ directly using the modulation algorithm

Consider the zeros of the z-transform of an unnormalized modulating signal $\nu(\phi_l[n])$, for a fixed index l. The zeros possess additional structure in the sense that every zero in Eq. (3.4.27) is a multiple of the parameter σ_l^{-1}. This additional structure is exploited next.

Vieta's formula [26], which gives the relationship between the coefficients of a polynomial and functions of its zeros, states that, using the notation of Eqs. (3.4.22) and (3.4.23),

$$
\begin{aligned}
\beta_1 + \cdots + \beta_N &= -\frac{b_{N-1}}{b_N} \\
(\beta_1\beta_2 + \beta_1\beta_3 + \cdots + \beta_1\beta_N) + (\beta_2\beta_3 + \beta_2\beta_4 + \cdots + \beta_2\beta_N) + \cdots + \beta_{N-1}\beta_N &= \frac{b_{N-2}}{b_N} \\
\vdots \quad &\quad \vdots \quad \vdots \\
\beta_1\beta_2\cdots\beta_N &= (-1)^N \frac{b_0}{b_N}
\end{aligned}
.
$$

(3.4.32)

Consider the zeros of $V^{(l)}(z)$ for a fixed index l. Each of the zeros are of the form $\beta_j^{-1} = \sigma_j^{-1}\sigma_l^{-1}$, as defined in Eq. (3.4.25). Therefore a factor of the form $\left(\sigma_l^{-1}\right)^k$, for some positive power k, can be factored out of each of the equations above when specialized to $V^{(l)}(z)$. Without loss of generality

assume that $l = 1$. The left-hand side of the set of equations above reduces to

$$\sigma_1^{-1} \left(\sigma_2^{-1} + \cdots + \sigma_N^{-1}\right)$$
$$\left(\sigma_1^{-1}\right)^2 \left(\left(\sigma_2^{-1}\sigma_3^{-1} + \sigma_2^{-1}\sigma_4^{-1} + \cdots + \sigma_2^{-1}\sigma_N^{-1}\right) + \left(\sigma_3^{-1}\sigma_4^{-1} + \cdots + \sigma_3^{-1}\sigma_N^{-1}\right) + \cdots + \sigma_{N-1}^{-1}\sigma_N^{-1}\right)$$
$$\vdots$$
$$\left(\sigma_1^{-1}\right)^N \left(\sigma_2^{-1} \cdots \sigma_N^{-1}\right).$$

(3.4.33)

This representation is interpreted as the coefficients of the z-transform of a signal $u\left(\phi_1\left[n\right]\right)$ modulated by $\tilde{\phi}_1\left[n\right] = \left(\sigma_1^{-1}\right)^n$ for $0 \leq n \leq N - 1$. Denote the z-transform of $u\left(\phi_1\left[n\right]\right)$ as $U^{(1)}\left(z\right)$, then $U^{(1)}\left(z\right)$ has zeros at $\sigma_2^{-1}, \ldots, \sigma_N^{-1}$. Generalizing to an arbitrary index l, $1 \leq l \leq N$, define

$$U^{(l)}\left(z\right) = \prod_{k=1,\, k\neq l}^{N} \left(1 - \sigma_k^{-1} z^{-1}\right) \tag{3.4.34}$$

to be the z-transform of $u\left(\phi_l\left[n\right]\right)$, a signal that when modulated by $\tilde{\phi}_l\left[n\right]$ produces $\nu\left(\phi_l\left[n\right]\right)$. That is

$$\nu\left(\phi_l\left[n\right]\right) = u\left(\phi_l\left[n\right]\right) \circ \tilde{\phi}_l\left[n\right] \tag{3.4.35}$$

for $1 \leq l \leq N$. Next, the product of $N - 1$ binomial terms in Eq. (3.4.34) is interpreted as $N - 1$ convolutions of length two sequences in the sample domain where each length two sequence has the form, for $1 \leq k \leq N$,

$$g_k\left[n\right] = \mathcal{Z}^{-1}\left\{1 - \sigma_k^{-1} z^{-1}\right\} \tag{3.4.36}$$
$$= \delta\left[n\right] - \sigma_k^{-1} \delta\left[n - 1\right]. \tag{3.4.37}$$

These length two sequences have many advantages over those defined in Eq. (3.4.30), as they are reused in the process of generating each dual basis signal. Therefore the convolutional expression for $u\left(\phi_l\left[n\right]\right)$ is given by, for $1 \leq l \leq N$,

$$u\left(\phi_l\left[n\right]\right) = g_1\left[n\right] * \cdots * g_{l-1}\left[n\right] * g_{l+1}\left[n\right] * \cdots * g_N\left[n\right]. \tag{3.4.38}$$

Directly using Eqs. (3.4.35) and (3.4.38) in order to generate the unnormalized modulating signals $\left\{\nu\left(\phi_l\left[n\right]\right)\right\}_{1:N}$ is shown next in Figure 3.4.2.
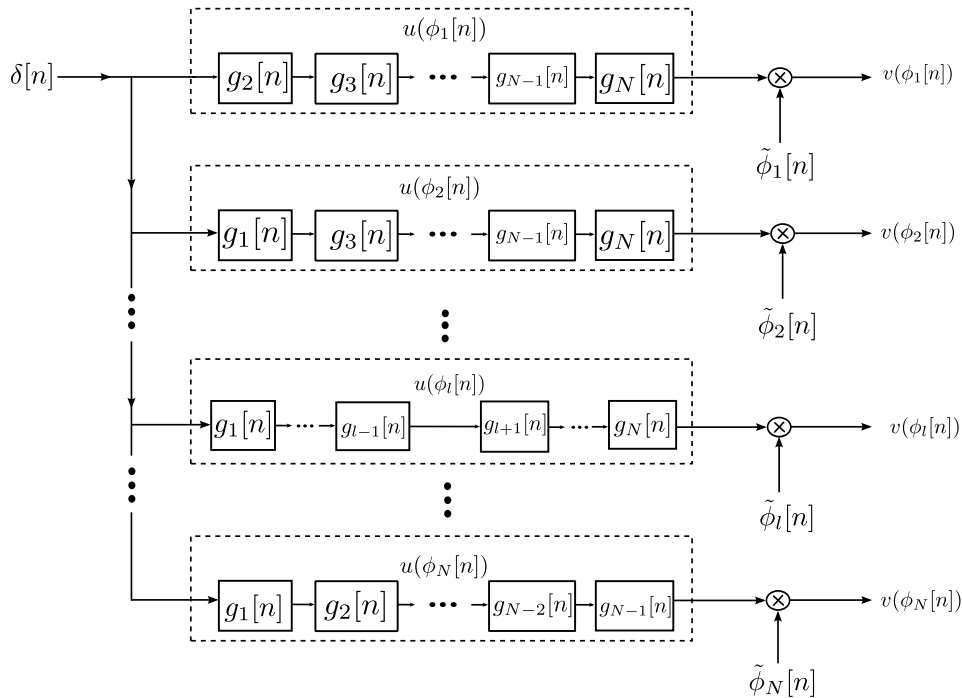
Figure 3.4.2: The signal flow graph for generating the unnormalized modulating signals $\{\nu\left(\phi_k\left[n\right]\right)\}_{1:N}$ using the modulated signals $\{u\left(\phi_k\left[n\right]\right)\}_{1:N}$

Combining the signal flow graphs in Figures 3.4.1 and 3.4.2, note that the modulating terms cancel, i.e., $\phi_l\left[n\right]\tilde{\phi}_l\left[n\right] = 1$, for $0 \leq n \leq N - 1$ and $1 \leq l \leq N$. An expression for the dual signals $\{\psi_l\left[n\right]\}_{1:N}$ is obtained then by directly using knowledge of the parameters $\{\sigma_k\}_{1:N}$ used to define the general exponential basis instead of pairs of these parameters. That is, for $1 \leq l \leq N$,

$$\psi_l\left[n\right] = \gamma_l \cdot u\left(\phi_l\left[n\right]\right),\tag{3.4.39}$$

where γ_l is the scalar in Eq. (3.4.20) and $u\left(\phi_l\left[n\right]\right)$ is given in Eq. (3.4.38). A signal flow graph implementation of this system is depicted in Figure 3.4.3. Again, the figure shows the procedure used to calculate the normalization parameters $\{\gamma_k\}_{1:N}$.
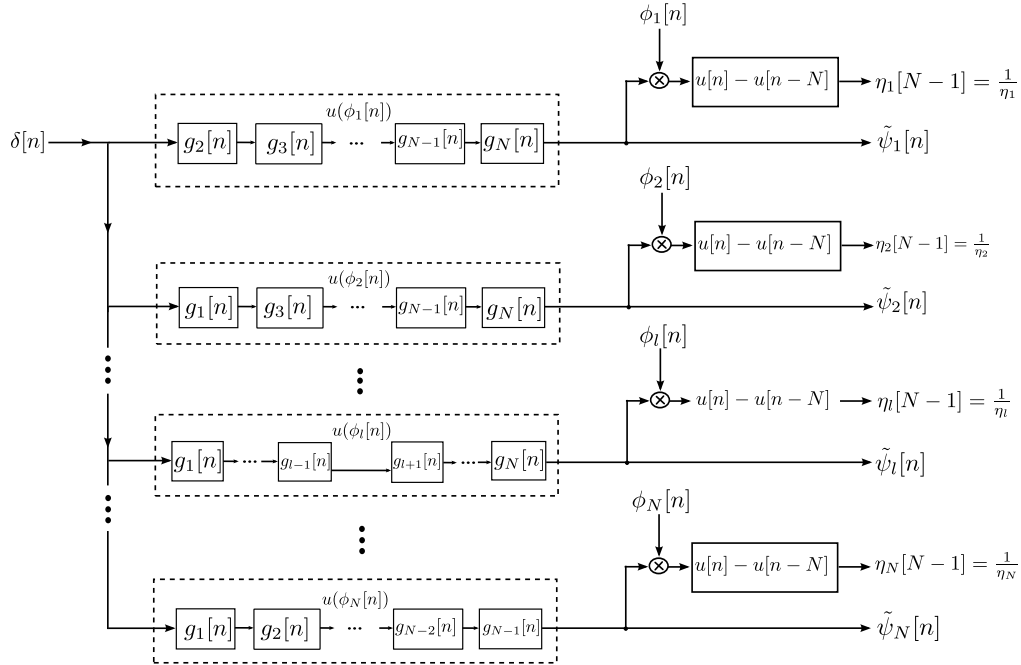
Figure 3.4.3: The simplified signal flow graph for generating the unnormalized dual exponential basis signals $\left\{ \tilde{\psi}_k\left[n\right] \right\}_{1:N}$ and the normalization constants $\left\{ \gamma_k \right\}_{1:N}$

One final simplification of the signal flow graph representation of dual basis generation shown in Figure 3.4.3 involves the re-use of the sub-system components $\left\{ g_k\left[n\right] \right\}_{1:N}$. Consider the polynomial in Eq. (3.4.34) where we reformulate it to be a rational function given by, for $1 \leq l \leq N$,

$$U^{(l)}\left(z\right) = \frac{\prod\limits_{k=1}^{N}\left(1 - \sigma_k^{-1}z^{-1}\right)}{\left(1 - \sigma_l^{-1}z^{-1}\right)}. \tag{3.4.40}$$

The numerator of this function is invariant to l, meaning it may be used for each $U^{(l)}\left(z\right)$, $1 \leq l \leq N$. Utilizing this fact gives another implementation for generating the dual basis signals and is shown in Figure 3.4.4.
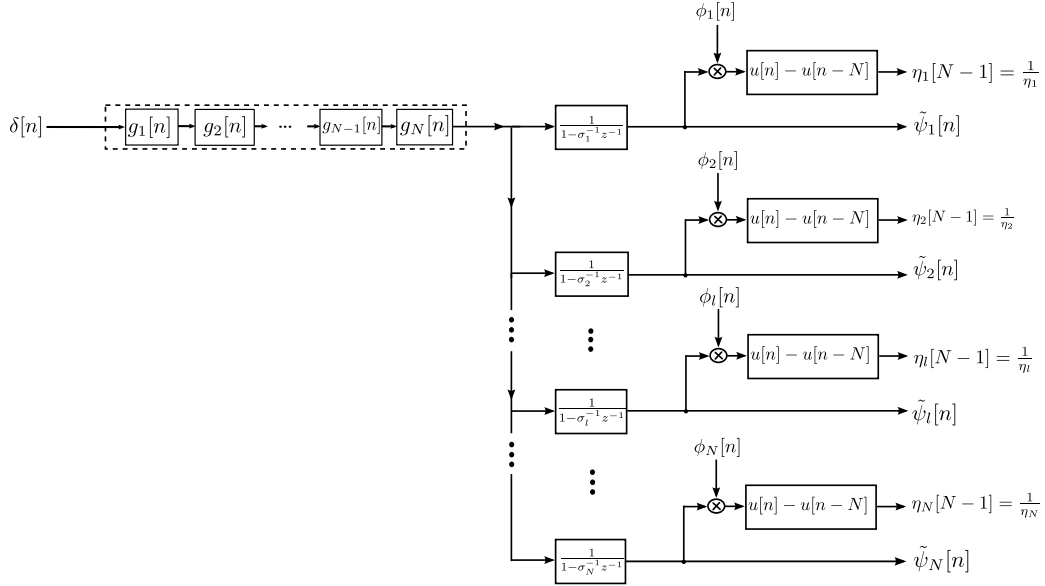
Figure 3.4.4: The signal flow graph optimized for computational savings for generating the unnormalized dual exponential basis signals $\left\{ \tilde{\psi}_k [n] \right\}_{1:N}$ and the normalization constants $\{\gamma_k\}_{1:N}$

3.5   Pole-Zero Interpretation of Polynomial Expansion

This section describes an alternative algorithm for generating the unnormalized dual signal, $\tilde{\psi}_k [n]$, using a modification of the algorithm presented in Section 3.4.4 for generating the dual exponential basis using polynomial expansion. The discussion in this section is specialized to the real exponential basis, although similar arguments hold for the general exponential basis. This algorithm uses a straightforward argument based upon the pole-zero representation of signals and systems. To begin, consider the set of N infinite duration exponential signals, $\left\{ \check{\phi}_k [n] \right\}_{1:N}$, where $\check{\phi}_k [n] = \sigma_k^n u [n]$, for $1 \leq k \leq N$. The process of designing a filter bank to detect which $\check{\phi}_k [n] \in \left\{ \check{\phi}_k [n] \right\}_{1:N}$ is the input to the filter by looking at an output snapshot is described next. Once this process has been stated, the system function for each channel in the filter bank is shown to be equivalent to the unnormalized dual signals found by polynomial expansion. The primary difference in this section and the derivation of the algorithm in Section 3.4.4 is that for finite-length exponential basis signals the arguments used here do not hold.

Define the structure of the l^{th} channels system function by $F^{(l)} (z)$, where $F^{(l)} (z)$ has a zero at each of the decay rates in the infinite duration exponential basis, except for decay rate σ_l. The impulse response of the l^{th} channels system function is denoted by $f^{(l)} [n]$. Therefore, the z-transform representation of the l^{th} channel is given by

$$\mathcal{Z} \left\{ f^{(l)} [n] \right\} = \prod_{k=1, \, k \neq l}^{N} \left(1 - \sigma_k z^{-1} \right). \tag{3.5.1}$$

The implication of the $N - 1$ zeros of $F^{(l)} (z)$, corresponding to $N - 1$ of the decay rates in the

infinite duration exponential basis $\left\{ \check{\phi}_k \left[n \right] \right\}_{1:N}$, is that $f^{(l)} \left[n \right]$ has a non-zero length corresponding to $0 \leq n \leq N - 1$. The pole-zero representation of $F^{(l)} \left(z \right)$ is shown in Figure 3.5.1(a), for a fixed index l. Each infinitely long exponential basis signal has the z-transform structure given by $\check{\Phi}_k \left(z \right) = \frac{1}{1 - \sigma_k z^{-1}}, 1 \leq k \leq N$. Therefore, for $k \neq l$, $F^{(l)} \left(z \right) \cdot \check{\Phi}_k \left(z \right)$ results in a pole-zero cancellation at σ_k. Equivalently, the convolution of $f^{(l)} \left[n \right]$ and $\check{\phi}_k \left[n \right]$ results in a signal of non-zero length for $0 \leq n \leq N - 2$. The pole-zero representation of this case is depicted in Figure 3.5.1(b). For the case where $k = l$, $F^{(l)} \left(z \right) \cdot \check{\Phi}_k \left(z \right)$ does not result in a pole-zero cancellation. The convolution of $f^{(l)} \left[n \right]$ and $\check{\phi}_k \left[n \right]$ results in an infinitely long non-zero output due to the pole at σ_k in the pole-zero representation. The pole-zero plot for this case is depicted in Figure 3.5.1(c).

(a) Pole-zero representation of the l^{th} channels system function $F^{(l)} \left(z \right)$

(b) Pole-zero representation of the output of the l^{th}-channel using the input $\check{\phi}_k \left[n \right]$ when $k \neq l$

(c) Pole-zero representation of the output of the l^{th}-channel using the input $\check{\phi}_k \left[n \right]$ when $k = l$
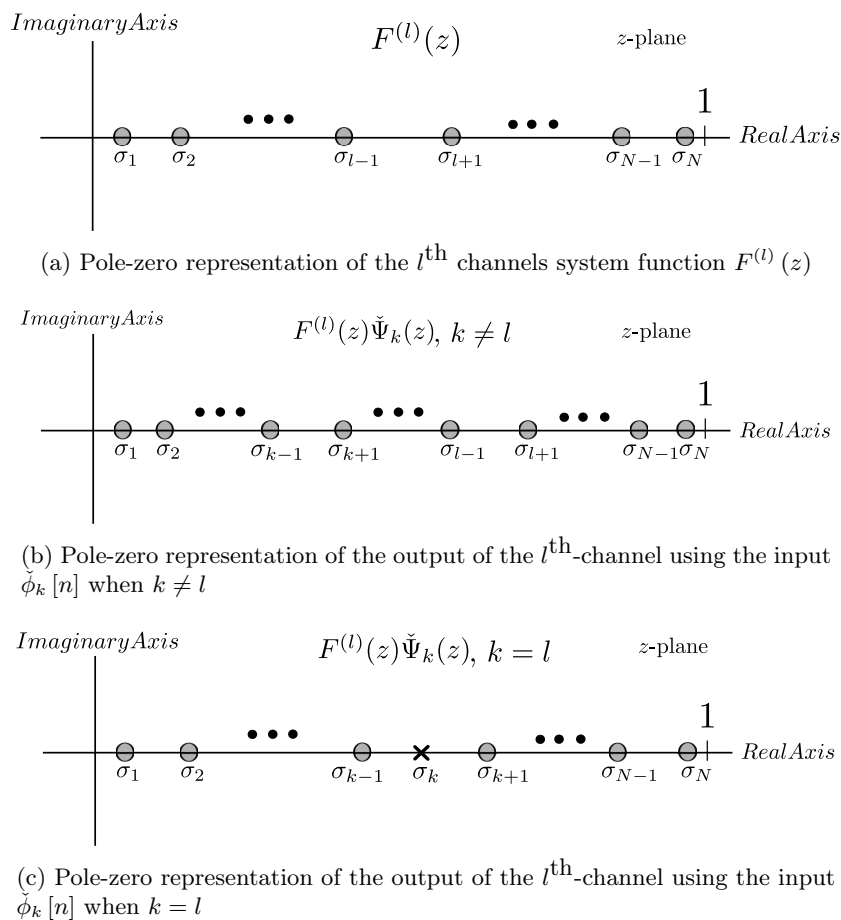
Figure 3.5.1: Pole-zero representations for understanding the dual exponential basis generation by polynomial expansion

To detect which of the infinitely long exponential signals is the input into the filter bank is easily obtained by looking at the snapshot of outputs at time $n = N - 1$. The index of the channel with non-zero output at this sample corresponds to the index of the infinitely long exponential signal that was the input into the filter bank. Note that without an appropriate scaling, the value of the output at this sample is not the expansion coefficient, but just an indication of presence. This means that if the linear combination of two infinitely long exponential signals are put into the filter bank,

then the two corresponding channels would have non-zero output at $n = N - 1$. However, the two output values would not necessarily indicate which exponential component had greater magnitude. Note that the l^{th} channel of the filter bank described above, i.e., $F^{(l)}(z)$, is structurally identical to $U^{(l)}(z)$ in Section 3.4.4.

3.6 Desirable Properties of Transient Spectra

The expansion coefficients resulting from a change of basis often have some physical significance with respect to the structure of basis signals selected. For example, one decomposition commonly used in a broad range of engineering disciplines is the Discrete Fourier Transform. The DFT has been successfully used across numerous disciplines; the importance of the DFT needs no elaboration. The orthogonal basis signals used in the DFT are harmonically related complex exponentials, and the expansion coefficients correspond, in amplitude, to the specific frequencies present in a sampled signal. The Fourier spectrum computed by the DFT has many properties that have aided it in its widespread success, only a subset of which are described in what follows. In discussing the DTT, many of the same properties exhibited by the DFT are desired.

As we saw in Section 3.1, the expansion coefficients for an orthogonal basis are easily determined through the standard inner product. Consequently, the dual signals of a purely imaginary complex exponential basis have a simple structure—they are also harmonically related complex exponentials. Specifically, each dual signal is the complex conjugate of its corresponding complex exponential basis signal, to within a scaling. As a result, both the basis signals and the dual basis signals remain numerically stable, even for large values of N. Parsevals theorem provides a simple relationship between the energy content of a signal in the standard sampling basis and the complex exponential basis. For signals with only a few large DFT coefficients, Parsevals theorem is helpful for determining an approximation to a signal while preserving as much energy content as possible, laying the foundation for many transform coding methods.

With regard to systems theory, the convenient result that an exponential signal is an eigenfunction of an LTI system means that the action of an LTI system, described by linear convolution, is restricted to a possibly complex scaling when the input is a complex exponential. This fact allows the Fourier spectrum to be interpreted as a spectrum of eigenvalues. Additionally, any finite-length signal can be represented using a complex exponential basis, where the correct "amount" of each complex exponential is the expansion coefficient, i.e., the DFT coefficient. Using this fact, linear convolution may be performed by multiplying the DFT of an input signal with the DFT of an LTI system's impulse response, assuming the degrees of freedom of the output are preserved in the resulting spectrum. Using an FFT algorithm for this procedure provides an alternative to the computationally intensive linear convolution by multiplying the possibly zero-padded DFTs of the signals involved. This procedure, referred to as Fast Convolution, motivates the design and implementation of numerous LTI systems directly in the complex exponential basis. Many physical systems, such as the human auditory system, are easily understood in the frequency domain, therefore designing directly in this basis is often advantageous.

The ability to resolve two closely spaced complex exponentials is well understood; the longer the interval of support, the higher the spectral resolution will be. Additionally, the ability to reduce spectral leakage, which occurs when a complex exponential being analyzed doesn't correspond to one of the basis signals, is mitigated by pre-processing the signal via windowing in the sample domain. When a signal acquired contains additive noise, the signal-to-noise ratio is invariant with respect to the number of samples acquired. This provides a meaningful indication of the distortion in the samples accounted for by the noise.

In parallel with the discussion above for the complex exponential basis, many similar properties will be presented for the exponential basis. In order to meaningfully discuss the transient spectrum, the dual exponential basis signals must be available and thus their structure was derived in Section 3.4. The remainder of the chapter discusses the numerical stability, transient resolution, spectral leakage, and the effects of additive noise with respect to the transient spectrum.

3.7 Ill-Conditioning of the Matrix Inversion Algorithm

In addition to miss-modeling and physical measurement errors, numerical errors in the implementation of algorithms for solving systems of linear equations often arise. Such errors are inherent in finite precision computer representations, and can play a significant role in the final error of an obtained solution. Previously, the algorithm for generating the dual exponential basis using matrix inversion labeled the process of inverting a Vandermonde structured matrix as an ill-conditioned procedure, potentially resulting in large computational errors. This section describes the sensitivity and conditioning of this matrix inversion problem as it was formulated in Section 3.4.2. In this section the SVD is used to analyze a system of equations in the form of Eq. (2.1.5) in order to understand the effect of noise on the matrix $V(\sigma^T)$. Note that $V\left(\sigma^T\right)$ is structurally equivalent to the matrix Φ which was used to find the dual signals in the matrix inversion algorithm. For the remainder of this section, A denotes a general square matrix, which includes the class of square Vandermonde matrices.

Every linear transformation possesses a Moore-Penrose pseudo-inverse, given by

$$A^\dagger = V\Pi^{-1}U^H = \sum_{k=1}^{r} \pi_k^{-1} v_k u_k^H. \tag{3.7.1}$$

When the matrix A is full rank, the inverse and the pseudo-inverse are identical. The Frobenius norm of a matrix, a straightforward generalization of the ℓ_2-norm of a vector, is defined by

$$\|A\|_F^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} |A_{ij}|^2 = \sum_{k=1}^{\mathrm{rank}(A)} \pi_k^2 = \mathrm{tr}\left(A^H A\right) \tag{3.7.2}$$

where $\mathrm{tr}(A) = \sum_{k=1}^{N} [A]_{kk}$. A second matrix norm, known as either the ℓ_2 or spectral norm, is given

by

$$\|A\|_2 = \max_{x, \|x\|_2=1} \|Ax\|_2 \,. \tag{3.7.3}$$

The spectral norm has the form of maximizing the gain of the matrix A, measured by $\frac{\|Ax\|_2}{\|x\|_2}$, over all possible directions x. It is straightforward to verify that the spectral norm of the matrix A is bounded by $\|A\|_2 \leq \pi_1$, where equality is achieved when $x = v_1$, the first right singular vector.

These tools enable us to determine the smallest perturbation that makes A singular in both the Frobenius and spectral norm sense. Because the solution to the system of equations given by $Ax = b$ is not unique when A is singular, when A is close to singular the solution is sensitive to small errors in b, even though it is unique. Consider a perturbation or error in b of δb. The solution is then perturbed to $x + \delta x$, where

$$\|\delta x\|_2 \leq \left\|A^\dagger\right\|_2 \cdot \|\delta b\|_2 \,. \tag{3.7.4}$$

Using the pseudo-inverse, we find that the spectral norm of $A^\dagger$ is bounded above by π_N^{-1}, therefore

$$\|\delta x\|_2 \leq \pi_N^{-1} \cdot \|\delta b\|_2 \,. \tag{3.7.5}$$

Note that the bound of a small error in b is inversely proportional to the smallest singular value; this bound is typically very large for Vandermonde structured matrices. However, we also know that $\|x\|_2$ is at least as large as $\pi_1^{-1} \|b\|_2$. Therefore the relative change in the solution, $\frac{\|\delta x\|_2}{\|x\|_2}$, is upper bounded by the ratio $\frac{\|\delta b\|_2}{\|b\|_2}$. When this bound is met the system is said to be well-conditioned, because the relative perturbation in the solution x is never larger than the relative perturbation in the measurement vector b. In general, we have that

$$\frac{\|\delta x\|_2}{\|x\|_2} \leq \frac{\pi_1}{\pi_N} \frac{\|\delta b\|_2}{\|b\|_2} = \|A\|_2 \left\|A^\dagger\right\|_2 \frac{\|\delta b\|_2}{\|b\|_2}. \tag{3.7.6}$$

In order to quantify the extent by which a small error in the right hand vector b can effect the solution to the linear system of equations, a metric known as the condition number is established. The definition of the condition number of a square matrix A is given by

$$\kappa_m \left(A\right) = \|A\|_2 \left\|A^\dagger\right\|_2 = \frac{\pi_1}{\pi_N} \tag{3.7.7}$$

where the condition number $\kappa_m \left(A\right) = 1$ is ideal. Note that the condition number does not depend on the right hand side of the system of equations $Ax = b$, but only on the matrix A. Also, note that the condition number of a matrix A, given this definition for the condition number, is the same as the condition number of $A^\dagger$. This means that the condition number of Φ is the same as Ψ, i.e., $\kappa_m \left(\Phi\right) = \kappa_m \left(\Psi\right)$.

From a numerical standpoint, solving for the inverse of a Vandermonde matrix is usually a bad approach, even though the matrix is guaranteed nonsingular. The condition number of Van-

dermonde systems has been shown to grow exponentially with the system size, yielding terribly ill-conditioned problems even for relatively small system complexities. [9] For example, the exponential basis with decay rates $\{\sigma_k\}_{1:17} = \{0.1, 0.15, \cdots, 0.9\}$ yields $\kappa_m(\Phi) = 8.1512 \times 10^{14}$.

3.8 Transient Spectral Leakage

In many scenarios, the amplitude coefficients of a transient signal are desired when a finite set of possible decay rates present in the signal is either unknown or unobtainable. When this occurs, one approach is to use the DTT with any set of any N distinct decay rates. In this case, the resulting transient spectrum will generally exhibit spectral leakage. This section introduces the effects of spectral leakage by way of an example, followed by a comparison to the spectral leakage exhibited in the complex exponential basis. In addition, the effects of noise on the transient spectrum are discussed.

Consider computing the transient spectrum of the signal $x_1[n] = 3(0.41)^n$ using an exponential basis with decay rates $\{\sigma_k\}_{1:17} = \{0.1, 0.15, \cdots, 0.9\}$. Note that for this example, the decay rate of the signal is not in the exponential basis, i.e., $0.41 \notin \{\sigma_k\}_{1:17}$. The magnitude of the resulting transient spectrum is depicted in Figure 3.8.1. One desirable attribute that was discussed in Section 3.6 is observed in this example, namely, the behavior of a broadened peak around the true decay rate. The maximum magnitude found in this example corresponds to decay rate $\sigma_7 = 0.4$ and the second highest is found corresponding to decay rate $\sigma_8 = 0.45$. For decay rates further from σ_7 and σ_8, the expansion coefficients decay to 0.

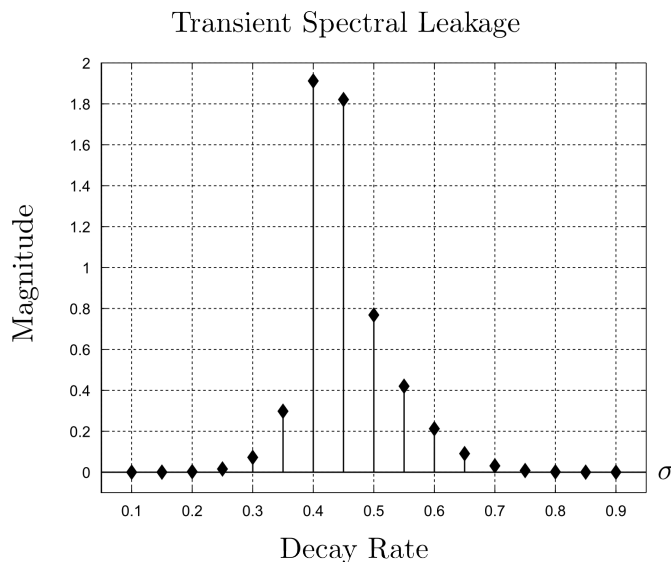Figure 3.8.1: Transient Spectrum of a single exponential signal depicting the effects of spectral leakage

A natural question arising from the transient spectrum in the previous example is whether or not any techniques exist to reduce the effects of spectral leakage. For example, a common pre-processing technique used in the DFT is to change the shape of the window the signal is multiplied by in the

sample domain, or equivalently a circular convolution in the frequency domain, which introduces a bias into the spectrum. A technique which reduces the spectral leakage for the transient spectrum is proposed in Section 4.1. Instead of modifying the signal to be analyzed, a relaxation of the biorthogonality constraints leads to reduced spectral leakage. Spectral leakage also arises due to the effects of additive noise contaminating the transient structured signal, such as in Eq. (2.1.7).

In order to explore the effects additive noise, we define the signal-to-noise ratio (SNR) of a finite-length noisy transient signal as

$$\text{SNR}_{\text{dB}}\left(\bar{x}_d\left[n\right]\right) = 10 \cdot \log_{10}\left(\frac{\displaystyle\sum_{n=0}^{N-1}\left(x_d\left[n\right]\right)^2}{\displaystyle\sum_{n=0}^{N-1}\left(\eta\left[n\right]\right)^2}\right). \tag{3.8.1}$$

Consider the signal $\bar{x}_1\left[n\right] = \phi_k\left[n\right] + \eta\left[n\right]$, for some k, $1 \leq k \leq N$. The denominator of the logarithm in $\text{SNR}_{\text{dB}}\left(\bar{x}_1\left[n\right]\right)$ is invariant to k while the numerator is monotonically increasing as σ_k approaches 1, meaning that the SNR monotonically increases as σ_k approaches 1 for a fixed amplitude coefficient. Given a fixed noise power, the transient spectral results are more accurate for slowly decaying exponentials for this reason; they contain more energy. However, with transient structured signals, the longer the interval of support for which sample values are obtained, the worse the SNR becomes. This behavior is demonstrated next as the data size grows large.

The following demonstrates one of the inherently difficult problems faced when dealing with noisy transient signals: the SNR, as calculated in Eq. (3.8.1), does not remain constant as the data size increases. Specifically, consider the SNR in the limit of the data record. A general transient signals instantaneous energy decays towards 0 as $n \to \infty$ while the noise power remains fixed. In the limit we find that

$$\lim_{N\to\infty}\text{SNR}_{\text{dB}}\left(\bar{x}_d\left[n\right]\right) \leq \lim_{N\to\infty} 10 \cdot \log_{10}\left(\frac{\displaystyle\sum_{n=0}^{N-1}\left(\alpha_{MAX}\left(\sigma_{MAX}\right)^n\right)^2}{\displaystyle\sum_{n=0}^{N-1}\left(\eta\left[n\right]\right)^2}\right) \tag{3.8.2}$$

$$\leq \lim_{N\to\infty} 10 \cdot \log_{10}\left(\frac{\dfrac{\alpha_{MAX}^2}{1-\sigma_{MAX}^2}}{\displaystyle\sum_{n=0}^{N-1}\left(\eta\left[n\right]\right)^2}\right) \tag{3.8.3}$$

$$= -\infty \tag{3.8.4}$$

where $\alpha_{MAX} = \displaystyle\sum_{k=1}^{d}\left|\alpha_k\right|$ and $\sigma_{MAX} = \max_k\left\{\sigma_k\right\}_{1:d}$. This result demonstrates an important property when working with transient signals: it is often best to use a shorter data record rather than longer,

because the overall SNR decreases to the extent that the recorded sample values contain strictly the effects of noise. When signal quantization is taken into account this phenomena results, in many cases, relatively quickly.

3.9 Dual Basis Generation for the Complex Exponential Basis

In this chapter three algorithms have been presented for generating a dual basis given a finite general exponential basis, two of which were derived explicitly for generating a dual exponential basis. This section shows that these two algorithms for dual basis generation, i.e., by modulation and by polynomial expansion, reduce to producing the expected dual signals used in the DFT.

The finite complex exponential basis signals are denoted by $\{e_k[n]\}_{0:N-1}$, where

$$\{e_k[n]\}_{0:N-1} = \{e_0[n], \cdots, e_{N-1}[n]\}, \quad e_k[n] = e^{-j\frac{2\pi}{N}nk}, \ \ 0 \leq n \leq N-1. \tag{3.9.1}$$

In what follows, the dual basis signals for the complex exponential basis are derived in three ways. First, a straightforward argument using the standard inner product is used to uncover the structure of the dual basis signals from first principles. This is equivalent to understanding the change of basis undergone by the DFT, discussed in Section 3.6. Second, the algorithm for dual basis generation by modulation is used to derive the correct structure of the unnormalized modulating signals for the complex exponential basis. Finally, the algorithm for dual basis generation by polynomial expansion is shown to be consistent by producing the correct dual basis signal structure.

We have previously seen that $\{e_k[n]\}_{0:N-1}$ is an orthogonal basis. Therefore the dual signal structure is given by a scalar multiple of the complex conjugate of the original basis signal, specifically, denote the set of dual signals by $\{d_k[n]\}_{0:N-1}$ where

$$\{d_k[n]\}_{0:N-1} = \{d_0[n], \cdots, d_{N-1}[n]\}, \quad d_k[n] = \frac{1}{N}e^{j\frac{2\pi}{N}nk}, \ \ 0 \leq n \leq N-1. \tag{3.9.2}$$

By inspection, computing the dual basis by matrix inversion leads to the same result. Using Eq. (3.4.11) to generate the dual basis signals by matrix inversion inverts Φ^H when the structure of Φ is the canonical DFT matrix. One well known property of the DFT matrix is that it is a unitary operator when appropriately scaled by $\frac{1}{\sqrt{N}}$, thus the dual signals are a scaled version of the complex conjugates of the complex exponential basis signals.

Next, we find the unnormalized modulating signals, $\{\nu(e_k[n])\}_{0:N-1}$, that satisfy Eq. (3.4.21) specialized to the complex exponential basis. This is given by, for $0 \leq n \leq N-1$,

$$d_k[n] = \gamma_k \cdot e_k[n] \circ \nu(e_k[n]), \tag{3.9.3}$$

where $\gamma_k^{-1} = N$, for $0 \leq k \leq N-1$. To find the unnormalized modulating signal $\nu(e_k[n])$, for a fixed index k, the complex signal $\nu(e_k[n]) = x_1[n] + jx_2[n]$ is found such that $e_k^*[n] = e_k[n](x_1[n] + jx_2[n])$. Setting up the system of equations to determine $x_1[n]$ and $x_2[n]$ yield, for

$0 \leq n \leq N - 1$,

$$
\begin{bmatrix} \cos\left(\frac{2\pi}{N}nk\right) & \sin\left(\frac{2\pi}{N}nk\right) \\ -\sin\left(\frac{2\pi}{N}nk\right) & \cos\left(\frac{2\pi}{N}nk\right) \end{bmatrix} \begin{bmatrix} x_1\,[n] \\ x_2\,[n] \end{bmatrix} = \begin{bmatrix} \cos\left(\frac{2\pi}{N}nk\right) \\ \sin\left(\frac{2\pi}{N}nk\right) \end{bmatrix}. \tag{3.9.4}
$$

The solution to this system of equations is given by $x_1\,[n] = \cos^2\left(\frac{2\pi}{N}nk\right) - \sin^2\left(\frac{2\pi}{N}nk\right)$ and $x_2\,[n] = 2\sin\left(\frac{2\pi}{N}nk\right)\cos\left(\frac{2\pi}{N}nk\right)$. Converting to polar representation, $\nu\left(e_k\,[n]\right) = e^{j2\left(\frac{2\pi}{N}\right)nk}$. Thus the signal modulated against $e_k\,[n]$, for a fixed index k, is $\left(e_k\,[n]\right)^{-2}$. Therefore, for $0 \leq k \leq N - 1$,

$$
\nu(e_k\,[n]) = e^{j2\left(\frac{2\pi}{N}\right)nk}, \quad 0 \leq n \leq N - 1. \tag{3.9.5}
$$

By substituting Eq. (3.9.5) into Eq. (3.9.3), it is straightforward to verify that the method of generating the dual basis for a complex exponential basis using modulation is valid.

Finally, we show that the dual complex exponential basis signals, $\{d_k\,[n]\}_{0:N-1}$, that result from polynomial expansion are also generated from the dual basis generation algorithm by polynomial expansion. Solving for $U^{(l)}\left(z\right) = \mathcal{Z}\left\{u\left(e_l\,[n]\right)\right\}$, for a fixed index l, yields

$$
U^{(l)}\left(z\right) = \prod_{k=0,k\neq l}^{N-1}\left(1 - \left(\sigma_k\right)^{-1}z^{-1}\right) \tag{3.9.6}
$$

$$
= \prod_{k=0,k\neq l}^{N-1}\left(1 - e^{j\frac{2\pi}{N}(k)}z^{-1}\right) \tag{3.9.7}
$$

$$
= \sum_{k=0}^{N-1} e^{j\frac{2\pi}{N}k}z^{-k}. \tag{3.9.8}
$$

This procedure can also be verified using Vieta's formula as described in Section 3.4.4.

Chapter 4

Approximate Transient Spectral Analysis

In order to compute the DTT, as defined in Section 3.4, a number of practical issues related to numerical stability, transient resolution, and spectral leakage must be addressed. This chapter proposes various algorithms for computing approximate transient spectra which mitigate many of these issues. Of course, nothing is gained without consequence, and each algorithm presents a fundamental tradeoff between conflicting properties. For example, just as with harmonic spectral analysis, in order to reduce spectral leakage, a biased transient spectrum is developed in Section 4.1, however the approximate spectrum is no longer uniquely invertible. As another example, recall the algorithm for generating dual signals using polynomial expansion where the zeros of the defined polynomials were both real and strictly greater than one. The resulting dual signal coefficients inherently alternate in sign and grow numerically unstable as the size of the exponential basis is increased. Section 4.3 proposes an unconstrained optimization problem whose solution is a set of numerically stable approximate dual signals, even for large basis sizes. This stability is gained at the expense of limiting transient resolution. An approach for creating an orthogonal basis closest to an exponential basis in a least squares sense is presented in Section 4.6. Finally, an approximate transient spectrum for continuous-time transient signals in the form of Eq. (2.1.1) based upon orthogonal polynomials is considered in Section 4.7.

4.1 Parametric Biorthogonal Constraint Relaxations

The techniques proposed in this section utilize various parametric relaxations to the biorthogonality constraints resulting in a set of approximate dual signals, denoted by $\left\{\hat{\psi}_k\left[n\right]\right\}_{1:N}$. The approximate dual signals are then substituted into the DTT in place of the dual exponential signals, resulting in an approximate transient spectrum. Several reasons for relaxing the biorthogonality constraints exist: to decrease computational errors when computing the transient spectrum for large N, to improve transient resolution in the presence of additive noise, and/or to reduce spectral leakage. Two of these properties are highlighted for each relaxation: transient resolution and spectral leakage.

Three relaxations of the biorthogonality constraints are presented, where each relaxation is parameterized by a single parameter θ. This parameter is then restricted to a range of values for

which desirable spectral resolution guarantees can be made. Recall that the transient spectrum of a signal of the form

$$x_1[n] = \alpha_a(\sigma_a)^n, \quad 0 \le n \le N-1, \tag{4.1.1}$$

where $\sigma_a \notin \{\sigma_k\}_{1:N}$ and $\alpha_a \ne 0$ has been previously demonstrated to contain a broadened peak in the expansion coefficients corresponding to the decay rates nearest to σ_a. For signals with multiple exponential components, such spectral leakage could possibly interfere with the identification of other decay rates. Each relaxation addresses this issue using an example in which the approximate spectrum is computed for a signal consisting of two closely spaced exponential components. Enforcing spectral resolution requires that the two components produce two visually distinguishable peaks in the approximate spectrum. Though in general the decay rates of a signal being analyzed are unknown, this analysis will provide the ratio of amplitudes guaranteed to be resolvable for each biorthogonality relaxation as it relates to the parameter θ.

4.1.1 Uniform Relaxation

Define the approximate DTT analysis equation using a uniform relaxation of the biorthogonality constraints to be

$$\mathcal{A}_U[k] = \sum_{n=0}^{N-1} x[n]\hat{\psi}_{U,k}[n] = \left\langle x, \hat{\psi}_{U,k} \right\rangle, \quad 1 \le k \le N, \tag{4.1.2}$$

where the uniformly relaxed approximate dual basis, $\left\{\hat{\psi}_{U,k}[n]\right\}_{1:N}$, satisfies the uniformly relaxed biorthogonality constraints given by

$$\langle \phi_l, \hat{\psi}_{U,k} \rangle = \begin{cases} 1, & l = k \\ \theta_U, & l \ne k \end{cases}, \quad 1 \le k, l \le N, \tag{4.1.3}$$

for $\theta_U \in \mathcal{I}_U$, for some interval $\mathcal{I}_U$. Only positive values of θ_U are considered in limiting the domain of θ_U based upon spectral resolution. Towards this end, consider $x_2[n] = \alpha_1\phi_{k_1}[n] + \alpha_2\phi_{k_2}[n] = \alpha_1\sigma_{k_1}^n + \alpha_2\sigma_{k_2}^n$ with $\alpha_1, \alpha_2 > 0$, $|k_1 - k_2| > 1$ (non-adjacent), and $k_1, k_2 \in \{1, \dots N\}$. The uniformly relaxed approximate transient spectrum of $x_2[n]$ is given by

$$\mathcal{A}_U[k] = \alpha_1\langle \phi_{k_1}, \hat{\psi}_{U,k} \rangle + \alpha_2\langle \phi_{k_2}, \hat{\psi}_{U,k} \rangle \tag{4.1.4}$$

$$= \begin{cases} \alpha_1 + \alpha_2\theta_U, & k = k_1 \\ \alpha_2 + \alpha_1\theta_U, & k = k_2 \\ (\alpha_1 + \alpha_2)\theta_U, & \text{otherwise} \end{cases}. \tag{4.1.5}$$

Enforcing meaningful spectral peaks for the limiting case of two exponential components requires distinct peaks in the spectrum when $|k_2 - k_1| = 2$. Without loss of generality, assume that $k_2 > k_1$.

In order to resolve the two spectral peaks, θ_U is constrained such that the following conditions hold:

$$\mathcal{A}_U[k_1] > \mathcal{A}_U[k_1 - 1] \iff \alpha_1 + \alpha_2\theta_U > (\alpha_1 + \alpha_2)\theta_U$$
$$\mathcal{A}_U[k_1] > \mathcal{A}_U[k_1 + 1] \iff \alpha_1 + \alpha_2\theta_U > (\alpha_1 + \alpha_2)\theta_U$$
$$\mathcal{A}_U[k_1 + 2] > \mathcal{A}_U[k_1 + 1] \iff \alpha_2 + \alpha_1\theta_U > (\alpha_1 + \alpha_2)\theta_U$$
$$\mathcal{A}_U[k_1 + 2] > \mathcal{A}_U[k_1 + 3] \iff \alpha_2 + \alpha_1\theta_U > (\alpha_1 + \alpha_2)\theta_U.$$

Figure 4.1.1 depicts the resulting approximate spectrum. Specifically, Figure 4.1.1(a) and Figure 4.1.1(b) depict the approximate transient spectrum of the exponential components $\alpha_1\phi_{k_1}[n]$ and $\alpha_2\phi_{k_2}[n]$, respectively. Figure 4.1.1(c) depicts the result of adding the two previous spectra, equivalent to the approximate transient spectrum of the signal $x_2[n]$.

(a) Approximate DTT of the signal $\alpha_1\phi_{k_1}[n]$ using a uniform relaxation of the biorthogonality constraints

(b) Approximate DTT of the signal $\alpha_2\phi_{k_2}[n]$ using a uniform relaxation of the biorthogonality constraints

(c) Approximate DTT of the signal $x_2[n]$ using a uniform relaxation of the biorthogonality constraints

Figure 4.1.1: Spectral resolution of $x_2[n]$ using a uniform relaxation of the biorthogonality constraints in the limiting case

Taking the symmetry of the desired spectral constraints into account, the domain of θ_U is restricted to the interval $\mathcal{I}_U = (0, 1)$. Note that the value of θ_U is invariant to the amplitude of the transient spectral components to be resolved. This property is not true for the remaining two constraint relaxations. The manifestation of bias in this spectrum is evident in looking at, for example, $\mathcal{A}_U[k_1]$, where the spectral amplitude is $\alpha_1 + \alpha_2\theta_U$ and the contribution of $\alpha_2\theta_U$ is the result of spectral leakage. Note that the approximate spectrum is asymptotically unbiased, meaning the true transient spectrum is found as $\theta_U \to 0$.

The effect of the uniform relaxation of the biorthogonality constraints on spectral leakage for

an exponentially decaying signal where the decay rate does not match one of the exponential basis signals is now presented. Specifically, to illustrate this effect, define an exponential signal of the form $x_1[n] = \alpha_1(\sigma)^n$, $0 \leq n \leq N-1$, where $\sigma \notin \{\sigma_k\}_{1:N}$ and $\alpha_1 \neq 0$. The set of exponential basis signals $\{\phi_k[n]\}_{1:N}$ are used where the decay rates are uniformly spaced between 0.1 and 0.9 and we select $N = 17$. Additionally, let $x_1[n]$ have a decay rate of $\sigma = 0.38$. The resulting approximate transient spectrum, $\mathcal{A}_U[k]$ is shown in Figure (4.1.2) for $\theta_U = 0.05$ and $\theta_U = 0.2$. For comparison the transient spectrum using the DTT is shown as well. The magnitudes of the transient spectra are also included. As is demonstrated in the figure, the uniform relaxation of the biorthogonality constraints does not produce a significant advantage in terms of spectral leakage, and actually increases the spectral leakage at further decay rates from the decay rate present.
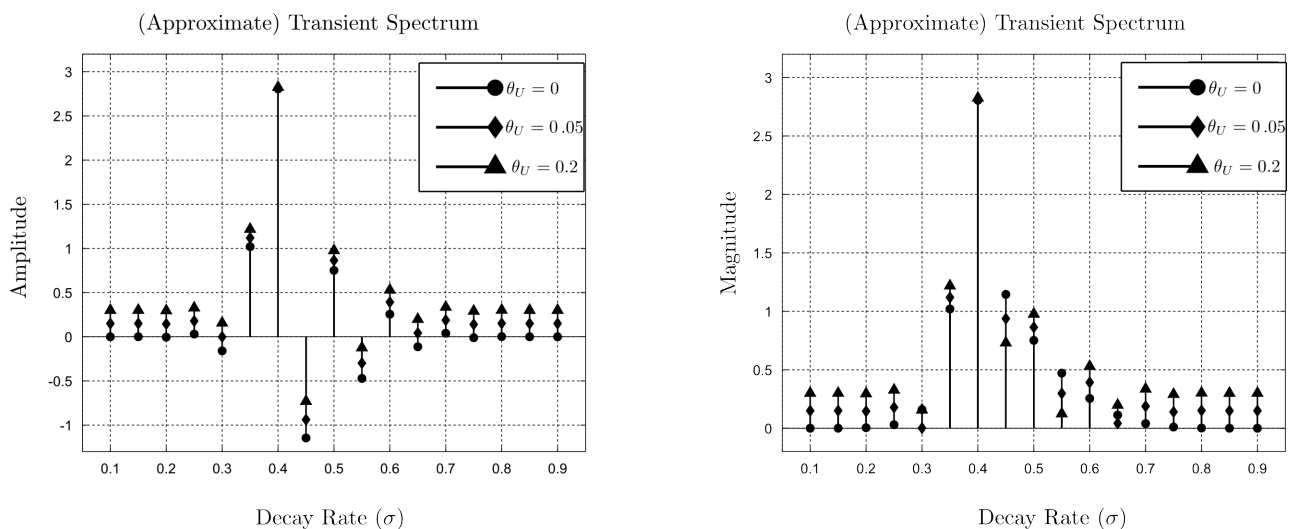
Figure 4.1.2: Approximate DTT of $x_1[n]$ using a uniform relaxation of the biorthogonality constraints for different values of θ

4.1.2　Linear Relaxation

Define the approximate DTT analysis equation using a linear relaxation of the biorthogonality constraints to be

$$\mathcal{A}_L[k] = \sum_{n=0}^{N-1} x[n]\,\hat{\psi}_{L,k}[n] = \left\langle x, \hat{\psi}_{L,k} \right\rangle, \quad 1 \leq k \leq N, \tag{4.1.6}$$

where the linearly relaxed approximate dual basis, $\left\{ \hat{\psi}_{L,k}[n] \right\}_{1:N}$, satisfies the linearly relaxed biorthogonality constraints given by

$$\langle \phi_l, \hat{\psi}_{L,k} \rangle = \begin{cases} 1, & l = k \\ \theta_L\left(\frac{1}{|l-k|}\right), & l \neq k \end{cases}, \quad 1 \leq k, l \leq N, \tag{4.1.7}$$

for $\theta_L \in \mathcal{I}_L$, for some interval $\mathcal{I}_L$. Only positive values of θ_L are considered in limiting the domain of θ_L based upon spectral resolution. Towards this end, consider $x_2[n] = \alpha_1 \phi_{k_1}[n] + \alpha_2 \phi_{k_2}[n] =$

$\alpha_1 \sigma_{k_1}^n + \alpha_2 \sigma_{k_2}^n$ with $\alpha_1, \alpha_2 > 0$, $|k_1 - k_2| > 1$ (non-adjacent), and $k_1, k_2 \in \{1, \dots N\}$. The linearly relaxed approximate transient spectrum of $x_2[n]$ is given by

$$
\begin{aligned}
\mathcal{A}_L[k] \;&=\; \alpha_1 \langle \phi_{k_1}, \hat{\psi}_{L,k} \rangle + \alpha_2 \langle \phi_{k_2}, \hat{\psi}_{L,k} \rangle \tag{4.1.8} \\[4pt]
&= \begin{cases}
\alpha_1 + \frac{\alpha_2}{|k_2 - k_1|} \theta_L, & k = k_1 \\[6pt]
\alpha_2 + \frac{\alpha_1}{|k_2 - k_1|} \theta_L, & k = k_2 \\[6pt]
\left(\frac{\alpha_1}{|k - k_1|} + \frac{\alpha_2}{|k - k_2|} \right) \theta_L, & \text{otherwise}
\end{cases} \tag{4.1.9}
\end{aligned}
$$

Enforcing meaningful spectral peaks for the limiting case of two exponential components requires distinct peaks in the spectrum when $|k_2 - k_1| = 2$. Without loss of generality assume that $k_2 > k_1$. In order to resolve the two spectral peaks, θ_L is constrained such that the following conditions hold:

$$
\mathcal{A}_L[k_1] > \mathcal{A}_L[k_1 - 1] \quad \Longleftrightarrow \quad \alpha_1 + \frac{1}{2}\alpha_2 \theta_L > \left(\alpha_1 + \frac{1}{3}\alpha_2 \right) \theta_L
$$

$$
\mathcal{A}_L[k_1] > \mathcal{A}_L[k_1 + 1] \quad \Longleftrightarrow \quad \alpha_1 + \frac{1}{2}\alpha_2 \theta_L > \left(\alpha_1 + \alpha_2 \right) \theta_L
$$

$$
\mathcal{A}_L[k_1 + 2] > \mathcal{A}_L[k_1 + 1] \quad \Longleftrightarrow \quad \alpha_2 + \frac{1}{2}\alpha_1 \theta_L > \left(\alpha_2 + \alpha_1 \right) \theta_L
$$

$$
\mathcal{A}_L[k_1 + 2] > \mathcal{A}_L[k_1 + 3] \quad \Longleftrightarrow \quad \alpha_2 + \frac{1}{2}\alpha_1 \theta_L > \left(\alpha_2 + \frac{1}{3}\alpha_1 \right) \theta_L.
$$

Figure 4.1.3 depicts the resulting approximate spectrum. Specifically, Figure 4.1.1(a) and Figure 4.1.1(b) depict the approximate transient spectrum of the exponential components $\alpha_1 \phi_{k_1}[n]$ and $\alpha_2 \phi_{k_2}[n]$, respectively. Figure 4.1.1(c) depicts the result of adding the two previous spectra, equivalent to the approximate transient spectrum of the signal $x_2[n]$.

(a) Approximate DTT of the signal $\alpha_1\phi_{k_1}[n]$ using a linear relaxation of the biorthogonality constraints

(b) Approximate DTT of the signal $\alpha_2\phi_{k_2}[n]$ using a linear relaxation of the biorthogonality constraints

(c) Approximate DTT of the signal $x_2[n]$ using a linear relaxation of the biorthogonality constraints

Figure 4.1.3: Spectral resolution of $x_2[n]$ using a linear relaxation of the biorthogonality constraints in the limiting case

The second spectral constraint restricts the domain to $\theta_L \in \left(0, \frac{2\alpha_1}{2\alpha_1+\alpha_2}\right)$ while the third spectral constraint, by symmetry, restricts the domain to $\theta_L \in \left(0, \frac{2\alpha_2}{2\alpha_2+\alpha_1}\right)$. Exploiting monotonicity and symmetry, the first and fourth spectral constraints are met if the second and third are. Therefore, the domain of θ_L is restricted to

$$\mathcal{I}_L = \left(0, \min\left(\frac{2\alpha_1}{2\alpha_1 - \alpha_2}, \frac{2\alpha_2}{2\alpha_2 - \alpha_1}\right)\right).$$

For example, if $\alpha_1 = 5\alpha_2$ then $\mathcal{I}_L \approx (0, 0.2857)$, however if $\alpha_1 = \alpha_2$ then $\mathcal{I}_L = (0, 2)$. The manifestation of bias in this spectrum is evident in looking at, for example, $\mathcal{A}_L[k_1]$, where the spectral amplitude is $\alpha_1 + \frac{1}{2}\alpha_2\theta_L$ and the contribution of $\frac{1}{2}\alpha_2\theta_L$ is the result of spectral leakage. Note that the approximate transient spectrum is asymptotically unbiased, meaning the true transient spectrum is found as $\theta_L \to 0$.

The effect of the linear relaxation of the biorthogonality constraints on spectral leakage for an exponentially decaying signal where the decay rate does not match one of the exponential basis signals is now presented. Specifically, to illustrate this effect, define an exponential signal of the form $x_1[n] = \alpha_1(\sigma)^n$, $0 \le n \le N - 1$, where $\sigma \notin \{\sigma_k\}_{1:N}$ and $\alpha_1 \ne 0$. The set of exponential basis signals $\{\phi_k[n]\}_{1:N}$ are used where the decay rates are uniformly spaced between 0.1 and 0.9 and we select $N = 17$. Additionally, $x_1[n]$ has a decay rate of $\sigma = 0.38$. The resulting approximate

transient spectrum $\mathcal{A}_L[k]$ is shown in Figure (4.1.4) for $\theta_L = 0.05$ and $\theta_L = 0.2$. For comparison the transient spectrum using the DTT is shown as well. The magnitudes of the transient spectra are also included. As is demonstrated in the figure, the linear relaxation of the biorthogonality constraints does not produce a significant advantage in terms of spectral leakage, and actually increases the spectral leakage at further decay rates from the decay rate present.
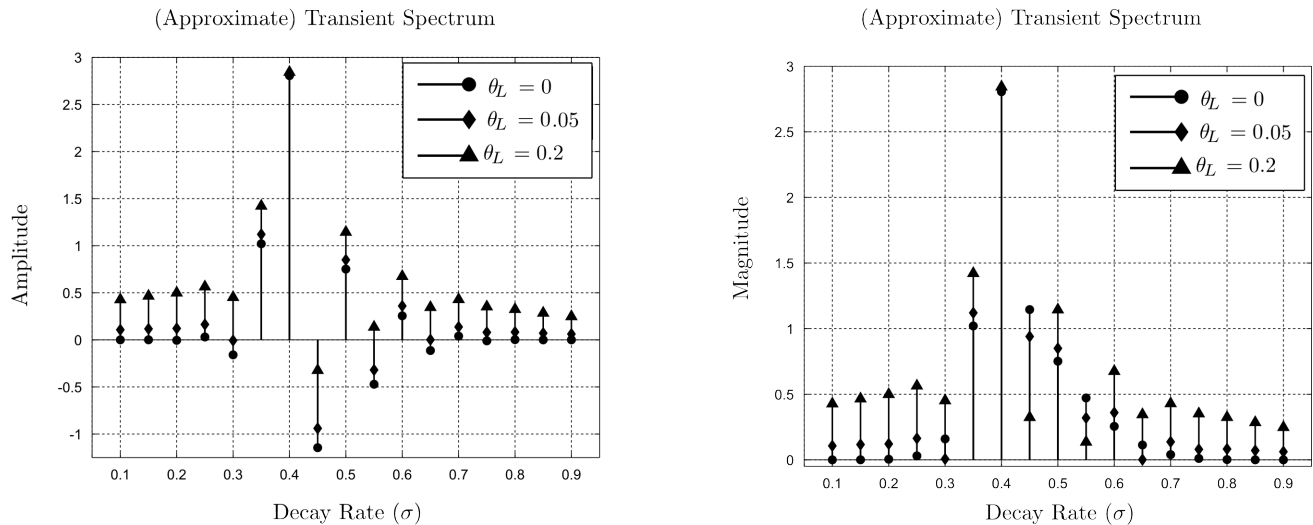
Figure 4.1.4: Approximate DTT of $x_1[n]$ using a linear relaxation of the biorthogonality constraints for different values of θ

4.1.3   Exponential Relaxation

Define the approximate DTT analysis equation using an exponential relaxation of the biorthogonality constraints to be

$$\mathcal{A}_E[k] \;\;=\;\; \sum_{n=0}^{N-1} x[n]\,\hat{\psi}_{E,k}[n] = \left\langle x,\, \hat{\psi}_{E,k} \right\rangle, \quad 1 \le k \le N, \qquad (4.1.10)$$

where the exponentially relaxed approximate dual basis, $\left\{ \hat{\psi}_{E,k}[n] \right\}_{1:N}$, satisfies the exponentially relaxed biorthogonality constraints given by

$$\langle \phi_l,\, \hat{\psi}_{E,k} \rangle = \begin{cases} 1, & l = k \\ \theta_E^{|l-k|}, & l \ne k \end{cases}, \quad 1 \le k, l \le N, \qquad (4.1.11)$$

for $\theta_E \in \mathcal{I}_E$, for some interval $\mathcal{I}_E$. Only positive values of θ_E are considered in limiting the the domain of θ_E based upon spectral resolution. Towards this end, consider $x_2[n] = \alpha_1 \phi_{k_1}[n] + \alpha_2 \phi_{k_2}[n] = \alpha_1 \sigma_{k_1}^n + \alpha_2 \sigma_{k_2}^n$ with $\alpha_1, \alpha_2 > 0$, $|k_1 - k_2| > 1$ (non-adjacent), and $k_1, k_2 \in \{1, \dots N\}$.

The exponentially relaxed approximate transient spectrum of $x_2[n]$ is given by

$$
\mathcal{A}_E[k] = \alpha_1\langle\phi_{k_1}, \hat{\psi}_{E,k}\rangle + \alpha_2\langle\phi_{k_2}, \hat{\psi}_{E,k}\rangle \tag{4.1.12}
$$

$$
= \begin{cases} \alpha_1 + \alpha_2\theta_E^{|k_2-k_1|}, & k = k_1 \\ \alpha_2 + \alpha_1\theta_E^{|k_2-k_1|}, & k = k_2 \\ \alpha_1\theta_E^{|k-k_1|} + \alpha_2\theta_E^{|k-k_2|}, & \text{otherwise} \end{cases} . \tag{4.1.13}
$$

Enforcing meaningful spectral peaks for the limiting case of two exponential components requires distinct peaks in the spectrum when $|k_2 - k_1| = 2$. Without loss of generality, assume that $k_2 > k_1$. In order to resolve the two spectral peaks, θ_E is constrained such that the following conditions hold:

$$
\mathcal{A}_E[k_1] > \mathcal{A}_E[k_1 - 1] \iff \alpha_1 + \alpha_2\theta_E^2 > (\alpha_1 + \alpha_2\theta_E^2)\theta_E
$$
$$
\mathcal{A}_E[k_1] > \mathcal{A}_E[k_1 + 1] \iff \alpha_1 + \alpha_2\theta_E^2 > (\alpha_1 + \alpha_2)\theta_E
$$
$$
\mathcal{A}_E[k_1 + 2] > \mathcal{A}_E[k_1 + 1] \iff \alpha_2 + \alpha_1\theta_E^2 > (\alpha_1 + \alpha_2)\theta_E
$$
$$
\mathcal{A}_E[k_1 + 2] > \mathcal{A}_E[k_1 + 3] \iff \alpha_2 + \alpha_1\theta_E^2 > (\alpha_2 + \alpha_1\theta_E^2)\theta_E.
$$

Figure 4.1.5 depicts the resulting approximate spectrum. Specifically, Figure 4.1.5(a) Figure 4.1.5(b) depict the approximate transient spectrum of the exponential components $\alpha_1\phi_{k_1}[n]$ and $\alpha_2\phi_{k_2}[n]$, respectively. Figure 4.1.5(c) depicts the result of adding the two previous spectra, equivalent to the approximate transient spectrum of the signal $x_2[n]$.

(a) Approximate DTT of the signal $\alpha_1 \phi_{k_1}[n]$ using an exponential relaxation of the biorthogonality constraints

(b) Approximate DTT of the signal $\alpha_2 \phi_{k_2}[n]$ using an exponential relaxation of the biorthogonality constraints

(c) Approximate DTT of the signal $x_2[n]$ using an exponential relaxation of the biorthogonality constraints

Figure 4.1.5: Spectral resolution of $x_2[n]$ using an exponential relaxation of the biorthogonality constraints in the limiting case

The second spectral constraint restricts the domain to

$$\theta_E \in \left\{ \left(0, \min\left(1, \frac{\alpha_1}{\alpha_2}\right)\right) \cup \left(\max\left(1, \frac{\alpha_1}{\alpha_2}\right), \infty\right) \right\}. \tag{4.1.14}$$

By symmetry, the third spectral constraint restricts the domain to

$$\theta_E \in \left\{ \left(0, \min\left(1, \frac{\alpha_2}{\alpha_1}\right)\right) \cup \left(\max\left(1, \frac{\alpha_2}{\alpha_1}\right), \infty\right) \right\}. \tag{4.1.15}$$

The first spectral constraint factors into the form $(1 - \theta_E)\left(\alpha_1 + \alpha_2 \theta_E^2\right) > 0$ therefore it restricts $\theta_E \in (0, 1)$. By symmetry the fourth constraints make the same restriction. Therefore we define the domain of θ_E to be $\mathcal{I}_E = \left(0, \min\left(\frac{\alpha_1}{\alpha_2}, \frac{\alpha_2}{\alpha_1}\right)\right)$. The manifestation of bias in this spectrum is evident in looking at, for example, $\mathcal{A}_E[k_1]$, where the spectral amplitude is $\alpha_1 + \alpha_2 \theta_E^2$ and the contribution of $\alpha_2 \theta_E^2$ is the result of spectral leakage. Note that the approximate transient spectrum is asymptotically unbiased, meaning the true transient spectrum is found as $\theta_E \to 0$.

The effect of the exponential relaxation of the biorthogonality constraints on spectral leakage for an exponentially decaying signal where the decay rate does not match one of the exponential basis signals is now presented. Specifically, we illustrate this effect using an exponential signal of the form $x_1(n) = \alpha_1 \sigma^n$, $0 \leq n \leq N - 1$, where $\sigma \notin \{\sigma_k\}_{1:N}$ and $\alpha_1 \in \mathbb{R}$. The set of exponential basis signals $\{\phi_k[n]\}_{1:N}$ are used where the decay rates are uniformly spaced between 0.1 and 0.9

and we select $N = 17$. Additionally, $x_1[n]$ has a decay rate of $\sigma = 0.38$. The resulting approximate transient spectrum, $\mathcal{A}_E[k]$, is shown in Figure 4.1.6 for $\theta_E = 0.05$ and $\theta_E = 0.2$. For comparison the transient spectrum using the DTT is shown as well. The magnitude of the transient spectra are also included. As is demonstrated in the figure, the spectral peaks closest to decay rate σ are improved (heightened) and the spectral values of other decay rates are generally decreased, especially the values of high leakage, which is advantageous. This parametric technique is the only one with a clear advantage with respect to spectral leakage.
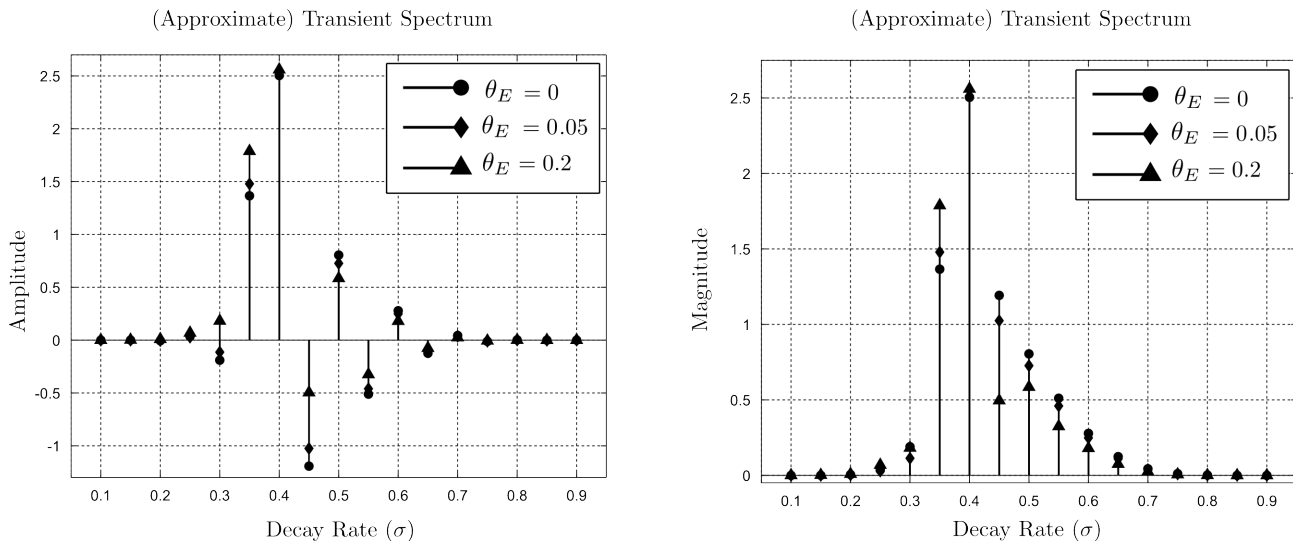
Figure 4.1.6: Approximate DTT of $x_1[n]$ using an exponential relaxation of the biorthogonality constraints for different values of θ

4.1.4 Implementation

The previous three subsections contained examples of approximate transient spectra computed using various approximate dual signals without discussing how these signals were obtained. The algorithms for generating dual exponential bases in Chapter 3 can easily be extended to generate these signals using only an additional matrix multiplication. Let the columns of the matrix $\hat{\Psi}$ represent the approximate dual signals $\{\hat{\psi}_k[n]\}_{1:N}$ and denote $G(\theta)$ as one of the relaxed biorthogonality constraint matrices given in Table 4.1, where θ is the design parameter. Then, the relaxed biorthogonality constraints are written as

$$\Phi\hat{\Psi}^H = G(\theta) \quad \text{or} \quad \langle \phi_i, \hat{\psi}_j \rangle = [G(\theta)]_{ij}. \tag{4.1.16}$$

Therefore the approximate dual signals are the columns of $\hat{\Psi}$ given by

$$\hat{\Psi} = \left(\Phi^H\right)^{-1} G(\theta). \tag{4.1.17}$$

Any of the algorithms proposed in Chapter 3 for computing $(\Phi^H)^{-1}$ may be used. The main result of this section is not necessarily the ability to produce these approximate dual signals, but to use

the spectral resolution guarantee equations to understand the spectral resolution of a stable set of dual signals produced through a well conditioned algorithm. An example of this use is presented in Section 4.3.

Uniform Relaxation	Linear Relaxation	Exponential Relaxation
$[G_U(\theta)]_{ij} = \begin{cases} 1, & i = j \\ \theta_U, & i \neq j \end{cases}$	$[G_L(\theta)]_{ij} = \begin{cases} 1, & i = j \\ \theta_L\left(\frac{1}{\|i-j\|}\right), & i \neq j \end{cases}$	$[G_E(\theta)]_{ij} = \begin{cases} 1, & i = j \\ \theta_E^{\|i-j\|}, & i \neq j \end{cases}$

Table 4.1: Parametric relaxation structures for the biorthogonality constraints

4.2   The Equal Energy Exponential Basis

In developing the parametric biorthogonal constraint relaxations in Sections 4.1.1, 4.1.2, and 4.1.3, the parameters α_1 and α_2 in $x_2[n]$ are interpreted as scalings of the amplitudes of two exponential basis signals where each basis signal contains a different amount of energy. That is, the derivation for spectral resolution used an exponential basis where each of the exponential basis signals $\phi_k[n]$ has equal initial amplitude, i.e., $\phi_k[0] = 1$ for $1 \leq k \leq N$. This section extends all of the results for transient spectra thus far using the equal initial amplitude exponential basis, $\{\phi_k[n]\}_{1:N}$, to an exponential basis where each basis signal has equal energy.

The energy in an exponential basis signal $\phi_k[n]$, for a fixed index k, is

$$\mathcal{E}(\phi_k[n]) = \sum_{n=0}^{N-1} (\sigma_k^n)^2 \tag{4.2.1}$$

$$= \frac{\sigma_k^{2N} - 1}{\sigma_k^2 - 1}. \tag{4.2.2}$$

Therefore, the energy in the basis signal monotonically increases as σ_k approaches 1. In order to create a set of exponential basis signals with equal energy, define the equal energy exponential basis, $\left\{\tilde{\phi}_k[n]\right\}_{1:N}$, as

$$\left\{\tilde{\phi}_k[n]\right\}_{1:N} = \left\{\frac{1}{\sqrt{\mathcal{E}(\phi_k[n])}}\phi_k[n]\right\}_{1:N}. \tag{4.2.3}$$

The equal energy exponential basis signals are each a scaled version of the equal initial amplitude exponential signals using the scalar $\tilde{\alpha}_k = 1/\sqrt{\mathcal{E}(\phi_k[n])}$. This definition allows a straightforward extension of the results for an equal initial amplitude exponential basis to the equal energy exponential basis. For example, the bounds on transient resolution derived in the previous section can be extended to the equal energy exponential basis by incorporating the scalars $\{\tilde{\alpha}_k\}_{1:N}$ into the expansion coefficients, consequently adjusting the range of resolution.

4.3 Unconstrained Optimization Formulation

The approximate transient spectra derived in Section 4.1 utilize approximate dual basis signals in the DTT which are capable of larger basis sizes than the dual exponential basis before numerical errors cause noticeable distortion in the resulting spectra. However, the procedure for generating these bases required first generating the dual exponential bases. Therefore, for large values of N, the numerical stability gained by using these approximate dual signals is lost. The algorithm proposed in this section generates a numerically stable approximate dual exponential basis as the solution to an unconstrained optimization problem. Again, in order to gain numerical stability, a trade-off with spectral resolution is made. This algorithm uses a non-parametric approximation of the biorthogonality constraints and includes an approximate enforcement of orthogonality between the approximate dual signals. To accomplish this, an optimization problem is formulated by trading off between these two forms of orthogonality.

4.3.1 Quadratic Penalty Formulation

Given the exponential basis $\{\phi_k[n]\}_{1:N}$, the inner product structure of the dual basis is un-constrained, meaning the dual signals do not satisfy both the biorthogonality and orthogonality constraints simultaneously. That is, for $1 \leq i,\ k \leq N$, the $2N^2$ constraints

$$\langle \psi_i, \psi_k \rangle = \delta_{ik} \qquad \text{(orthogonality)} \tag{4.3.1}$$

$$\langle \phi_i, \psi_k \rangle = \delta_{ik} \qquad \text{(biorthogonality)} \tag{4.3.2}$$

have no solution for $\{\psi_k[n]\}_{1:N}$. We then settle for finding an approximate dual basis that is closest to meeting these constraints, for some definition of closest. It has been observed experimentally that approximate orthogonality of the dual exponential signals significantly improves the behavior of the transient spectrum when analyzing a signal that contains additive noise, e.g., $\bar{x}_d[n]$. In order to relax these constraints, a quadratic penalty formulation[1] is proposed, i.e., instead of enforcing the constraints directly, we penalize solutions that do not meet them and optimize over the space consisting of all sets of N length N signals to find the set of approximate dual signals that are closest to meeting the constraints for the given penalization values. [4] This yields an optimization problem given by

$$\{\check{\psi}_k[n]\}_{1:N} \in \arg \min_{\{\psi_k[n]\}_{1:N}} \varrho_1 \sum_{i=1}^{N} \sum_{k=1}^{N} \left[\Lambda^{(1)}\right]_{ik} \cdot \left(\langle \phi_i, \psi_k \rangle - \delta_{ik}\right)^2$$

$$+ \varrho_2 \sum_{i=1}^{N} \sum_{k=1}^{N} \left[\Lambda^{(2)}\right]_{ik} \cdot \left(\langle \psi_i, \psi_k \rangle - \delta_{ik}\right)^2, \tag{4.3.3}$$

---

[1]A quadratic penalty formulation is also referred to in some texts as a Lagrangian method, while in other texts a Lagrangian method is defined by only a linear penalty.

where the optimal approximate dual signals are denoted by $\left\{\check{\psi}_k[n]\right\}_{1:N}$.

This formulation allows an emphasis to be placed on meeting a specific constraint with respect to the other constraints using the $2N^2$ Lagrangian multipliers, $\left\{\Lambda^{(k)}\right\}_{1:2}$. Even for reasonable values of N, the number of penalty coefficients required is quite large and a parametric penalty structure is often useful, but not necessary. A few simple parametric structures are listed in Table 4.2, where each penalty structure is parametrized by the design parameter λ. In each of the structures given, selecting $0 < \lambda < 1$ emphasizes penalizing deviations from the constraints given by $\langle \phi_k, \psi_k \rangle = \langle \psi_k, \psi_k \rangle = 1$, for $1 \le k \le N$. The coefficients $\{\varrho_k\}_{1:2}$ are used to emphasize the set of biorthogonality constraints with respect to the set of orthogonality constraints.

Equal Penalty	Linearly Decreasing Penalty	Exponentially Decreasing Penalty
$\left[\Lambda_{Eq}^{(k)}\right]_{ij} = \begin{cases} 1, & i = j \\ \lambda, & i \neq j \end{cases}$	$\left[\Lambda_{L}^{(k)}\right]_{ij} = \begin{cases} 1, & i = j \\ \lambda^{\frac{1}{|i-j|}}, & i \neq j \end{cases}$	$\left[\Lambda_{E}^{(k)}\right]_{ij} = \begin{cases} 1, & i = j \\ \lambda^{|i-j|}, & i \neq j \end{cases}$

Table 4.2: Structures of Lagrangian coefficients for the biorthogonal and orthogonal constraints in the quadratic penalty optimization formulation

The examples presented in this thesis using the solution to the optimization problem in Eq. (4.3.3) use a non-linear conjugate gradient algorithm to find the approximate dual basis. This algorithm is briefly described in Section 4.3.2.

Let the objective function in Eq. (4.3.3) be defined by $\mathcal{F} = \mathcal{C}_1 + \mathcal{C}_2$, where $\mathcal{C}_1$ represents the N^2 biorthogonality constraints, with associated Lagrangian multipliers $\Lambda^{(1)}$, given by

$$\mathcal{C}_1 = \varrho_1 \sum_{i=1}^{N} \sum_{k=1}^{N} \left[\Lambda^{(1)}\right]_{ik} \cdot \left(\langle \phi_i, \psi_k \rangle - \delta_{ik}\right)^2, \tag{4.3.4}$$

and where $\mathcal{C}_2$ represents the N^2 orthogonality constraints, with associated Lagrangian multipliers $\Lambda^{(2)}$, and is given by

$$\mathcal{C}_2 = \varrho_2 \sum_{i=1}^{N} \sum_{k=1}^{N} \left[\Lambda^{(2)}\right]_{ik} \cdot \left(\langle \psi_i, \psi_k \rangle - \delta_{ik}\right)^2. \tag{4.3.5}$$

Any optimization technique utilizing gradient descent can be used to at least find a local minimum of this formulation. Gradient descent algorithms require the gradient of the objective function, $\nabla \mathcal{F}$, to move from one solution to the next. The derivative of the objective function, with respect to one approximate dual signal $\psi_k[n]$, for a fixed index k, is given by

$$\frac{\partial}{\partial \psi_k} \mathcal{F} = \frac{\partial}{\partial \psi_k} \mathcal{C}_1 + \frac{\partial}{\partial \psi_k} \mathcal{C}_2 \tag{4.3.6}$$

where the derivatives with respect to a single signal ψ_k are given by

$$\frac{\partial}{\partial \psi_k} \mathcal{C}_1 \;\; = \;\; \varrho_1 \sum_{i=1}^{N} 2 \left[\Lambda^{(1)} \right]_{ik} \left(\langle \phi_i, \, \psi_k \rangle - \delta_{ik} \right) \cdot \phi_i \tag{4.3.7}$$

and

$$\frac{\partial}{\partial \psi_k} \mathcal{C}_2 = \varrho_2 \left(\sum_{i=1, \, i \neq k}^{N} 2 \left[\Lambda^{(2)} \right]_{ik} \left(\langle \psi_i, \, \psi_k \rangle - \delta_{ik} \right) \cdot \psi_i \; + \; 4 \left[\Lambda^{(2)} \right]_{kk} \left(\langle \psi_k, \, \psi_k \rangle - \delta_{kk} \right) \cdot \psi_k \right). \tag{4.3.8}$$

The following subsection details the non-linear conjugate gradient optimization algorithm used in this thesis. For a detailed treatment of this, or other possible gradient descent methods, see [4, 3, 13].

4.3.2 Conjugate Gradient Method

A simple unconstrained formulation of the non-linear method of Conjugate Gradient (CG) optimization is to solve the problem given by

$$\hat{x} \in \arg \min_{x} \mathcal{F}(x) \tag{4.3.9}$$

where $\mathcal{F}(x)$ is the objective function. When the form of $\mathcal{F}(x) = \|Ax - b\|_2^2$ then a CG algorithm attempts to find a solution to the system of equations $A^T A x = A^T b$, which may be found from either taking the gradient and solving, i.e., $\nabla_x \mathcal{F}(x) = 2A^T(Ax - b) = 0$, or from the orthogonality principle of least squares.

The non-linear CG algorithm used in this thesis is presented in summary next, a more thorough treatment may be found in [1]. Given a multi-objective function $\mathcal{F}(x)$, where x represents N length N signals, an initial feasible solution x_0 is assumed. In this thesis, the initial solution used is the dual exponential basis as found by polynomial expansion. The initial step is to move to a solution in the direction of steepest descent, i.e. negative of the gradient of $\mathcal{F}(x)$. This direction is given by

$$\Delta x_0 = -\nabla \mathcal{F}(x_0) \tag{4.3.10}$$

where the partial derivatives of $\mathcal{F}(x)$ with respect to a single dual signal is given in Eqs. (4.3.7) and (4.3.8). Next, the size of the step in this direction is computed by a line search of the form

$$\alpha_0 = \arg \min_{\alpha, \, \alpha \in \mathcal{R}} \mathcal{F}(x_0 + \alpha \Delta x_0) \tag{4.3.11}$$

where α_0 is the step size and the region $\mathcal{R}$ is selected using the Armijo Rule. The current solution is updated to x_1 by the formula

$$x_1 = x_0 + \alpha_0 \Delta x_0. \tag{4.3.12}$$

Once the first step has been taken, an iterative procedure begins where all following directions are conjugate directions, denoted by s_n. Note that $s_0 = \Delta x_0$. This iterative procedure is described in

Algorithm 4.1. The procedure is stopped when the change in the objective function is less than a predetermined threshold in making the next move.

Algorithm 4.1 Update step of the non-linear conjugate gradient algorithm

For each iteration t:

1. Find the direction of steepest descent: $\Delta x_t = -\nabla \mathcal{F}(x_t)$.

2. Determine the value β_t by the Fletcher-Reeves formula: $\beta_t = \frac{\Delta x_t^T \Delta x_t}{\Delta x_{t-1}^T \Delta x_{t-1}}$

3. Update the conjugate direction s_t: $s_t = \Delta x_t + \beta_t s_{t-1}$

4. Perform a line search to determine the step size α_t: $\alpha_t = \arg \min_{\alpha,\, \alpha \in \mathcal{R}} f(x_t + \alpha s_t)$

5. Move to the new solution: $x_{t+1} = x_t + \alpha_t s_t$.

4.4 Bounds on Transient Resolution

The tradeoff made in Section 4.1 resulted in a reduction of the space of transient signals for which spectral resolution is guaranteed. The derivation of the spectral resolution guarantees relied upon the symmetric properties of the relaxed biorthogonality constraints. The solution to the optimization problem in Eq. (4.3.3) results in a set of approximate dual signals, $\{\check{\psi}_k[n]\}_{1:N}$, for which the resulting biorthogonality structure cannot be guaranteed to be symmetric. This is because the feasible space includes all possible sets of N signals where each signal is of length N. If we denote the standard inner product of $\phi_i[n]$ with $\check{\psi}_j[n]$ as $[G_L]_{ij}$, then, for $i \neq j$, we have, in general,

$$\langle \phi_i, \check{\psi}_j \rangle = [G_L]_{ij} \neq [G_L]_{ji} = \langle \phi_j, \check{\psi}_i \rangle. \tag{4.4.1}$$

As a result, the biorthogonality structure resulting from $\{\check{\psi}_k[n]\}_{1:N}$ is generally unlike the biorthogonality structures defined in Sections 4.1.1, 4.1.2, and 4.1.3, for which simple expressions of spectral resolution guarantees were found. However, the analysis tools and expressions developed in those sections allow bounds on spectral resolution to be given for any approximate dual basis, including $\{\check{\psi}_k[n]\}_{1:N}$. The remainder of this section describes a procedure to provide spectral resolution guarantees for any solution to the optimization problem in Eq. (4.3.3).

Define G_B to be structured as either the equal penalty, the linearly decreasing penalty, or the exponentially decreasing penalty, as defined in Table 4.2, where λ_B was previously a design

parameter. Instead, we now solve for λ_B such that

$$\lambda_B = \arg\min_{\lambda} \lambda \tag{4.4.2}$$

$$\text{s.t. } [G_B(\lambda)]_{ij} \geq \left|[G_L]_{ij}\right| \tag{4.4.3}$$

for $1 \leq i, j, \leq N$. The value of λ_B that minimizes the right hand side of Eq. (4.4.2) provides an indication of what range of transient coefficients are guaranteed to be resolved in the same way that θ did in the aforementioned three sections. Often the spectral resolution bounds given by λ_B are only tight for a few values of G_L. This means that signals with larger ratios of transient coefficients are sometimes able to be resolved, but the guarantee of such resolution is not possible. Therefore, the bound on λ_B is typically due to restrictions between only a few specific neighboring decay rates. Consider the simple example when $\lambda_B = 0$ is found. Then, the resulting approximate dual signals correspond to the exact dual signals. This solution is found when $\varrho_2 = 0$.

Next, a tractable example is considered for which the unconstrained optimization formulation is solved using parameters: $\{\sigma_k\}_{1:5} = \{0.1, 0.3, 0.5, 0.7, 0.9\}$, $\varrho_1 = 10^4$, the equal penalization structure is used for the biorthogonal constraints with $\lambda = 10$, $\varrho_2 = 1$, and the equal penalization structure is used for the pairwise orthogonality constraints with $\lambda = \frac{1}{2}$. The initial solution chosen is the exact dual basis, which was computed using the algorithm in Section 3.4.4. The numerical results of the inner product structure $[G_L]_{ij} = \langle \phi_i, \check{\psi}_j \rangle$ are displayed in Table 4.3.

1	0.32734	-0.17458	-0.058744	0.02121
0.21947	1	0.70888	-0.10421	-0.003078
-0.08127	0.59807	1	0.50362	-0.082594
-0.050135	-0.1401	0.67168	1	0.093969
0.027698	-0.0063279	-0.15898	0.14038	1

Table 4.3: Inner product structure of the approximate dual signals from the quadratic penalty optimization formulation

The inner product structure in this example is bounded from above, in absolute value, by an exponential bound with $\lambda_B = 0.7088$, i.e., $(0.70888)^{|i-j|} \geq \left|\langle \phi_i, \check{\psi}_j \rangle\right|$ for $1 \leq i, j \leq 5$. From Section 4.1.3 we have that $0.70888 \in \left(0, \min\left(\frac{\alpha_1}{\alpha_2}, \frac{\alpha_2}{\alpha_1}\right)\right)$. Without loss of generality assume that $\alpha_2 > \alpha_1$. Using this exponential bound on the inner product structure given in Table 4.3, transient signals with amplitudes related up to $\alpha_2 = 1.4107 \cdot \alpha_1$ are guaranteed to be resolved using the approximate dual basis $\{\check{\psi}_k[n]\}_{1:5}$. Notice that the exponential bound is only tight for a few values in the table, and for some it is a large over-estimate, meaning signals of greater amplitudes may be resolved. However, signals with transient coefficients larger than this ratio cannot be guaranteed to be resolved.

4.5 Visualization of Approximate Dual Signals

This chapter has proposed several algorithms for generating approximate dual exponential sig-

nals for use in the DTT. For each algorithm, a fundamental tradeoff between conflicting properties was identified and then discussed. In order to demonstrate the numerical instability of the dual exponential basis, as well as the stability of the approximate dual basis from solving the optimization problem in Eq. (4.3.3), Figure 4.5.1 depicts the exponential basis signals, dual exponential signals, approximate dual signals using a linear relaxation of the biorthogonality constraints, and the approximate dual signals found from the unconstrained optimization formulation. Each axis of this figure depicts one dimension of the corresponding length $N = 3$ vectors. The exponential basis is defined using the decay rates $\{\sigma_k\}_{1:3} = \{0.3, 0.5, 0.8\}$. The numerical instability of the dual basis is clearly evident, even given the low model complexity. Note that the approximate dual basis produced by the linear relaxation to the biorthogonality constraints appear to be less numerically unstable than the dual basis, but still relatively unstable. The approximate dual signals resulting from the unconstrained optimization formulation are both numerically stable and seen to be approximately orthogonal to the exponential basis signals as well as to one another, as desired.
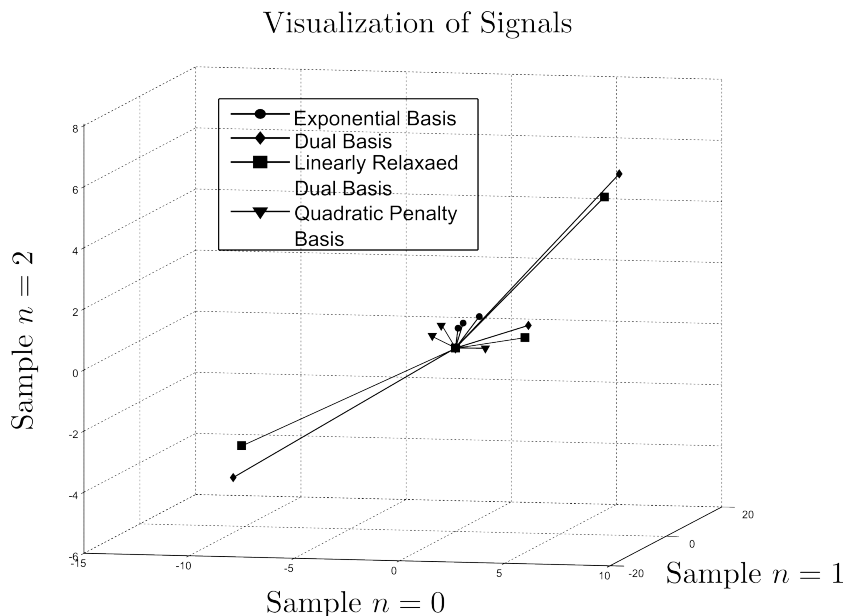
Figure 4.5.1: Visualization of the numerical instability of various dual bases

4.6   Inner Product Shaping

Inner product shaping is used to create an orthogonal basis, under the standard inner product, using which we may compute an approximate transient spectrum. Specifically, the result of inner product shaping on the exponential basis $\{\phi_k [n]\}_{1:N}$ is an orthogonal basis, denoted $\{\bar{\psi}_k [n]\}_{1:N}$, where the orthogonal basis signals are chosen to minimize the distance from exponential basis signals in a least squares sense. For a more thorough and general treatment of inner product shaping see [7]. The following discussion is specialized to an exponential basis.

The inner product shaping formulation considered in this thesis is given by

$$\left\{\bar{\psi}_l[n]\right\}_{1:N} = \arg \min_{\{\psi_k[n]\}_{1:N}} \sum_{k=1}^{N} \langle \phi_k - \psi_k, \phi_k - \psi_k \rangle \tag{4.6.1}$$

$$\text{s.t.} \quad \langle \psi_k, \psi_l \rangle = \delta_{kl}.$$

An equivalent formulation of the inner product shaping problem, using matrices, results in

$$\bar{\Psi} = \arg \min_{\Psi} \operatorname{tr}\left((\Phi - \Psi)^H (\Phi - \Psi) \right) \tag{4.6.2}$$

$$\text{s.t.} \quad \Psi^H \Psi = I_N.$$

Note that approximately enforcing orthogonality of the approximate dual signals was observed in Section 4.3 to improve the transient spectra found when the signal to be analyzed contains noise. The resulting signals, $\left\{\bar{\psi}_k[n]\right\}_{1:N}$, are known as the Orthonormal Least Squares Vectors (OLSV).

In order to solve the formulation in Eq. (4.6.2), we perform a unitary change of basis using the linear map U such that

$$\Phi_U = \Phi U \quad \text{and} \quad \Psi_U = \Psi U. \tag{4.6.3}$$

Note that unitary matrices preserve inner products, which can be seen by definition: $\langle Ux, Uy \rangle = (Uy)^H U x = y^H U^H U x = \langle x, y \rangle$. Therefore, the constraints in Eq. (4.6.2) become $\Psi_U^H \Psi_U = I_N$. Further, it is straight forward to verify that the objective function is equivalent to $\operatorname{tr}\left((\Phi_U - \Psi_U)^H (\Phi_U - \Psi_U) \right)$. We shall denote the optimal Ψ_U and Φ_U as $\hat{\Psi}_U$ and $\hat{\Phi}_U$, respectively. The following relation then holds to convert the optimal solution back from this new basis,

$$\hat{\Psi} = \hat{\Psi}_U U^H \tag{4.6.4}$$

The SVD may be used in order to find a suitable unitary change of basis. Denote the SVD of Φ by $\Phi = W \Pi V^H$ where both W and V are unitary matrices. By selecting $\Phi_U = \Phi V = W \Pi$, we find that the unitary map U in Eq. (4.6.3) is equivalent to V in the SVD of Φ.

Now that we have constructed Φ_U, we may solve for $\hat{\Psi}_U$ in the optimization problem given by

$$\hat{\Psi}_U = \arg \min_{\Psi_U} \operatorname{tr}\left((\Phi_U - \Psi_U)^H (\Phi_U - \Psi_U) \right) \tag{4.6.5}$$

$$\text{s.t.} \quad \Psi_U^H \Psi_U = I_N.$$

Re-writing the objective function in terms of basis signals yields

$$\sum_{k=1}^{N} \langle \phi_{U,k} - \psi_{U,k}, \phi_{U,k} - \psi_{U,k} \rangle = N + \sum_{k=1}^{N} \pi_k^2 - 2 \sum_{k=1}^{N} \langle \phi_{U,k}, \psi_{U,k} \rangle \tag{4.6.6}$$

Next, from the Cauchy-Schwartz inequality, the relation

$$\langle \phi_{U,k}, \psi_{U,k} \rangle \leq |\langle \phi_{U,k}, \psi_{U,k} \rangle| \leq \langle \psi_{U,k}, \psi_{U,k} \rangle^{1/2} \langle \phi_{U,k}, \phi_{U,k} \rangle^{1/2} = \pi_k. \qquad (4.6.7)$$

is maximized when the inequalities are made equality. This occurs when $\hat{\psi}_{U,k} = w_k$, where w_k is the k^{th} column of W. This implies that $\hat{\Psi}_U = W$. Therefore the solution to the optimization problem is given by

$$\bar{\Psi} = WV^T.$$

An example is considered next which demonstrates the use of inner product shaping to produce an approximate transient spectrum. Consider the signal $x_1[n] = 1(0.35)^n$ and the exponential basis with decay rates $\{\sigma_k\}_{1:17} = \{0.1, 0.15, 0.2, \cdots, 0.9\}$. The resulting approximate transient spectrum is depicted in Figure 4.6.1. The figure exhibits a peak corresponding to the correct decay rate, however the amplitude is incorrect and there is significant spectral leakage throughout the spectrum. In fact, it has been observed for general transient signals $x_d[n]$ with $d > 1$ the peaks generally correspond to incorrect decay rates due large amounts of spectral leakage.
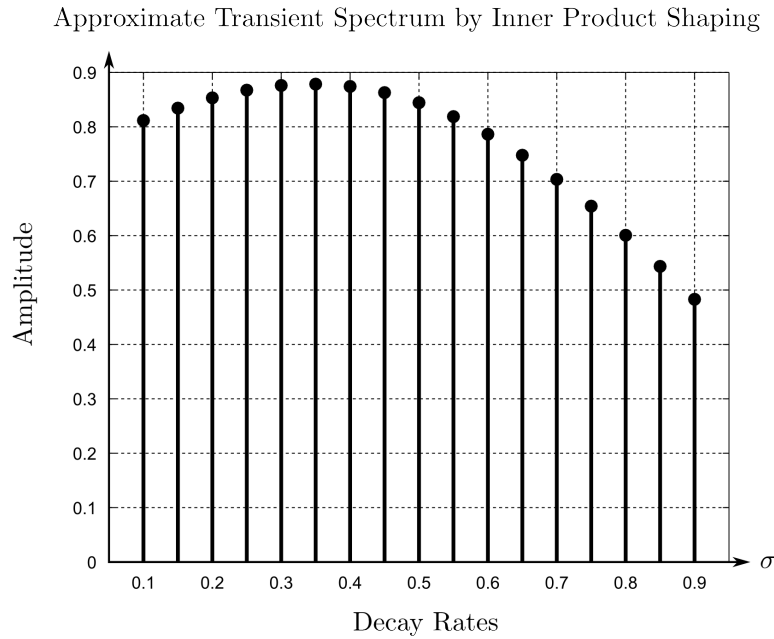
Figure 4.6.1: Spectrum of $x_1[n]$ using inner product shaping on the exponential basis

Another well known algorithm for generating an orthogonal basis given a non-orthogonal basis is the Gram-Schmidt procedure. The inner product shaping algorithm has a distinct advantage over the Gram-Schmidt process of generating an orthogonal basis given a non-orthogonal basis in that it is not order dependent. The Gram-Schmidt algorithms initial iteration uses the first original basis signal as it's first orthogonal basis signal and subsequent signals reflect only the space not spanned by the first signal. The inner product shaping algorithm does not have an ordering of original basis signals, meaning that there is no optimal ordering to be solved for, as in the case of Gram-Schmidt.

4.7 Polynomial Based Algorithms

This section considers the generation of an approximate transient spectrum for continuous-time signals. Neither an exponential basis nor a dual exponential basis is used directly in this section. The general outline of this approach is to select a known orthogonal polynomial structure, orthogonal over the domain $[0, 1]$, and perform a substitution of variables such that the resulting functions are polynomials of real, decaying exponentials. The domain is then mapped such that these exponential polynomials are orthogonal over $[0, \infty)$. This technique is known in the literature as the Orthogonal Exponential Transform. [17]

This type of algorithm is now presented by an example using Jacobi polynomials. Jacobi polynomials are used to define a specific structure of functions, which are both orthogonal and polynomials of decaying exponentials, to play the role of an approximate dual signal in order to compute transient spectral coefficients. Towards this end, consider the n^{th}-order Jacobi polynomial, defined by

$$J_n\left(a, c | x\right) = \frac{x^{1-c}\left(1 - x\right)^{c-a}\Gamma\left(c\right)}{\Gamma(c + n)} \cdot \frac{d^n}{dx^n}\left(x^{c+n-1}\left(1 - x\right)^{a+n-c}\right) \tag{4.7.1}$$

where $\Gamma\left(\cdot\right)$ is the Gamma function and a and c are design parameters. Jacobi polynomials also satisfy the recurrent relationships below. [16]

$$\frac{d}{dx}J_n(a, c | x) = -\frac{n\left(n + a\right)}{c}J_{n-1}\left(a + 2, c + 1 | x\right) \tag{4.7.2}$$

$$xJ_n\left(a, c | x\right) = \frac{c - 1}{2n + a}\left[J_n\left(a - 1, c - 1 | x\right) - J_{n+1}\left(a - 1, c - 1 | x\right)\right]. \tag{4.7.3}$$

The orthogonality of two Jacobi polynomials, for $x \in [0, 1]$, requires a weighting function of the form

$$w_J\left(x\right) = x^{c-1}\left(1 - x\right)^{a-c}. \tag{4.7.4}$$

The Jacobi orthogonality equation is given by

$$\int_0^1 x^{c-1}\left(1 - x\right)^{a-c}J_m\left(a, c | x\right)J_n\left(a, c | x\right)dx = \frac{n!\left[\Gamma\left(c\right)\right]^2\Gamma\left(n + a - c + 1\right)}{\left(a + 2n\right)\Gamma\left(a + n\right)\Gamma\left(c + n\right)} \cdot \delta_{mn}, \tag{4.7.5}$$

and is only defined for $a > 0$ and $a + 1 > c$. Next, we simplify the expression by selecting $a = c$ in order to vanish the term $(1 - x)$ in the weighting function. This simplifies the expression to

$$\int_0^1 x^{a-1}J_m\left(a, a | x\right)J_n\left(a, a | x\right)dx = \frac{n!\left[\Gamma\left(a\right)\right]^2\Gamma(n + 1)}{\left(a + 2n\right)\left[\Gamma\left(a + n\right)\right]^2} \cdot \delta_{mn}. \tag{4.7.6}$$

Performing a change of variables where $x \to e^{-t}$ requires a change in the domain; when t ranges from 0 to ∞, e^{-t} goes from 1 to 0, the same domain as x. Therefore, substituting $x = e^{-t}$ and thus

$dx = -e^{-t}dt$ into Eq. (4.7.6) results in

$$\int_0^\infty e^{-at} J_m\left(a, a|e^{-t}\right) J_n\left(a, a|e^{-t}\right) dt = \frac{n! \left[\Gamma\left(a\right)\right]^2 \Gamma(n+1)}{(a+2n)\left[\Gamma\left(a+n\right)\right]^2} \cdot \delta_{mn}. \tag{4.7.7}$$

We now consider selecting a value for a. Note that for any order n, $J_n\left(a, a|0\right) = 1$. Therefore it is desirable to have the Jacobi polynomials multiplied by the expression e^{-t}, this way the resulting functions approach 0 as t goes to infinity. One simple choice results from selecting $a = 2$. Continuing with this choice the expression simplifies to

$$\int_0^\infty \left(e^{-t} J_m\left(2, 2|e^{-t}\right)\right) \left(e^{-t} J_n\left(2, 2|e^{-t}\right)\right) dt = \frac{1}{2(1+n)^3} \cdot \delta_{mn}. \tag{4.7.8}$$

Using this simplified equation, the structure of functions which take the role of the dual signals in creating an approximate transient spectrum based on Jacobi polynomials is defined. Define the order $(k+1)$ function to have the structure

$$\acute{\psi}_{k+1}\left(t\right) = (-1)^k \sqrt{2\left(1+k\right)^3} e^{-t} J_k\left(2, 2|e^{-t}\right) \tag{4.7.9}$$

for $k \geq 0$. Substituting these functions into the continuous inner product, it is straightforward to verify that

$$\langle \acute{\psi}_k, \acute{\psi}_l \rangle = \int_0^\infty \acute{\psi}_k\left(t\right) \acute{\psi}_l^*\left(t\right) dt = \delta_{kl}. \tag{4.7.10}$$

Therefore a continuous-time transient signal $x_d\left(t\right)$, as in Eq. (2.1.1), may be represented by

$$x_d\left(t\right) = \sum_{k=1}^\infty \beta_k \acute{\psi}_k\left(t\right) \tag{4.7.11}$$

where we compute β_k as the continuous inner product of $x_d\left(t\right)$ with the approximate dual function $\acute{\psi}_k\left(t\right)$ as

$$\beta_k = \int_0^\infty x_d\left(t\right) \acute{\psi}_k\left(t\right) dt. \tag{4.7.12}$$

[2] showed that $x_d\left(t\right)$ may be represented in the following manner.

$$\begin{aligned} x_d\left(t\right) &= \sum_{k=1}^\infty \sum_{l=1}^k \beta_k c_{kl} e^{-lt} \\ &= \sum_{l=1}^\infty \alpha_l e^{-lt}. \end{aligned}$$

Therefore the approximate transient spectrum using Jacobi polynomials is defined as the α_l values as a function of l. In practice, an infinite number of coefficients cannot be computed, so a finite approximation is used instead.

It has been observed that the spectrum produced using the approximate dual signals formed

by Jacobi polynomials generally does not exhibit meaningful spectral results. When the transient spectrum is generated for a linear combination of exponentials with integer decay rates the transient spectrum is often well behaved. On the contrary, when the decay rates are not integer, the resulting spectrum may have wildly varying behavior. The following example demonstrates these phenomena. Consider the transient signal $x_1(t) = e^{-\lambda t}$ for $\lambda = 2$ and $\lambda = 2.1$. The resulting transient spectrum for $\lambda = 2$ is shown in Figure 4.7.1(a) while the spectrum for $\lambda = 2.1$ is shown in Figure 4.7.1(b). In the second case, it is clear that the transient spectrum produced does not exhibit a broadened peak between $\lambda = 2$ and $\lambda = 3$ and provides no indication of the true decay rate of $x_1(t)$.

(a) Spectrum of $x_1(t) = e^{-2t}u(t)$ using the OET with Jacobi Polynomials

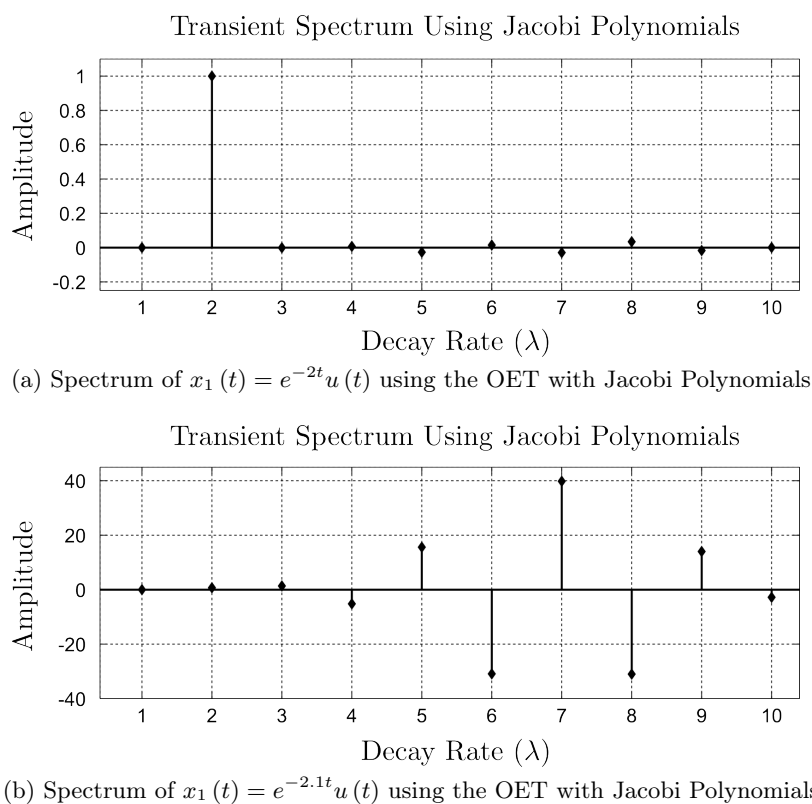(b) Spectrum of $x_1(t) = e^{-2.1t}u(t)$ using the OET with Jacobi Polynomials

Figure 4.7.1: Spectra produced using the OET using Jacobi polynomials to form a basis

Chapter 5

Overdetermined Recovery

The successful recovery of the decay rates present in a transient signal given only samples of the lowpass filtered transient signal was previously defined in Chapter 2 under Definitions 2 and 4. These definitions were motivated by noting that the process of sampling a causal, continuous-time transient signal, i.e., the signal in Eq. (2.1.1), to produce the transient time-series in Eq. (2.1.2) unavoidably included aliasing effects due to the infinite bandwidth of the continuous-time signal. Furthermore, most physical sampling systems have a lowpass characteristic and often include an anti-aliasing LPF as the first stage of the sampling process. Many signal processing systems also include a sample rate conversion system in which a signal is lowpass filtered as a part of this conversion process, e.g., oversampled noise shaping. This chapter develops a parameter recovery framework for transient signals inspired by these types of scenarios. The general approach is to solve for a larger number of decay rates than the true order of the transient signal in such a way that from this larger set of parameters the correct parameters are identifiable. This procedure is defined as overdetermined recovery.

In order to use the overdetermined recovery framework, the order of a transient model must be known *a priori*. The parameter recovery algorithms presented in Appendix B each require the model order to be established before determining values for the decay rates. Each recovery algorithm assumes a fixed model order of d, meaning that each method produces d values for the decay rates as a function of the available data. However, it is often the case that the order of a transient signal acquired is not known *a priori*, resulting in a need to estimate it. To do this, a typical approach begins by postulating several potential model orders. Based on these educated guesses, one computes some error criterion that indicates which model order to select. The process of determining a model order encompasses the use of intuition and insight into the structure of the problem—a prime example of the art in engineering. In many cases there is no obvious solution, and thus an engineering judgement must be made. Selecting too low a model order produces a smooth estimate of the Fourier spectrum of the data while selecting too high a model order increases the resolution and introduces spurious detail into the frequency representation. The issues associated with selecting an appropriate model order demonstrate the classic tradeoff scenario between increased resolution and decreased variance.

The algorithms for model order selection presented in this chapter may be used as a guideline for initial order selection. The techniques presented are known to work well for computer generated, synthetic transient signals, but may or may not work well with actual data depending on how well such data is modeled by a transient structured signal. Section 5.1 begins by providing a qualitative insight into why the structure of the SVD is useful for de-noising a set of transient data. Taking advantage of this representation, Section 5.2 proposes an algorithm which uses the SVD in conjunction with the Eckart-Young theorem in order to use the numerical rank of a data matrix to determine the transient model order as a function of the available data.

In Section 5.3, several typical examples of the overdetermined recovery framework are highlighted using different lowpass filters, recovery algorithms, data lengths, and overdetermined model orders. These examples provide insight into the recovery process, where the correct decay rates are visually identified from the larger set of parameters, as described in Section 5.4. Finally, an alternating projection algorithm is proposed in Section 5.5 in order to solve the transient recovery problem for the amplitude coefficients utilizing both the DFT and the DTT.

5.1 SVD-like Transient Representation

The singular value decomposition, which was defined in Section 3.4.3, is an extremely powerful tool for determining the order of a transient signal given data. The exposition in this section leads directly to an algorithm for determining the order of a transient model as a function of available data using an SVD-like representation. The following discussion specifically pertains to the transient structure of signals. However, a straightforward generalization may be used to extend this insight to other structures, e.g., sinusoidal or damped sinusoidal signal models. Note that the model order determination algorithms in this chapter rely on access to the transient signal prior to the lowpass filter. If this data is unavailable, Section 5.5 presents an algorithm that attempts to recover a transient signal prior to the lowpass filter, for which the following model order determination may be used.

Suppose, for example, that a transient signal is acquired over the interval $0 \leq n \leq N-1$, where the true order d is unknown. In order to use a parameter recovery algorithm, an estimated model order $\hat{d}$ is chosen such that the data, either $\{x_d\,[n]\}_{0:N-1}$ or $\{\bar{x}_d\,[n]\}_{0:N-1}$, can be modeled as a sum of $\hat{d}$ decaying exponentials with distinct decay rates and real amplitudes. We first present a structural understanding of why the SVD is useful for identifying a model order, $\hat{d}$, in the noiseless case. An understanding of the noiseless case can then easily be extended to include when the available data contains additive noise, which is discussed in the next section.

We begin the process of determining the model order by defining an order $\hat{d}$ Toeplitz data

matrix, which will provide much insight into this problem, by

$$
X^{\left(\hat{d}\right)} = \begin{bmatrix} x_d\left[\hat{d}-1\right] & x_d\left[\hat{d}-2\right] & \cdots & x_d\left[0\right] \\ x_d\left[\hat{d}\right] & x_d\left[\hat{d}-1\right] & \cdots & x_d\left[1\right] \\ \vdots & \vdots & \ddots & \vdots \\ x_d\left[N-1\right] & x_d\left[N-2\right] & \cdots & x_d\left[N-\hat{d}\right] \end{bmatrix}. \tag{5.1.1}
$$

The range of values taken by the integer $\hat{d}$ will momentarily be restricted. First, consider the case of a simple transient signal with a single decay rate, i.e., $x_1\left[n\right]$. In this case, the structure of the transient signal is known to be $x_1\left[n\right] = \alpha_1\left(\sigma_1\right)^n$, for $0 \le n \le N-1$. Plugging this model for $x_1\left[n\right]$ into the definition of $X^{\left(\hat{d}\right)}$ yields

$$
X^{\left(\hat{d}\right)} = \left(\alpha_1\right) \begin{bmatrix} \sigma_1^{\hat{d}-1} & \cdots & \sigma_1 & 1 \\ \sigma_1^{\hat{d}} & \cdots & \sigma_1^2 & \sigma_1 \\ \vdots & \ddots & \vdots & \vdots \\ \sigma_1^{N-1} & \cdots & \sigma_1^{N-\hat{d}+1} & \sigma_1^{N-\hat{d}} \end{bmatrix}, \tag{5.1.2}
$$

which is valid for $\hat{d} > 1 = d$. It is straightforward to verify that the matrix in Eq. (5.1.2) factors into an outer product in the following manner:

$$
X^{\left(\hat{d}\right)} = \alpha_1 f\left(\sigma_1\right) h\left(\sigma_1\right)^T \tag{5.1.3}
$$

where the dimension $N - \hat{d} + 1$ vector $f\left(\sigma_1\right)$ is defined by

$$
f\left(\sigma_1\right) = \left[1, \sigma_1, \sigma_1^2, \cdots, \sigma_1^{N-\hat{d}}\right]^T \tag{5.1.4}
$$

and the dimension $\hat{d}$ vector $h\left(\sigma_1\right)$ is defined by

$$
h\left(\sigma_1\right) = \left[\sigma_1^{\hat{d}-1}, \sigma_1^{\hat{d}-2}, \cdots, \sigma_1, 1\right]^T. \tag{5.1.5}
$$

The decomposition of the matrix $X^{\left(\hat{d}\right)}$ into this form makes clear that the matrix has unit rank. By exploiting linearity, when $x_d\left[n\right] = \sum_{k=1}^{d} \alpha_k\left(\sigma_k\right)^n$, for $0 \le n \le N-1$, the expression for $X^{\left(\hat{d}\right)}$ is expanded to be

$$
X^{\left(\hat{d}\right)} = \sum_{k=1}^{d} \alpha_k f\left(\sigma_k\right) h\left(\sigma_k\right)^T, \tag{5.1.6}
$$

or equivalently as

$$X^{(\hat{d})} = F(\sigma)\,\mathcal{A}(\alpha)\,H(\sigma)^T \tag{5.1.7}$$

$$= [f(\sigma_1), f(\sigma_2), \cdots, f(\sigma_d)] \begin{bmatrix} \alpha_1 & 0 & \cdots & 0 \\ 0 & \alpha_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha_d \end{bmatrix} \begin{bmatrix} h(\sigma_1)^T \\ h(\sigma_2)^T \\ \vdots \\ h(\sigma_d)^T \end{bmatrix}. \tag{5.1.8}$$

From this factorization, specifically the structures of $F(\sigma)$, $H(\sigma)$, and $\mathcal{A}(\alpha)$, it is straightforward to verify that $\mathrm{rank}\left(X^{(\hat{d})}\right) = d$. Consequently, when initially selecting values for $\hat{d}$, any prior knowledge on the transient model order should be used to select a value $\hat{d}$ which is believed to satisfy $d \leq \hat{d} \leq N - d$. Generally, because the value of d is unknown, several values of $\hat{d}$ are often attempted in hopes that one may fall into this range.

The matrix decomposition in Eq. (5.1.7) is similarly structured to the SVD of the matrix $X^{(\hat{d})}$. It is important to understand that these two decompositions are not equivalent. This is easily understood by considering either the matrix $F(\sigma)$ or $H(\sigma)$. For example, the matrix $F(\sigma)$ is guaranteed not to be unitary. In fact, when the matrix $X^{(\hat{d})}$ is populated with a transient structured signal, $F(\sigma)$ has columns corresponding to transient structured signals, which we have shown to be non-orthogonal. A decomposition of this structure would be extremely useful for transient model order determination, but unfortunately no known method decomposes a general Toeplitz matrix into this form.

5.2   Model Order Determination Algorithms

The factorization of the Toeplitz data matrix in Eq. (5.1.7) is structured similarly to the SVD factorization in Eq. (3.4.17). This section combines this similarity with the fact that the procedure to compute an SVD is well-known in order to discuss a commonly used algorithm for determining the model order of available transient data and to propose a new method. [23]

In Eq. (5.1.2) the decomposition of the Toeplitz data matrix $X^{(\hat{d})}$, populated using the signal $x_1[n]$, was not necessary for determining the true model order d. In fact, only the rank of that matrix needed to be determined. Now the application of the SVD becomes clear. First, choose an initial integer value for $\hat{d}$ which is believed to satisfy $d \leq \hat{d} \leq N - d$. Next, form the matrix $X^{(\hat{d})}$ for this initial value of $\hat{d}$. The rank of $X^{(\hat{d})}$ is easily determined by computing the SVD and counting the total number of non-zero singular values. For the case of noiseless transient data, the transient model order $\hat{d} = d$ is determined once the following condition is satisfied:

$$\mathrm{rank}\left(X^{(\hat{d}+1)}\right) = \hat{d}. \tag{5.2.1}$$

In fact, due to the noiseless assumption, any value of $\check{d}$ such that $\check{d} > d$ may be used where the rank

of $X^{(d)}$ is d and therefore the model order can be determined immediately.

In practice, the Toeplitz data matrix is often populated with noisy transient data corresponding to the underlying transient signal corrupted by measurement noise, i.e., $\{\bar{x}_d[n]\}_{0:N-1}$. Using the intuition developed in the noiseless case, this procedure is now extended to the case of a noisy Toeplitz data matrix of the form

$$\bar{X}^{(\hat{d})} = \begin{bmatrix} x_d\left[\hat{d}-1\right] & \cdots & x_d[0] \\ x_d\left[\hat{d}\right] & \cdots & x_d[1] \\ \vdots & \ddots & \vdots \\ x_d[N-1] & \cdots & x_d\left[N-\hat{d}\right] \end{bmatrix} + \begin{bmatrix} \eta\left[\hat{d}-1\right] & \cdots & \eta[0] \\ \eta\left[\hat{d}\right] & \cdots & \eta[1] \\ \vdots & \ddots & \vdots \\ \eta[N-1] & \cdots & \eta\left[N-\hat{d}\right] \end{bmatrix}. \tag{5.2.2}$$

In general, $\bar{X}^{(\hat{d})}$ will always have full rank for any valid choice of $\hat{d}$. Therefore, the stopping criterion in the previously described algorithm must be altered in order to overcome this obstacle.

The most straightforward extension of the previous stopping criterion is to determine the numerical rank in place of the rank of the data matrix. We previously defined the rank of a matrix by the total number of non-zero singular values in the SVD representation of that matrix. However, taking into consideration numerical representation used in computation and the effects of additive noise, the numerical rank of a matrix is an important quantity in practice. The numerical rank of a general $M \times N$ matrix A, denoted $\text{rank}_n(A)$, is defined by

$$\text{rank}_n(A) = \left| \{\pi_k > \max(M,N) * \epsilon \|A\|_2\}_{1:\min(M,N)} \right| \tag{5.2.3}$$

where $|\cdot|$ represents the cardinality of a set of singular values and ϵ represents the relative machine precision for the hardware the algorithm is computed on, e.g., $\epsilon = 2.22 \times 10^{-16}$. Therefore, the modified stopping criterion is to find the value of $\hat{d}$ for which

$$\text{rank}_n\left(\bar{X}^{(\hat{d}+1)}\right) = \hat{d}. \tag{5.2.4}$$

Note that alternative thresholds to $\max(M,N) * \epsilon \|A\|_2$ may be used, but the observation that the smallest singular values of noiseless data matrices are quite small should be taken into account when designing such a threshold.

Next, an alternative model order estimation algorithm is proposed by considering the application of the Eckart-Young theorem, as discussed in Appendix B.2, to the matrix $\bar{X}^{(\hat{d})}$ in Eq. (5.2.2). An inherent assumption in this procedure is that the smallest non-zero singular value in the SVD representation of $\bar{X}^{(\hat{d})}$ corresponds to the noise in the samples and not one of the exponential components. If we denote the result of this de-noising process by $G^{(\hat{d})}$, then the matrix $E^{(\hat{d})} = \bar{X}^{(\hat{d})} - G^{(\hat{d})}$ corresponds to the smallest unit rank reducing perturbation of the noisy data matrix in a Frobenius norm sense. This estimate of the noise matrix is generally not Toeplitz structured, although we know from Eq. (5.2.2) that the true noise matrix is. Therefore, we produce a new

Toeplitz noise estimate by

$$\hat{E}^{(\hat{d})} = h\left(E^{(\hat{d})}\right) \tag{5.2.5}$$

where h is the diagonal-averaging operator in Eq. (B.3.4). Using this estimate for the noise, we create an estimate of the underlying noiseless transient data as

$$\hat{X}^{(\hat{d})} = \bar{X}^{(\hat{d})} - \hat{E}^{(\hat{d})}.$$

This estimate of the de-noised data is guaranteed to remain Toeplitz. $\hat{X}^{(\hat{d})}$ can then be used in the same way $\bar{X}^{(\hat{d})}$ was previously for model order determination, i.e., a comparison of the number of singular values larger than a given threshold. This proposed method is summarized by determining the model order of the data given by

$$\bar{X}^{(\hat{d})} - h\left(\bar{X}^{(\hat{d})} - f\left(\bar{X}^{(\hat{d})}\right)\right) \tag{5.2.6}$$

where f is the unit rank reduction operator from Eq. (B.3.5).

Consider the signal $\bar{x}_4[n]$ over the interval $0 \leq n \leq 28$, using decay rates $\{\sigma_k\}_{1:4} = \{0.3, 0.5, 0.7, 0.9\}$ and amplitude coefficients $\{\alpha_k\}_{1:4} = \{3, -5, 7, 2\}$. Figure 5.2.1 depicts the singular values of the matrix $X^{(8)}$ from Eq. (5.1.1), the average singular values of the matrix $\bar{X}^{(8)}$ from Eq. (5.2.2), and the average singular values of the proposed algorithm in Eq. (5.2.6) with $\hat{d} = 8$, where the averages are over 100,000 trials for various SNRs. Note that in all 3 cases, the singular values drop off sharper for the proposed algorithm, making the underlying model order $d = 4$ easier to identify.

(a) Average singular values for three transient model order estimation
matrices with an SNR of 10 dB

(b) Average singular values for three transient model order estimation
matrices with an SNR of 0 dB

(c) Average singular values for three transient model order estimation
matrices with an SNR of −10 dB

Figure 5.2.1: Singular values of a noisy transient signal using the proposed model order estimation
algorithm for different SNR values

5.3   Illustrating Overdetermined Recovery

Consider the following example in which the overdetermined recovery framework is applied
using the Tufts-Kumaresan parameter recovery algorithm on the data $\{g_4\,[n]\}_{0:29}$ for several values
of $\hat{d}$. The signal $g_4\,[n]$ is produced by lowpass filtering the transient signal $x_4\,[n]$ with decay rates
$\{\sigma_k\} = \{0.3,\, 0.5,\, 0.7,\, 0.9\}$ and amplitude coefficients $\{\alpha_k\}_{1:4} = \{-1,\, 2,\, 7,\, 3\}$ by a truncated sinc
filter of length 65536 and cutoff frequency $\omega_c = 0.8\pi$. Figure 5.3.1 depicts the resulting decay rates
where in parts (a), (b), (c), and (d) the values 4, 8, 12, and 16 are used for $\hat{d}$, respectively. This
example shows that as the value of $\hat{d}$ increases from the model order d, the desired decay rates are

89

easy to visually identify in the z−plane. However, it becomes increasingly difficult to identify the decay rates when the value of $\hat{d}$ becomes too high.
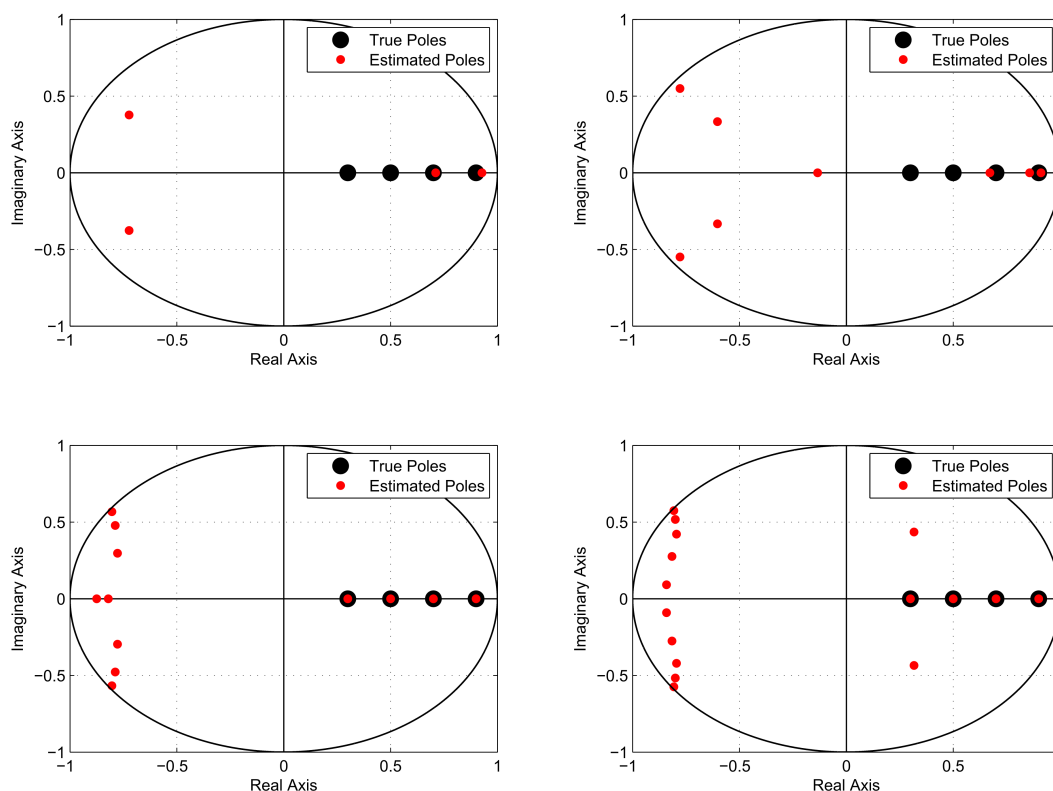
Figure 5.3.1: Pole-zero representation of the estimated decay rates resulting from the Tufts-Kumaresan method for various values of $\hat{d}$

The decay rates resulting from the overdetermined recovery framework are presented in the following three examples for the same signal $x_4[n]$ as above, but for different LPFs, data records, parameter recovery algorithms, and overdetermined model orders.

Example. Figure 5.3.2 depicts the decay rates resulting from the overdetermined recovery framework where the signal $x_4[n]$ is lowpass filtered using a length 512 Parks-McClellan LPF with cutoff frequency $\omega_c = \frac{3\pi}{4}$, the overdetermined model order $\hat{d} = 13$ is used, the data record length is $N = 34$, and the covariance method of linear prediction is used for recovering the decay rates.

Figure 5.3.2: Overdetermined recovery results after lowpass filtering $x_4[n]$ by a Parks-McClellan LPF

Example. Figure 5.3.3 depicts the decay rates resulting from the overdetermined recovery framework where the signal $x_4[n]$ is lowpass filtered using an 18^{th}-order Butterworth LPF with cutoff frequency $\omega_c = \frac{\pi}{2}$, the overdetermined model order $\hat{d} = 20$ is used, the data record length is $N = 44$, and the extended method of Prony is used for recovering the decay rates.

Figure 5.3.3: Overdetermined recovery results after lowpass filtering $x_4[n]$ by a Butterworth LPF

Example. Figure 5.3.4 depicts the decay rates resulting from the overdetermined recovery framework where the signal $x_4[n]$ is lowpass filtered using a 10^{th}-order Elliptic LPF with cutoff frequency

$\omega_c = \frac{3\pi}{10}$, the overdetermined model order $\hat{d} = 23$ is used, the data record length is $N = 50$, and Cadzows method for recovering the decay rates.
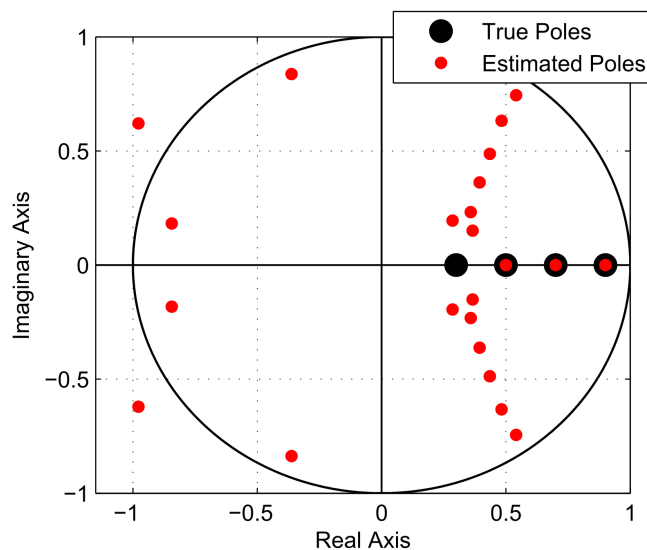

Figure 5.3.4: Overdetermined recovery results after lowpass filtering $x_4[n]$ by an Elliptic LPF

A simple qualitative explanation of the overdetermined recovery procedure is now proposed with regard to the spectral matching property of the Yule-Walker method of all-pole modeling. Consider breaking the frequency spectrum of the signal $g_d[n]$ into the passband and stopband regions of the lowpass filter. The frequency spectrum over the passband can be completely characterized using d poles placed along the real line between 0 and 1 in the z-plane. However, if we only use d poles to model the entire frequency spectrum, then the poles also attempt to model the stopband attenuation, and thus will not model the passband exactly. By allowing $\hat{d} - d$ additional poles to model the stopband, the original d poles may completely characterize the passband. At a high level, this explanation hints at the fact that the estimated poles should be easily separated into two groups: the decay rates and the extraneous solutions. A systematic way to differentiate between the decay rates and the extraneous solutions is considered in Section 5.4. The preceding examples demonstrated typical results of the overdetermined recovery framework where the two sets of poles were easy to visually separate in the z-plane: the decay rates were found on the real axis between 0 and 1, while the extraneous solutions were found in regions corresponding to the lowpass filters high frequency attenuation. Note that when a lowpass filter has a narrow passband, the decay rates corresponding to quickly decaying exponentials are not recoverable.

5.4 Overdetermined Recovery Framework

The key criterion in selecting an order for the overdetermined recovery framework, defined as d^*, is that the underlying model order has already been established and we need to select d^* to be meaningfully larger than the true order d, with respect to the size of the LPF's passband. The general procedure of overdetermined recovery simply uses the parameter recovery algorithms in Appendix B with the overdetermined order d^* where an additional step is needed to sort through the resulting parameter estimates in order to identify the correct decay rates, $\{\sigma_k\}_{1:d}$. This additional step is discussed in the remainder of this section.

Note that an overdetermined parameter recovery algorithm produces a set of d^* decay rate estimates for which none of the estimates are guaranteed to be real or positive. Generally, when the wrong number of recovered parameters are on the positive real axis in the z-plane, one of two problems most likely occurred: (1) the passband of the LPF may be too narrow to recover a quickly decaying component and/or (2) the value of d^* may have been chosen either too high or too low for the given LPFs passband. A useful heuristic is that the larger the stopband of the LPF, the higher the overdetermined order d^* should be. Often several choices for d^* must be attempted in a trial-and-error type procedure. Section 5.4.1 describes process of separating the decay rates from the extraneous solutions where the locations of the extraneous solutions are unconstrained and the value of d^* has been chosen appropriately for the passband of the LPF. Subspace based parameter recovery algorithms, e.g. the Tufts-Kumaresan method or Cadzows method, are capable of producing $d^* - d$ sets of d^* parameter estimates. In the noiseless case, each estimate has the same d decay rates, but the extraneous solutions are often different. Section 5.4.2 identifies one set of parameter estimates for which the extraneous solutions are guaranteed to be within the unit circle.

5.4.1 Classifying Roots

Recall the Toeplitz data matrix from Eq. (5.1.1) with the structure

$$X^{(d^*+1)} = \begin{bmatrix} x_d\,[d^*] & x_d\,[d^*-1] & \cdots & x_d\,[0] \\ x_d\,[d^*+1] & x_d\,[d^*] & \cdots & x_d\,[1] \\ \vdots & \vdots & \ddots & \vdots \\ x_d\,[N-1] & x_d\,[N-2] & \cdots & x_d\,[N-d^*-1] \end{bmatrix}. \tag{5.4.1}$$

The homogeneous solution to the linear system of equations given by

$$X^{(d^*+1)}\theta = 0, \tag{5.4.2}$$

where $\theta = [1, \theta_1, \cdots, \theta_{d^*}]^T$, lies in the nullspace of the matrix $X^{(d^*+1)}$, and consequently in the nullspace of the matrix $\left[X^{(d^*+1)}\right]^T X^{(d^*+1)}$ too. If the matrix $X^{(d^*+1)}$ contains noise, i.e., $\bar{X}^{(d^*+1)}$, then, with high probability, the data matrix has an empty nullspace. In this case, the right singular vectors corresponding to the $d^* - d$ smallest singular values are defined to be the approximate nullspace of $\bar{X}^{(d^*+1)}$. In either of these cases, the $d^* - d$ solutions are easily obtainable via the SVD.

Next, define the monic, order d^*, characteristic polynomial using the coefficients from any solution θ by

$$\Theta(z) = 1 + \sum_{k=1}^{d^*} \theta_k z^{-k}. \tag{5.4.3}$$

Section 5.3 illustrated that for a broad range of cutoff frequencies, d of the zeros of $\Theta(z)$ corresponded to the desired decay rates $\{\sigma_k\}_{1:d}$. However, the characteristic polynomial also has $d^* - d$ additional zeros, which we defined to be the extraneous solutions. Therefore, the order d^* characteristic polynomial can be factored such that

$$\Theta(z) = \Theta_1(z)\,\Theta_2(z) \tag{5.4.4}$$

where the two sub-characteristic polynomials are defined by

$$\Theta_1(z) = \sum_{k=0}^{d} p_k z^{-k} = \prod_{k=1}^{d} \left(1 - \sigma_k z^{-1}\right), \tag{5.4.5}$$

where we select $p_0 = 1$, and

$$\Theta_2(z) = \sum_{k=0}^{d^*-d} q_k z^{-k} = \prod_{k-1}^{d^*-d} \left(1 - \vartheta_k z^{-1}\right), \tag{5.4.6}$$

where we select $q_0 = 1$. The zeros of $\Theta_2(z)$ are the extraneous solutions.

Consider the case where d^* is selected such that $d^* - d > 1$. As a consequence, there are at least two vectors in the nullspace of $X^{(d^*+1)}$, or at least two vectors that span the approximate nullspace of $\bar{X}^{(d^*+1)}$, depending on whether the data contains noise or not. We proceed by assuming noiseless data, although the following procedure is similar in both cases where approximations are made accordingly.

We would like to separate the zeros of $\Theta(z)$ into $\{\sigma_k\}_{1:d}$ and $\{\vartheta_k\}_{1:d^*-d}$, or equivalently to factor $\Theta(z)$ into $\Theta_1(z)$ and $\Theta_2(z)$, using any solution θ. One procedure for doing so is to systematically discard the extraneous solutions for each θ. The decay rates appear in each solution while the extraneous solutions are typically different, therefore, a systematic procedure for identifying the decay rates is to identify the zeros that are real, in the range $(0, 1)$, and common to each of the characteristic polynomials formed by homogeneous solutions to Eq. (5.4.2).

5.4.2   Special Case of the Extraneous Solutions

When multiple homogeneous solutions θ exist for the system of equations in Eq. (5.4.2), one of these solutions has an identifiable property for which the locations of the extraneous solutions are guaranteed to be within the unit circle. [11] This subsection uses the correlation method of linear prediction in order to identify which solution guarantees this property. Specifically, we show that if

the solution θ with coefficients that minimize the sum of squares, i.e.,

$$\|\theta\|_2^2 = 1 + \sum_{k=1}^{d^*} |\theta_k|^2 , \tag{5.4.7}$$

is used to form the polynomial $\Theta(z)$, then the zeros of $\Theta_2(z)$ are found to lie strictly inside the unit circle. Note that each vector must be normalized such that $\theta_0 = 1$ for the comparison.

We previously factored the polynomial $\Theta(z)$ into the product of the two polynomials $\Theta_1(z)$ and $\Theta_2(z)$, therefore the coefficients in the vector θ are the result of a linear convolution between the coefficients of the two sub-polynomials. Using the notation in Eqs. (5.4.5) and (5.4.6), the coefficients of θ are given by

$$\theta_n = \sum_k p_k q_{n-k}. \tag{5.4.8}$$

We consider the the coefficients p_k to be fixed, as they correspond to the desired decay rates. By interpreting this convolution as linear prediction, where the input sequence is the signal with co-efficients from $\Theta_1(z)$, we select the coefficients q_n in order to minimize the sum of squares of the linear prediction error, θ_n. Thus the solution, i.e., coefficients for $\Theta_2(z)$, is the solution to the Yule-Walker equations in Eq. (B.4.3), where the autocorrelation sequence is structured as the biased autocorrelation estimator in Eq. (B.4.4). A well-known result states that the solution to this set of equations always results in a set of roots that have magnitude less than 1. Therefore, by selecting the solution θ with minimum ℓ_2-norm, the extraneous solutions are guaranteed to be found inside the unit circle. In Figure 5.3.1(a)-(d) the solution θ was chosen with minimum ℓ_2-norm, resulting in the parameter estimates being contained in the unit circle.

5.5   Iterative Transient Spectral Projection

This section proposes an alternating projection algorithm for the lowpass filtered transient pa-rameter recovery problem in order to recover the signal $x_d[n]$ from samples of the signal $g_d[n]$. This algorithm exploits two inherent properties of the underlying transient signal $x_d[n]$ in order to iteratively project a signal between two spaces in anticipation of converging to the underlying signal, for which a parameter recovery algorithm will identify the desired parameters. Each of these spaces satisfies one of the known properties of $x_d[n]$. To begin, define the space of all d^{th}-order transient structured signals as

$$\mathcal{V}_d = \{x_d[n] \ : \ x_d[n] \text{ of the form of Eq. (2.1.2)}\}$$

and the space $\mathcal{W}_{g_d}$ to be the space of all signals with the same Fourier transform as the signal $g_d[n]$ over the passband regime of the lowpass filter. Therefore, this algorithm can be summarized by

$$\text{find } x_d^*[n] \tag{5.5.1}$$
$$\text{s.t. } x_d^*[n] \in \mathcal{V}_d \cap \mathcal{W}_{g_d}.$$

Appendix A establishes the uniqueness of the solution to this formulation. To solve this formulation, the parametric estimator of a d^{th}-order transient signal, defined in Eq. (2.5.1), is iteratively updated as its parameter estimates are updated. Therefore, the t^{th} iteration of the parametric estimate of the underlying transient structured signal is given by

$$\hat{x}_d^{(t)}[n] = \sum_{k=1}^{d} \alpha_k^{(t)} \left(\sigma_k^{(t)}\right)^n, \tag{5.5.2}$$

where the parameters $\left\{\alpha_k^{(t)}\right\}_{1:d}$ and $\left\{\sigma_k^{(t)}\right\}_{1:d}$ are the t^{th} iterative estimates of the amplitude coefficients and decay rates, respectively. The overall methodology used here is to iteratively update the signal $g_d^{(t)}[n]$ by

$$g_d^{(t+1)}[n] = \mathcal{P}_{\mathcal{W}_{g_d}}\left\{\mathcal{P}_{\mathcal{V}_d}\left\{g_d^{(t)}[n]\right\}\right\} \tag{5.5.3}$$

until $g_d^{(t+1)}[n] = g_d^{(t)}[n]$, where $\mathcal{P}_{\mathcal{V}_d}\{\cdot\}$ and $\mathcal{P}_{\mathcal{W}_{g_d}}\{\cdot\}$ are projections into the spaces $\mathcal{V}_d$ and $\mathcal{W}_{g_d}$, respectively.

We initialize the algorithm by defining the projector $\mathcal{P}_{\mathcal{V}_d}$ to satisfy an input-output relationship given by

$$\hat{x}_d^{(0)}[n] = \mathcal{P}_{\mathcal{V}_d}\left\{g_d^{(0)}[n]\right\}. \tag{5.5.4}$$

Given the filtered transient data, denote $g_d^{(0)}[n] = g_d[n]$. The initial step of this algorithm is to first estimate the location of $d^* > d$ decay rates using the overdetermined recovery framework. Define the mapping $\mathcal{D}\{\cdot\}$ as this procedure. Then the initial step of this algorithm is then given by

$$\left\{\sigma_k^{(0)}\right\}_{1:d^*} = \mathcal{D}\left\{g_d^{(0)}[n]\right\}. \tag{5.5.5}$$

Using the real decay rates between 0 and 1 from this overdetermined recovery mapping as the decay rates of an exponential basis, the exponential basis is extended to include d^* total real decay rates. The additional decay rates may be chosen arbitrarily, though a good heuristic is to place a single decay rate relatively close to each side of the recovered decay rates. The DTT of $g_d^{(0)}[n]$, is then computed, and is denoted as

$$G_d^{(0)}[k] = \mathcal{T}\left\{g_d^{(0)}[n]\right\}. \tag{5.5.6}$$

Next, an adaptive transient filter $\bar{\mathcal{H}}[k]$, as discussed in Section 3.3, is defined by

$$\bar{\mathcal{H}}[k] = \begin{cases} 1, & \text{for the largest } d \text{ peaks of } \mathcal{A}[k] \\ 0, & \text{otherwise} \end{cases} \tag{5.5.7}$$

where $\mathcal{A}[k]$ is the DTT of the transient filters input. Therefore, the output of the projector $\mathcal{P}_{\mathcal{V}_d}$ is given by

$$\hat{x}_d^{(0)}[n] = \mathcal{T}^{-1}\left\{G_d^{(0)}[k]\,\bar{\mathcal{H}}[k]\right\}. \tag{5.5.8}$$

This parametric signal contains d real exponentials, which means that it sits in the space of all d^{th}-order transient signals, i.e., $\mathcal{V}_d$. However, the passband frequency is not guaranteed to match the known frequency content of the underlying transient signal. Therefore, in order to enforce this property, define the input-output relationship of the projector $\mathcal{P}_{\mathcal{W}_{g_d}}$ by

$$g_d^{(1)}[n] = \mathcal{P}_{\mathcal{W}_{g_d}} \left\{ \hat{x}_d^{(0)}[n] \right\}. \tag{5.5.9}$$

In this projection, the known frequency band is inserted in place of the estimated band while the out of band estimate is left unmodified. The result of this procedure is denoted $G_d^{(1)}[k]$ and is given by

$$G_d^{(1)}[k] = (1 - F[k])\,\hat{X}_d^{(0)}[k] + \mathcal{F}\left\{ g_d^{(0)}[n] \right\} \tag{5.5.10}$$

where $\mathcal{F}\{\cdot\}$ denotes the DFT, $F[k]$ represents the lowpass filter, and

$$\hat{X}_d^{(0)}[k] = \mathcal{F}\left\{ \hat{x}_d^{(0)}[n] \right\}. \tag{5.5.11}$$

Finally, the first iteration is completed by transforming the signal with correct frequency content over the passband region, but possibly no longer in $\mathcal{V}_d$, to the sample domain by

$$g_d^{(1)}[n] = \mathcal{F}^{-1}\left\{ G_d^{(1)}[k] \right\}. \tag{5.5.12}$$

The iterative update procedure to update the signal $g_d^{(t)}[n]$ to $g_d^{(t+1)}[n]$ is given in Algorithm 5.1.

Algorithm 5.1 Update procedure for the iterative transient spectral projection algorithm

1. $\hat{x}_d^{(t)}[n] = \mathcal{P}_{\mathcal{V}_d}\left\{ g_d^{(t)}[n] \right\}$:

    (a) $\left\{ \sigma_k^{(t)} \right\}_{1:d^*} = \mathcal{D}\left\{ g_d^{(t)}[n] \right\}$

    (b) $G_d^{(t)}[k] = \mathcal{T}\left\{ g_d^{(t)}[n] \right\}$

    (c) $\hat{x}_d^{(t)}[n] = \mathcal{T}^{-1}\left\{ G_d^{(t)}[k]\,\bar{\mathcal{H}}[k] \right\}$

2. $g_d^{(t+1)}[n] = \mathcal{P}_{\mathcal{W}_{g_d}}\left\{ \hat{x}_d^{(t)}[n] \right\}$:

    (a) $\hat{X}_d^{(t)}[k] = \mathcal{F}\left\{ \hat{x}_d^{(t)}[n] \right\}$

    (b) $G_d^{(t+1)}[k] = (1 - F[k])\,\hat{X}_d^{(t)}[k] + \mathcal{F}\left\{ g_d^{(0)}[n] \right\}$

    (c) $g_d^{(t+1)}[n] = \mathcal{F}^{-1}\left\{ G_d^{(t+1)}[k] \right\}$

Chapter 6

Conclusions and Future Directions

In this thesis, we have considered the recovery of amplitude coefficients and decay rates for transient structured signals under three stages of recovery. The accurate determination of these parameters is closely aligned with valuable information in many applications.

In Chapter 3, the first stage of parameter recovery was presented in which we determined the amplitude coefficients of a transient signal given both samples of the signal and the decay rates associated with the signal. This process led to the definition of the exponential basis and the concept of a transient transform. For finite-length signals, the dual exponential basis was derived using two novel algorithms such that the DTT became computationally feasible. Using these algorithms, numerical improvements were observed for a commonly solved ill-conditioned system of equations. The DTT then became the standard procedure to solve for the amplitude coefficients for the second stage of parameter recovery once the decay rates were recovered.

Chapter 4 extended the generation of the dual exponential basis such that several approximate dual bases were defined by trading off between properties. For example, the solution to an unconstrained optimization problem provided a stable approximate dual basis at the cost of reducing transient resolution. Section 4.4 provided bounds on guaranteed transient resolution for any approximate dual exponential basis. The use of Inner Product Shaping and orthogonal polynomials were unsuccessfully considered for creating a meaningful approximate transient spectrum.

The second stage of parameter recovery, for which the decay rates of a transient signal given samples of the signal, is presented in Appendix B. This appendix provides a tutorial on several existing parameter recovery algorithms for determining the decay rates. The DTT is then utilized to accurately determine the amplitude coefficients when desired as an alternative to Shanks Method.

In Chapter 5, the third and final stage of parameter recovery was presented in which we established a framework for determining the amplitude coefficients and decay rates of a transient signal given samples of the lowpass-filtered transient signal. We saw that this framework heavily relied upon selecting an overdetermined model order to be meaningfully larger than the underlying transient signals model order with respect to the severity of the LPF. An alternating projection algorithm was presented in order to recover the transient signal prior to the lowpass filter using the overdetermined recovery procedure. In order to use this algorithm, a novel model order deter-

mination algorithm was presented which may be used to identify the number of parameters to be determined.

6.1 Future Directions

The DTT uses an exponential basis in order to represent a finite-length signal with any set of decay rates for the exponential basis signals. Research for understanding the effects of different spacing structures of the decay rates is needed to provide a deeper sense of the conditioning issues of exponential signals. For example, clustering of the decay rates has a significantly undesirable effect on the numerical stability of the dual exponential basis. Understanding the consequences of different spacing structures may lead to improved stability of the DTT. In addition, it has been observed by simulation that different spacings of the decay rates often has noticeable effects on spectral leakage and resolution of the approximate DTTs in Chapter 4. Finally, identifying applications that can take advantage of transient filtering is a potentially rich area for future research encompassing a broad range of fields.

Another potential area for future research involves utilizing functional composition such that additional parameter recovery algorithms become of use. For example, the forward-backward method of linear prediction exploits the fact that a sample from a periodic time-series is linearly predictable using either previous or future sample values [15]. This results in recovered parameters which are guaranteed to be on the unit circle in the z-plane. Borrowing inspiration from this algorithm, it may be advantageous to utilize the information that the poles are located on the real axis to ensure that recovered parameters are all on the real axis as well. If the z-domain representation of the transient signal could be modified in such a way that the real line over $(0, 1)$ was uniquely mapped to the unit circle, then the forward-backward method would produce pole estimates on the unit circle that would then be mapped back to the real line. One way of implementing this transformation is to use the functional composition given by

$$\tilde{X}_d(z) \;=\; X_d\left(G\left(z\right)\right) \tag{6.1.1}$$

$$=\; \sum_n x_d[n]\left(\frac{1}{G(z)}\right)^n \tag{6.1.2}$$

where $X_d\left(z\right)$ is the z-transform of the transient signal $x_d\left[n\right]$, $G(z)$ is a warping function, and we denote $\tilde{x}_d\left[n\right]$ as the inverse z-transform of $\tilde{X}_d\left(z\right)$. If $G(z)$ is a rational function in z, then Eq. (6.1.1) represents a rational composition when both inner and outer functions are ratio's of polynomials in z. Table 6.1 lists several potential mapping functions and the image of one of the N^{th} roots of unity.

Warping Function $G(z)$	Image of one of the N^{th} roots of unity $e^{\frac{2\pi}{N}k}$
$G_1(z) = \frac{1}{2}z + \frac{1}{2}$	$\cos\left(\frac{\pi k}{N}\right)e^{j\frac{\pi k}{N}}$
$G_2(z) = \frac{1}{4}(z^{-1} + 2 + z)$	$\cos^2\left(\frac{\pi k}{N}\right)$
$G_3(z) = \frac{1}{2}(1+K)z + \frac{1}{2}(1-K)z^{-1}$	$\cos\left(\frac{2\pi k}{N}\right) + j(K)\sin\left(\frac{2\pi k}{N}\right)$
$G_4(z) = \frac{K}{z+a}$	$\dfrac{K\left(\cos\left(\frac{2\pi k}{N}\right)+a-j\sin\left(\frac{2\pi k}{N}\right)\right)}{\left(\cos\left(\frac{2\pi k}{N}\right)+a\right)^2+\sin^2\left(\frac{2\pi k}{N}\right)}$

Table 6.1: Warping functions and the image of one of the N^{th} roots of unity

Figure 6.1.1 depicts the relationship between computing the DFT of the signal $\tilde{x}_d[n]$ with the equivalent z-plane evaluation of $x_d[n]$ for different warping functions $G(z)$. Note that $G_3(z)$ is shown with a value of $K = \frac{1}{2}$ and $G_4(z)$ uses $K = -1.5$ and $a = -3$.
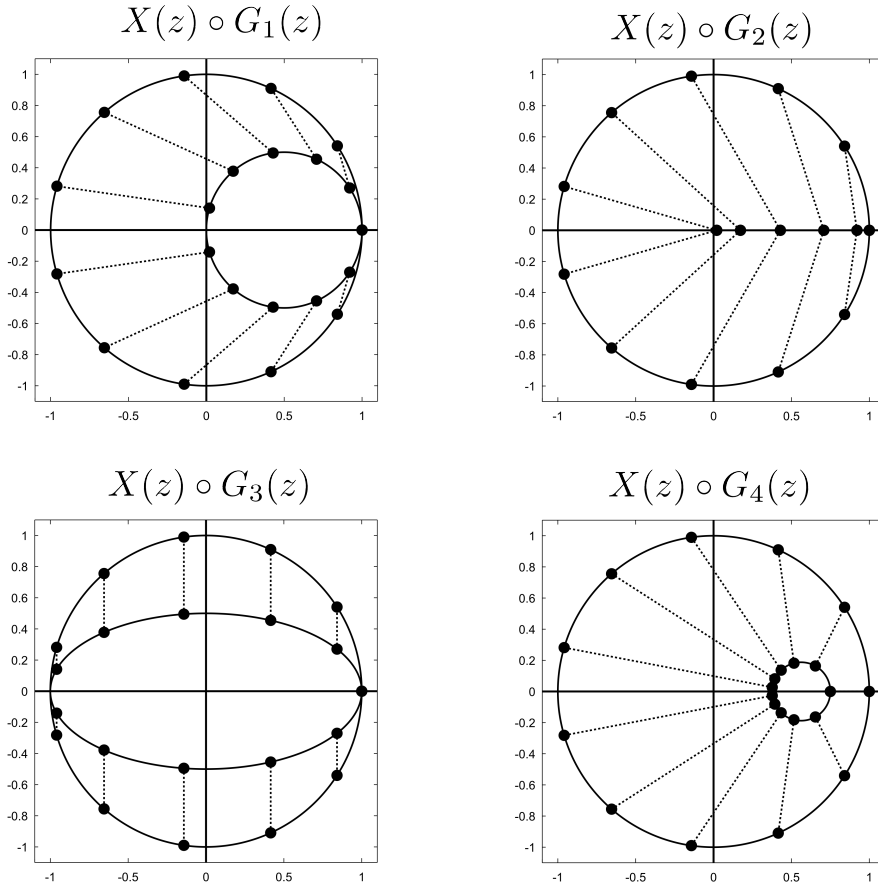

Figure 6.1.1: Mappings for rational functional composition

101

Appendix A

Proof of Uniqueness for ℓ_2 Minimization

The uniqueness of the solution for the linearly and decay rates in Eq. (2.5.6) is verified in the following proposition. In order to see this, we define the space of all d^{th}-order transient structured signals as

$$\mathcal{V}_d = \{x_d[n] \; : \; x_d[n] \text{ of the form of Eq. (2.1.2)}\}.$$

Proposition. *If $g_d[n] = f_{lp}[n] * x[n]$ and $x[n] \in \mathcal{V}_d$, then there exists a unique $x[n]$ that produces $g_d[n]$.*

Proof. Let $x^{(1)}[n]$, $x^{(2)}[n] \in \mathcal{V}_d$ where $x^{(1)}[n] \neq x^{(2)}[n]$. Denote the DTFTs of $x^{(1)}[n]$ and $x^{(2)}[n]$ as $X_1\left(e^{j\omega}\right) = \frac{N_1(e^{j\omega})}{D_1(e^{j\omega})}$ and $X_2\left(e^{j\omega}\right) = \frac{N_2(e^{j\omega})}{D_2(e^{j\omega})}$, respectively, where each $N_1\left(e^{j\omega}\right)$, $N_2\left(e^{j\omega}\right)$, $D_1\left(e^{j\omega}\right)$, and $D_2\left(e^{j\omega}\right)$ are polynomials in $e^{-j\omega}$ of maximum order d. Let $G_d\left(e^{j\omega}\right)$ be the DTFT of $g_d[n]$. By definition

$$G_d\left(e^{j\omega}\right) = \frac{N_1\left(e^{j\omega}\right)}{D_1\left(e^{j\omega}\right)} = \frac{N_2\left(e^{j\omega}\right)}{D_2\left(e^{j\omega}\right)}, \quad |\omega| < \omega_c.$$

Next, define $\Delta X\left(e^{j\omega}\right)$ as

$$
\begin{aligned}
\Delta X\left(e^{j\omega}\right) &= \frac{N_1\left(e^{j\omega}\right)}{D_1\left(e^{j\omega}\right)} - \frac{N_2\left(e^{j\omega}\right)}{D_2\left(e^{j\omega}\right)} \\
&= \frac{N_1\left(e^{j\omega}\right) D_2\left(e^{j\omega}\right) - N_2\left(e^{j\omega}\right) D_1\left(e^{j\omega}\right)}{D_1\left(e^{j\omega}\right) D_2\left(e^{j\omega}\right)} \\
&= \frac{\tilde{N}\left(e^{j\omega}\right)}{D_1\left(e^{j\omega}\right) D_2\left(e^{j\omega}\right)} \\
&= 0
\end{aligned}
$$

for $|\omega| < \omega_c$. Therefore $\tilde{N}\left(e^{j\omega}\right)$ may be written as a polynomial of maximum order $2d$, given by

$$\tilde{N}\left(e^{j\omega}\right) = \sum_{l=0}^{2d} c_l e^{-j\omega l}.$$

Let $\omega_1, \ldots, \omega_{2d+1}$ be $2d+1$ distinct samples of ω in $[0, \omega_c)$. Consider the set of equations

$$\tilde{N}\left(e^{j\omega_i}\right) = \sum_{l=0}^{2d} c_l \gamma_i^l, \quad i = 1, \ldots, 2d+1,$$

where $\gamma_i = e^{-j\omega_i}$. Expanding this into matrix form yields

$$
\begin{bmatrix}
\tilde{N}\left(e^{j\omega_1}\right) \\
\tilde{N}\left(e^{j\omega_2}\right) \\
\vdots \\
\tilde{N}\left(e^{j\omega_{2d+1}}\right)
\end{bmatrix}
=
\begin{bmatrix}
1 & \gamma_1 & \cdots & \gamma_1^{2d} \\
1 & \gamma_2 & \cdots & \gamma_2^{2d} \\
\vdots & \vdots & \ddots & \vdots \\
1 & \gamma_{2d+1} & \cdots & \gamma_{2d+1}^{2d}
\end{bmatrix}
\begin{bmatrix}
c_0 \\
c_1 \\
\vdots \\
c_{2d}
\end{bmatrix}
=
\begin{bmatrix}
0 \\
0 \\
\vdots \\
0
\end{bmatrix}.
$$

Written in matrix form this is $\mathbf{0} = \Gamma \mathbf{c}$ where Γ is a Vandermonde matrix of full rank and therefore has only the trivial vector in it's null-space. The only set of c_l that satisfy the above equation are $c_l = 0$ for $0 \leq l \leq 2d$. Thus we have that $\tilde{N}\left(e^{j\omega}\right) = 0$ for all ω, which implies

$$X_1\left(e^{j\omega}\right) = \frac{N_1\left(e^{j\omega}\right)}{D_1\left(e^{j\omega}\right)} = \frac{N_2\left(e^{j\omega}\right)}{D_2\left(e^{j\omega}\right)} = X_2\left(e^{j\omega}\right)$$

Therefore $x^{(1)}[n] = x^{(2)}[n]$ which contradicts the original assumption, so the input to produce $g_d[n]$ must be unique.

$\square$

One result of the above proposition is that Eq. (2.5.8) is simplified to

$$(\hat{\alpha}, \hat{\sigma}) = \arg\min_{\alpha, \sigma} \left\| G_d\left(e^{j\omega}\right) - X_d\left(e^{j\omega}\right) F_{lp}\left(e^{j\omega}\right) \right\|_2^2$$

as the uniqueness of the solution has been established.

Appendix B

Parameter Recovery Algorithms

Parameter modeling is a powerful tool with which a set of observed samples is fit to a predefined model of a specific structure using only a limited number of parameters. Often, the number of parameters determined is fewer than the number necessary for an exact representation of the data. Therefore the parametrization results in some form of modeling error. In this case, the model parameters are determined through minimization of an error function for some definition of error. Solving this formulation results in the set of model parameters that best represent the observed samples in some sense. Different error criteria result in different recovery algorithms and solutions.

The general procedure of parameter recovery is summarized into two primary stages: model selection and parameter determination. First, the structure and order of a parametric model need to be selected. Once the structure is fixed, the model order must be carefully chosen when it is not known *a priori*. Selecting too low an order leads to a poor representation of the data, while selecting too high an order, in general, leads to over fitting to the observed samples, including fitting to noise. Generally, if the data does not exactly fit the model assumed, diminishing reductions in modeling error are obtainable by continually increasing the model order. The order selection process can be thought of as a tradeoff between parameter parsimony and tolerable modeling error. Chapter 5 proposes an algorithm for estimating model order using insight based upon the transient signal structure, and where the model order is computed as a function of the available data. Once the model structure and order are fixed, the recovery algorithm used for determining the model parameters must be selected in order to minimize an error criterion. The assumptions made in formulating each recovery algorithm are very important in practice when deciding which recovery algorithm is appropriate for a given application; selecting an algorithm whose inherent assumptions best match known facts about the source from which the data is obtained often leads to better results. For the remainder of this chapter the model structure corresponding to the source of data is either the transient structure in Eq. (2.1.2) or Eq. (2.1.7). Additionally, the model order is assumed to be fixed to d. The algorithms presented in this appendix are directed at solving the problems stated in both Definitions 1 and 3 alike.

B.1 Prony's Method

After defining Eq. (2.1.5), we saw that directly solving for the set of decay rates, given a set of observed sample values, requires solving a non-linear system of equations. Prony's method reformulates this problem such that the decay rates are found by solving a linear system of equations instead. This method originates from recognizing that a signal with the structure in Eq. (2.1.2) is also the structure of the solution to a homogeneous Linear Constant Coefficient Difference Equation (LCCDE) of order d. [21] Specifically, Prony's method fits N observed samples to a pre-determined exponential model, containing d exponentials, through a three stage procedure, outlined as follows:

1. Solve a homogeneous LCCDE of order d for the coefficients $\{\theta_k\}_{1:d}$.

2. Use the coefficients found in stage one to form a d^{th} order polynomial whose roots give the decay rates $\{\sigma_k\}_{1:d}$.

3. Use the decay rates found in stage two to solve the linear system of equations in Eq. (2.1.5) for the amplitude coefficients $\{\alpha_k\}_{1:d}$.

The algorithms presented in this chapter only describe the recovery process for the decay rates, i.e., stages 1 and 2 above. Shanks method then uses these methods to solve stage 3 in a least squares manner. However, for transient structured signals, this involves solving the ill-conditioned system of equations in Eq. (2.1.5). Solutions obtained by using the algorithms proposed in Chapter 3 often have lower numerical error, and thus make an improvement over the stage 3 recovery problem.

To develop the insight used in Prony's method for stages 1 and 2, define a characteristic polynomial as

$$\Theta(z) = \prod_{k=1}^{d}(1 - \sigma_k z^{-1}) \tag{B.1.1}$$

$$= \sum_{m=0}^{d} \theta_m z^{-m} = 0, \tag{B.1.2}$$

where we choose $\theta_0 = 1$ without loss of generality. Next, the result of manipulating Eq. (2.1.2) by delaying by m, multiplying by θ_m, and summing over m, yields

$$\sum_{m=0}^{d} \theta_m x_d[n-m] = \sum_{k=1}^{d} \alpha_k \sigma_k^n \sum_{m=0}^{d} \theta_m \sigma_k^{-m} \tag{B.1.3}$$

$$= \sum_{k=1}^{d} \alpha_k \sigma_k^n \Theta(\sigma_k) = 0 \tag{B.1.4}$$

for $n \geq d$. For notational convenience, define $\theta = [\theta_1, \cdots, \theta_d]^T$. This equation says that the coefficients $\{\theta_k\}_{1:d}$ are exactly the same coefficients found by solving a d^{th} order homogeneous LCCDE for θ, starting from $n = d$. Prony's method is not restricted to solving for the parameters in Eq. (2.1.2) for only transient signals, but is also applicable for more general curve fitting with complex exponentials and complex amplitudes. The LCCDE is defined, over some interval of support, by

$$x_d[n] + \theta_1 x_d[n-1] + \theta_2 x_d[n-2] + \cdots + \theta_d x_d[n-d] = 0. \tag{B.1.5}$$

Different intervals of support lead to different algorithms and consequently different solutions. The zeros of the characteristic polynomial $\Theta(z)$, i.e., $\{\sigma_k\}_{1:d}$, are the decay rates in Eq. (2.1.2). Therefore if the coefficients of the LCCDE are correctly identified, then the decay rates of the signal $x_d[n]$ may be correctly determined. Two intervals of support are considered in the following two subsections.

B.1.1 Prony's Original Method

The original formulation of Prony's method defines the interval of support for Eq. (B.1.5) to be $d \leq n \leq N - 1$ where $N = 2d$. As a result, Prony's original method exactly fits $2d$ observed sample values to a linear combination of d exponentials. There is no modeling error either with respect to representing the data in this method. However, when the data doesn't exactly fit the transient structure, e.g., $\bar{x}_d[n]$, then modeling error with respect to the transient structure is unavoidable. Solving for the coefficients θ, in order to form the characteristic polynomial, requires the solution to the linear system of equations given by

$$-\begin{bmatrix} x_d[d] \\ x_d[d+1] \\ \vdots \\ x_d[2d-1] \end{bmatrix} = \begin{bmatrix} x_d[d-1] & x_d[d-2] & \cdots & x_d[0] \\ x_d[d] & x_d[d-1] & \cdots & x_d[1] \\ \vdots & \vdots & \ddots & \vdots \\ x_d[2d-2] & x_d[2d-3] & \cdots & x_d[d-1] \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_d \end{bmatrix}. \tag{B.1.6}$$

The solution $\{\theta_k\}_{1:d}$ is extended to include $\theta_0 = 1$ and is subsequently substituted into Eq. (B.1.2). Then a polynomial factoring algorithm solves for the zeros of $\Theta(z)$. It has been observed that the first two stages of Prony's original method perform poorly in estimating the correct values of $\{\sigma_k\}_{1:d}$ when the observed samples contain additive noise, i.e., $\{\bar{x}_d[n]\}_{0:2d-1}$. [8] This is attributed to the fact that this algorithm makes no separate estimates of the noise and the underlying signal, but fits exactly to the noise included in each sample. As a result, Prony's original method achieves successful recovery as in Definition 1 exactly and performs poorly in Definition 3. It is for Definition 3 that the extended Prony's method makes improvement upon this original formulation.

B.1.2 Extended Prony's Method

Consider the case where the number of exponentials to be modeled is fixed to d while $N > 2d$

samples of a transient signal are available. Prony's original method cannot utilize any of the additional $N - 2d$ samples without increasing the model order to $\hat{d} > d$. Extending the system of equations in Eq. (B.1.6) to include all available data while holding the model order fixed results in an overdetermined system of equations given by

$$
- \begin{bmatrix} x_d[d] \\ x_d[d+1] \\ \vdots \\ x_d[N-1] \end{bmatrix} = \begin{bmatrix} x_d[d-1] & x_d[d-2] & \cdots & x_d[0] \\ x_d[d] & x_d[d-1] & \cdots & x_d[1] \\ \vdots & \vdots & \ddots & \vdots \\ x_d[N-2] & x_d[N-3] & \cdots & x_d[N-d-1] \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_d \end{bmatrix}. \tag{B.1.7}
$$

Eq. (B.1.7) is compactly written as $-x_{d,0} = X\theta$, and the least squares solution, i.e., $\hat{\theta}_{LS} = -(X^T X)^{-1} X^T x_{d,0}$, is referred to as the Extended Prony Method solution. [20] The same procedure for producing the decay rates as the original Prony method is then used here, however the estimates for the decay rates may be different.

A common interpretation of the least squares solution to Eq (B.1.7) is to select the vector θ that minimizes $\|X\theta + x_{d,0}\|_2^2 = \|e\|_2^2$ where $e[n]$ is the parametric modeling error, $e[n] = x_d[n] - \hat{x}_d[n]$, for $0 \le n \le N - 1$. This method has been observed to outperform the original Prony method in the presence of noise, such as in the formulation of Definition 3.

B.2 Tufts-Kumaresan Method

The Tufts-Kumaresan method is a natural extension built upon the framework of the stage one extended Prony method, detailed above. [12] Towards developing this extension, Eq. (B.1.7) is re-written as

$$
-\bar{x}_{d,0} = \bar{X}\theta \tag{B.2.1}
$$

$$
0 = \begin{bmatrix} \bar{x}_{d,0} \mid \bar{X} \end{bmatrix} \begin{bmatrix} 1 \\ \theta \end{bmatrix} \tag{B.2.2}
$$

where $\begin{bmatrix} \bar{x}_{d,0} \mid \bar{X} \end{bmatrix}$ is a concatenated Toeplitz matrix and $\bar{x}_d[n]$ has been used in place of $x_d[n]$ to emphasize the advantage this method provides when the data contains additive noise. The key insight exploited in the Tufts-Kumaresan algorithm is to utilize the low rank of the concatenated data matrix in Eq. (B.2.2) when the samples do not contain additive noise in order to perform a de-noising like procedure. This is a more realistic formulation than in the extended Prony method because the sample values that populate $\bar{x}_{d,0}$ also populate $\bar{X}$. Previously, error was only assumed to be present in $\bar{x}_{d,0}$. This insight suggests an unstructured Total Least Squares (TLS) solution to Eq. (B.2.2). Systematically, this means solving the system of equations for the coefficients of the characteristic polynomial while taking into account that there is noise in both the data matrix $\bar{X}$ as well as the observation vector $\bar{x}_{d,0}$.

The advantage of this algorithm becomes apparent by writing Eq. (B.2.2) as

$$
\left[\bar{x}_{d,0} \mid \bar{X}\right]
\begin{bmatrix} 1 \\ \theta \end{bmatrix}
=
\begin{bmatrix}
\bar{x}_{d,0}\,[d] & \bar{x}_d\,[d-1] & \bar{x}_d\,[d-2] & \cdots & \bar{x}_d\,[0] \\
\bar{x}_{d,0}\,[d+1] & \bar{x}_d\,[d] & \bar{x}_d\,[d-1] & \cdots & \bar{x}_d\,[1] \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
\bar{x}_{d,0}\,[N-1] & \bar{x}_d\,[N-2] & \bar{x}_d\,[N-3] & \cdots & \bar{x}_d\,[N-d-1]
\end{bmatrix}
\begin{bmatrix} 1 \\ \theta_1 \\ \vdots \\ \theta_d \end{bmatrix}
= 0,
$$

(B.2.3)

and noting that in the noiseless case, i.e., the problem in Definition 1, the observed samples that populate the concatenated data matrix produce a matrix with a null space of dimension one, i.e., $\dim\left\{\mathcal{N}\left([x_{d,0} \mid X]\right)\right\} = 1$. Solving for the non-trivial vector in the null space results in a scalar multiple of the concatenated vector $\left[1 \mid \theta^T\right]^T$. Re-scaling of this non-trivial vector, such that the leading coefficient is 1, is an unnecessary step in forming the characteristic polynomial as the roots of a polynomial are invariant to scalings of the coefficients by a constant, though appropriate scaling may help numerical errors when solving for the roots of large polynomials. It is straightforward to verify that this algorithm yields the same solution as both versions of the Prony method in the noiseless case.

When the observed samples, $\{\bar{x}_d\,[n]\}_{0:N-1}$, contain additive noise, i.e., the problem in Definition 3, the joint matrix $\left[\bar{x}_{d,0} \mid \bar{X}\right]$ will have rank $d+1$ with high probability. Consequently, the only vector in the null space of this noisy concatenated matrix is the trivial solution, i.e., $\mathcal{N}\left(\left[\bar{x}_{d,0} \mid \bar{X}\right]\right) = 0$. To overcome this, the Tufts-Kumaresan algorithm finds the smallest perturbation matrix, in a Frobenius norm sense, of the joint matrix which yields a unit rank reduction. The Eckart-Young Theorem proves that the closest rank reduced matrix is easily obtainable through manipulating the SVD of the joint matrix. [23] To find the resulting matrix, consider the SVD of the joint matrix, given by

$$
\left[\bar{x}_{d,0} \mid \bar{X}\right] = \sum_{k=1}^{d+1} \pi_k u_k v_k^H
\tag{B.2.4}
$$

where $\pi_1 \geq \cdots \pi_d \geq \pi_{d+1} > 0$. The estimate of the noise is then given by the unit rank matrix $\pi_{d+1} u_{d+1} v_{d+1}^H$. The rank 1 reduced matrix, $\left[\bar{x}_{d,0} \mid \bar{X}\right]_d$, is then given by

$$
\left[\bar{x}_{d,0} \mid \bar{X}\right]_d = \sum_{k=1}^{d} \pi_k u_k v_k^T.
\tag{B.2.5}
$$

The smallest singular value is assumed to be due to additive noise and not one of the transient components in order to interpret this as a de-noising process. Further, the de-noising is guaranteed to be unique if $\pi_d > \pi_{d+1}$.

Reconstructing the de-noised matrix and solving for its null space is unnecessary; the null space of the rank reduced matrix is v_{d+1}^H. Therefore, computing the SVD of the joint matrix in Eq. (B.2.3) results in the coefficients necessary for the characteristic polynomial to be formed, i.e,

$v_{d+1}^H \propto [1 \mid \theta]^T$.

This algorithm generally out performs both versions of Prony's method when noise is present in the samples. One drawback of this algorithm is that the de-noising process does not consider the structure known for the joint matrix in the noiseless case. This is demonstrated by noting that, for all i,j such that $\left[\bar{x}_{d,0} \mid \bar{X}\right]_{i,j} = \left[\bar{x}_{d,0} \mid \bar{X}\right]_{i+1,j+1}$ before the de-noising process, there is no guarantee that the resulting rank-reduced matrix has the Toeplitz structure, which is a property of the rank d joint matrix formed by the underlying transient signal structure, i.e., $[x_{d,0} \mid X]$. This is equivalent to saying that the joint Toeplitz matrix $\left[\bar{x}_{d,0} \mid \bar{X}\right]$ is formed by adding the noiseless Toeplitz matrix $[x_{d,0} \mid X]$ and a Toeplitz noise matrix but the de-noising algorithm produces a non-Toeplitz estimate both the noise and underlying signal.

B.3   Cadzow's Method

Cadzow's method builds additionally upon the Tufts-Kumaresan framework by artificially enforcing properties that the noiseless transient signal, $x_d[n]$, is known to possess. [5] As discussed in Section B.2, the noiseless joint matrix $[x_{d,0} \mid X]$ has two intrinsic attributes which are exploited in this algorithm: a rank of d and a Toeplitz structure. The algorithm presented in this section uses an iterative procedure to alternate between enforcing these two attributes. A structured TLS projection to the closest Toeplitz matrix of rank d is ideally desired, however, enforcing both of these constraints simultaneously is difficult. The Tufts-Kumaresan method provided a projection of the noisy joint data matrix to the closest matrix of rank d, in a Frobenius norm sense, without imposing any structure on the resulting matrix. Generally, the resulting projection does not have the Toeplitz structure. The iterative procedure used by Cadzow's method is presented next, in which a method for imposing the Toeplitz structure without any constraints on the resulting matrices rank is presented.

First, define the mapping f as the projection of a noisy rank $d+1$ joint data matrix to the closest rank d matrix in the Frobenius norm sense. Therefore, the base iteration of this algorithm is given by

$$\left[\bar{x}_{d,0} \mid \bar{X}\right]_d^{(0)} = f\left(\left[\bar{x}_{d,0} \mid \bar{X}\right]\right) \tag{B.3.1}$$

$$= \sum_{k=1}^d \pi_k u_k v_k^T. \tag{B.3.2}$$

Next, in order to impose the Toeplitz structure on $\left[\bar{x}_{d,0} \mid \bar{X}\right]_d^{(0)}$, define the mapping h such that the resulting matrix has, for each of it's diagonal entries, the average value from the corresponding diagonal of the input matrix. This mapping will take any general input matrix and produce a Toeplitz matrix, e.g.,

$$h\left(\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}\right) = \begin{bmatrix} 3 & 4 & 3 \\ 4 & 3 & 4 \end{bmatrix}. \tag{B.3.3}$$

Using these two mappings, the iterative procedure from the n^{th} iteration to the $(n+1)^{\text{st}}$ is given by:

$$
\begin{aligned}
\left[\bar{x}_{d,0} \mid \bar{X}\right]_T^{(n)} &= h\left(\left(\left[\bar{x}_{d,0} \mid \bar{X}\right]_d^{(n)}\right)\right) \quad \text{(project to Toeplitz strucutred matrices)} & \text{(B.3.4)} \\
\left[\bar{x}_{d,0} \mid \bar{X}\right]_d^{(n+1)} &= f\left(\left[\bar{x}_{d,0} \mid \bar{X}\right]_T^{(n)}\right) \quad \text{(project to rank } d \text{ matrices)}. & \text{(B.3.5)}
\end{aligned}
$$

Note that after projecting to the space of Toeplitz structured matrices, the resulting matrix $\left[\bar{x}_{d,0} \mid \bar{X}\right]_T^{(n)}$, with high probability, has a rank of $d+1$. This iterative procedure has been shown to converge, as the number of iterations grows large, to a solution which possess both desired attributes, i.e., a rank of d and a Toeplitz structure. In practice, the number of iterations is often fixed to some number L. The result of the L^{th} iteration, $\left[\bar{x}_{d,0} \mid \bar{X}\right]_d^{(L)}$ is guaranteed to have a non-trivial null space but not guaranteed to have a Toeplitz structure. Given this matrix, the same procedure for finding the decay rates from the Tufts-Kumaresan method is used by computing the null-space. Cadzow's iterative method has been shown to work extremely well in practice for problems such as harmonic retrieval in noise. However this must be balanced with the increased complexity of performing an iterative alternating projection algorithm.

B.4   Autoregressive Algorithms

An autoregresive model assumes that a time-series current value depends linearly on its previous values. This section provides the general procedure for computing the stage one solution for the characteristic polynomial coefficients $\{\theta_k\}_{1:d}$ by solving different autoregressive models. [15] To being this exposition, define a d^{th} order linear predictor, $x_{lp,d}[n]$, as

$$
x_{lp,d}[n] = -\sum_{k=1}^{d} \theta_k x_d[n-k]. \tag{B.4.1}
$$

The structure of this linear predictor is known as the forward or causal predictor because it computes each sample as a linear combination of strictly previous samples. The error criterion used in linear prediction is defined to be the sum of squares of $e_{lp}[n]$ where

$$
e_{lp}[n] = x_d[n] - x_{lp,d}[n] = x_d[n] + \sum_{k=1}^{d} \theta_k x_d[n-k]. \tag{B.4.2}
$$

B.4.1   The Yule-Walker Method

The Yule-Walker method is an all-pole parameter modeling algorithm based upon a least squares inverse model. This method is straightforward to derive from the orthogonality principle, and thus a derivation of the Yule-Walker Equations, or Autocorrelation Normal Equations, is not presented

here, but can be found in [18]. The linear prediction coefficients, $\{\theta_k\}_{1:d}$, of a time-series are found by solving the system of equations given by

$$
\begin{bmatrix}
r_{\bar{x}_d \bar{x}_d}[1] \\
r_{\bar{x}_d \bar{x}_d}[2] \\
\vdots \\
r_{\bar{x}_d \bar{x}_d}[d]
\end{bmatrix}
=
\begin{bmatrix}
r_{\bar{x}_d \bar{x}_d}[0] & r_{\bar{x}_d \bar{x}_d}[1] & \cdots & r_{\bar{x}_d \bar{x}_d}[d-1] \\
r_{\bar{x}_d \bar{x}_d}[1] & r_{\bar{x}_d \bar{x}_d}[0] & \cdots & r_{\bar{x}_d \bar{x}_d}[d-2] \\
\vdots & \vdots & \ddots & \vdots \\
r_{\bar{x}_d \bar{x}_d}[d-1] & r_{\bar{x}_d \bar{x}_d}[d-2] & \cdots & r_{\bar{x}_d \bar{x}_d}[0]
\end{bmatrix}
\begin{bmatrix}
\theta_1 \\
\theta_2 \\
\vdots \\
\theta_d
\end{bmatrix},
\tag{B.4.3}
$$

where $r_{\bar{x}_d \bar{x}_d}[m]$ is the true autocorrelation sequence of $\bar{x}_d[n]$. In practice the autocorrelation sequence is rarely known, therefore it must be estimated using the finite amount of data available. One autocorrelation estimator, using the data $\{\bar{x}_d[n]\}_{0:N-1}$, is defined by

$$
\hat{r}_{\bar{x}_d \bar{x}_d}[m] = \frac{1}{N} \sum_{n=0}^{N-|m|-1} \bar{x}_d[n+|m|]\,\bar{x}_d[n].
\tag{B.4.4}
$$

This estimator is readily shown to be biased with $\mathcal{E}\{\hat{r}_{\bar{x}_d \bar{x}_d}[m]\} = \left(1 - \frac{|m|}{N}\right) r_{\bar{x}_d \bar{x}_d}[m]$. For a fixed lag m, the variance asymptotically approaches zero as the number of samples N grows large. A second autocorrelation estimator is defined by modifying the biased estimator to form an unbiased estimator structured as

$$
\tilde{r}_{\bar{x}_d \bar{x}_d}[m] = \frac{1}{N-|m|} \sum_{n=0}^{N-|m|-1} \bar{x}_d[n+|m|]\,\bar{x}_d[n].
\tag{B.4.5}
$$

The unbiased auto-correlation estimator yields sequences for which there is increasing statistical uncertainty at increasingly large lags. When either Eq. (B.4.4) or Eq. (B.4.5) is used, the algorithm is referred to as the Yule-Walker method of autoregressive parameter recovery using a biased or unbiased estimator, respectively.

Both of the autocorrelation estimators yield identical values for lag $m = 0$, i.e., $\hat{r}_{\bar{x}_d \bar{x}_d}[0] = \tilde{r}_{\bar{x}_d \bar{x}_d}[0]$. Despite the estimator in Eq. (B.4.4) being biased, it is often preferred because it produces an autocorrelation matrix that is guaranteed to be positive semi-definite. This results in a stable set of estimated pole locations, which can be seen by using the biased estimator to solve for the k-parameters of a lattice filter, in which the k-parameters are guaranteed to be strictly less than one. [18] The estimator in Eq. (B.4.5) may yield autocorrelation sequences which are invalid, in that they violate the the property that $r_{x_d x_d}[0] \geq r_{x_d x_d}[1]$. For example, consider the sequence $\{x_d[0], x_d[1], x_d[2]\} = \{1, 2, \frac{3}{2}\}$. The resulting unbiased autocorrelation sequence samples are $\{\tilde{r}_{x_d x_d}[0], \tilde{r}_{x_d x_d}[1]\} = \{2.4167, 2.5\}$.

B.4.2   Linear Prediction

For each value of n, the linear prediction error is an equation of d unknowns. Immediately, we have that $e_{lp}[0] = x_d[0]$ because the input sample at $n = 0$ has no preceding samples from which to be predicted. Writing the linear prediction error as a system of equations, over the interval

$0 \leq n \leq N + d - 1$, yields

$$
\begin{bmatrix}
e\,[0] \\
e\,[1] \\
\vdots \\
e\,[d] \\
\vdots \\
e\,[N-1] \\
\vdots \\
e\,[N+d-1]
\end{bmatrix}
=
\begin{bmatrix}
x_d\,[0] & 0 & \cdots & 0 \\
x_d\,[1] & x_d\,[0] & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
x_d\,[d] & x_d\,[d-1] & \cdots & x\,[0] \\
\vdots & \vdots & \ddots & \vdots \\
x_d\,[N-1] & x_d\,[N-2] & \cdots & x_d\,[N-d-1] \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & x_d\,[N-1]
\end{bmatrix}
\begin{bmatrix}
1 \\
\theta_1 \\
\vdots \\
\theta_d
\end{bmatrix} .
\qquad \text{(B.4.6)}
$$

This equation is compactly denoted as $e_{lp} = D^{(d)} \begin{bmatrix} 1 \\ \theta \end{bmatrix}$, and the error criterion to be minimized is simply given by

$$
\epsilon_{lp}^{(d)} = \sum_{n \in \mathcal{I}_{lp}} |e_{lp}\,[n]|^2 . \qquad \text{(B.4.7)}
$$

Different choices for the interval $\mathcal{I}_{lp}$ result in different linear prediction algorithms. When the signal $x_d\,[n]$ is only known over the interval $0 \leq n \leq N - 1$, the zeros present in the upper and lower anti-diagonals of the data matrix $D^{(d)}$ above represent an implicit assumption of rectangular windowing on an infinitely long underlying signal. It is often the case that this assumption, that the signal takes value 0 for $n < 0$ and $n \geq N$, is unrealistic. However, for transient structured signals, the trailing zeros are only marginally unrealistic, especially for data sets with large N. Using a discrete representation transient signals never reach the value 0, but when quantization is taken into account, noise free digital transient signals typically reach 0 relatively quickly.

The correlation method of linear prediction defines the interval $\mathcal{I}_{lp} = \{0 \leq n \leq N + d - 1\}$. It is straightforward to show, from Eq. (B.4.7), that this choice of interval produces the same solution as if we chose the interval $\mathcal{I}_{lp} = \{n \in (-\infty, \infty)\}$. Further, Eq. (B.4.6) uses $e_{lp}\,[n] = 0$ for $1 \leq n \leq N + d - 1$. Therefore, the least squares solution is given by

$$
\begin{bmatrix} 1 \\ \theta \end{bmatrix} = \left(\left[D^{(d)} \right]^T D^{(d)} \right)^{-1} \left[D^{(d)} \right]^T e_{lp}. \qquad \text{(B.4.8)}
$$

With this form of the solution, we find that $\left[D^{(d)} \right]^T D^{(d)}$ is structurally equivalent to the same Toeplitz matrix derived in Eq (B.4.3) when the biased estimator of Eq (B.4.4) is used. Therefore the Yule-Walker method of autoregressive parameter recovery using a biased estimator is structurally equivalent to the correlation method of linear prediction.

The covariance method of linear prediction avoids the pre- and post-windowing of the data which was assumed in the correlation method. [15] The use of the term covariance is misleading; this technique does not use covariance in the traditional sense, i.e., correlation with the product of

the means subtracted. The difference between this method and the correlation method comes from the choice of the region of support for the error term. The interval chosen for minimizing the sum of squares of the linear prediction error is selected to be $\mathcal{I}_{lp} = \{d \leq n \leq N - 1\}$. To understand the consequence of this choice, partition the matrix $D^{(d)}$ in Eq (B.4.6) as

$$D^{(d)} = \begin{bmatrix} L \\ T \\ U \end{bmatrix} \tag{B.4.9}$$

where the partitioned matrices are defined by

$$L^{(d)} = \begin{bmatrix} x_d[0] & 0 & \cdots & 0 \\ x_d[1] & x_d[0] & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ x_d[d-1] & x_d[d-2] & \cdots & 0 \end{bmatrix}, \; T^{(d)} = \begin{bmatrix} x_d[d] & x_d[d-1] & \cdots & x_d[0] \\ \vdots & \vdots & \ddots & \vdots \\ x_d[N-1] & x_d[N-2] & \cdots & x_d[N-d-1] \end{bmatrix}, \text{ and}$$

(B.4.10)

$$U^{(d)} = \begin{bmatrix} 0 & x_d[N-1] & \cdots & x_d[N-d-2] \\ 0 & 0 & \cdots & x_d[N-d-3] \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & x_d[N-1] \end{bmatrix}.$$

The windowing of the infinitely long underlying signal is only apparent in the lower triangular matrix L and the upper triangular matrix U. The intermediate Toeplitz matrix T does not have any windowing, and thus does not make any assumptions about the signal outside of finite set known. This gives rise to the following system of equations

$$\begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} x_d[d] & x_d[d-1] & \cdots & x_d[0] \\ \vdots & \vdots & \ddots & \vdots \\ x_d[N-1] & x_d[N-2] & \cdots & x_d[N-d-1] \end{bmatrix} \begin{bmatrix} 1 \\ \theta_1 \\ \vdots \\ \theta_d \end{bmatrix}. \tag{B.4.11}$$

This system of equations is denoted in matrix-vector form as $\tilde{e}_{lp} = T^{(d)} \begin{bmatrix} 1 \\ \theta \end{bmatrix}$. The system of equations in Eq. (B.4.11) is then re-structured to be equivalent to the system of equations in Eq. (B.1.7). Therefore the covariance method of linear prediction and the stage one equations of the extended Prony's method are structurally equivalent. It is due to this relationship that the characteristic polynomial defined in Prony's method is also referred to as the prediction polynomial as well.

Bibliography

[1] L. Adams and J. Nazareth. *Linear and Nonlinear Conjugate Gradient-related Methods*. SIAM, 1996.

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

[3] D. Bertsimas and J. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, 1997.

[4] S. Boyd and L. Vandenberghe. *Convex Optimization Convex Optimization Convex Optimization*. Cambridge University Press, 2004.

[5] J. A. Cadzow. Signal enhancement-a composite property mapping algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(1):49 –62, jan 1988.

[6] A. Z. Dodd. *The essential guide to telecommunications*. Prentice Hall, 2002.

[7] Y. C. Eldar. Least-squares inner product shaping least-squares inner product shaping least-squares inner product shaping. *Linear Algebra and its Applications*, 348(1-3):153–174, jun 2002.

[8] M. Kahn et. al. On the consistency of prony's method and related algorithms. *Journal of Computational and Graphical Statistics*, 1992.

[9] W. Gautschi. Norm estimates for inverses of vandermonde matrices. *Numer. Math.*, 23:337–347, 1975.

[10] G. Goertzel. An algorithm for the evaluation of finite trigonometric series. *The American Mathematical Monthly*, 65(1):34–35, Jan 1958.

[11] R. Kumaresan. On the zeros of the linear prediction-error filters for deterministic signals. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 1983.

[12] R. Kumaresan and D. Tufts. Estimating the parameters of exponentially damped sinusoids and pole-zero modeling in noise. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 30(6):833 – 840, dec 1982.

[13] L. Lasdon. *Optimization Theory for Large Systems*. Dover Books on Mathematics, 2011.

[14] Jae S. Lim. *Two-Dimensional Signal and Image Processing.* Prentice Hall, 1989.

[15] J. Makhoul. Linear prediction: A tutorial review. *Proceedings of the IEEE*, 63(4):561 – 580, april 1975.

[16] P. Morse and H. Feshbach. *Methods of Theoretical Physics.* McGraw-Hill, 1953.

[17] S. Cohn-Sfetcu M.R. Smith and H.A. Buckmaster. Decomposition of multicomponent exponential decays by spectral analytic techniques. *Technometrics*, 18(4):467–482, 1976.

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

[19] J. Paris. *A Framework for Non-Intrusive Load Monitoring and Diagnostics.* PhD thesis, MIT, 2006.

[20] D. Potts and M.Tasche. *An inverse probem of digital signal processing.* 2009.

[21] R. Prony. Essai éxperimental et analytique: sur les lois de la dilatabilité de la fluides élastique et sur celles de la force expansive de la vapeur de l'alkool, à différentes températures. *Journal de l'École Polytechnique Floréal et Plairial*, 1(22):24–76, 1795.

[22] L. Schenker. Pushbutton calling with a two-group voice-frequency code. *The Bell system technical journal*, 39(1):235–255, 1960.

[23] G. Stewart. On the early history of the singular value decomposition. *SIAM Review*, 1993.

[24] G. Strang. *Linear Algebra and Its Applications*, volume 4th. Brooks Cole, 2005.

[25] Lloyd N. Trefethen and David Bau III. *Numerical Linear Algebra.* SIAM, 1997.

[26] E. B. Vinberg. *A course in algebra.* American Mathematical Society, 2003.

Epilog

The unique research style within DSPG appears to me to be difficult to characterize, and may possibly be best described via an example. In the following discussion I have summarized my personal reflections on the progression of my masters thesis and its supporting research. In it I hope to convey a different story than one might expect, given the presentation of results in the thesis itself.

The results in this thesis were initially inspired by a question posed by Al during one of our first research meetings: "What can wave tunneling teach us about signal processing?" My initial response was an interesting and enjoyable dive into the literature about wave propagation and quantum physics, yet I returned to Al with more questions than answers about wave tunnelings relationship with signal processing. We spent the next several research meetings discussing variations of this question, which eventually led us to an analogy involving wave tunneling and wide-band signals passing through narrow-band channels. In essence, we agreed that because a propagating wave tunnels through an energy barrier by transforming into an evanescent wave within the barrier, it may be possible for a wide-band signal to use a similar transformation to "tunnel" through a lowpass channel. In the spirit of the Dr. Seuss cartoon in the DSPG library[1], we decided that this idea was worth pursuit. In doing so, we identified a more refined question that needed to be addressed before we tackled the wave tunneling question: Assuming we can transform a signal into an evanescent like signal, if this signal is processed by a lowpass filter is it even possible to recover the signal at the output of the filter? If the answer was no, then our analogy would have needed modification. This question clearly played an important part in the development of this thesis, as our starting point became the recovery of signal parameters post lowpass filtering of linear combination of real, decaying exponentials.

The remainder of my first year was spent learning the background and supporting material to similar questions. One of the first papers I read related to the topic of parameter identification was a 1795 publishing by Gaspard Riche de Prony, whose work laid the foundation for Appendix B. At this point I felt that a deeper familiarity with parameter modeling and linear prediction techniques would benefit me, so I began to catch up on the 200+ years of literature that followed that paper. During this time, at a weekly meeting with Al, we attempted to postulate what a filter that passes only slowly or quickly decaying exponentials would look like. At that time, I mistakenly

---

[1]"ONE plus ONE adds up to TWO if that is all you think it can do... But ONE plus ONE could equal THREE or anything else you'd like it to be!"

tried envisioning what characteristics the impulse response of such a system would have. However, the process of answering this question got me to begin to pull on a string that eventually led to the Transient Transform and algorithms for generating dual exponential bases. It wasn't for over an entire year before I returned to Al with exactly how this system would work.

The summer after my first year I interned at Bose Corporation, and was fortunate enough to be supervised by a former DSPG member, Steve Isabelle. My primary project at Bose was quite distant from my thesis research, however the modeling algorithms I had studied were of use in a side project I collaborated on. In a different context, I developed an overdetermined framework similar to the one presented in Chapter 5.

During the fall semester of my second year, in various group meetings and personal discussions with Sefa and Guolong, I developed a healthy curiosity about the applications of functional and/or polynomial composition to signal processing. As they independently developed results in this area, I briefly looked where the light was brightest and began thinking about using functional composition to warp my problem into a well-studied harmonic retrieval problem. I even got as far as some preliminary simulations before other directions began developing quickly into interesting results. I hope someone takes the functional composition approach at some point.

At the same time, I had been working on ways to solve the linear system of equations for the amplitude coefficients of a transient signal when the decay rates are known which result in as little computational error as possible, compared to built in software routines, e.g., inverse routines in MATLAB. This work was motivated by simulations I had been running using large synthetic data sets, in which the recovered amplitude coefficients were far enough from correct to raise concern. This work solidified the algorithms for dual basis generation proposed in this thesis, which eventually became the DTTT and the DTT. In discussing my findings with Al, we realized that the algorithms I developed were more broadly applicable for spectral analysis using a general exponential basis, making the DTT and the DFT special cases of my algorithms. Following this, several discussions with Tom eventually led me down an algebraic road which helped me to solidify interpretations for many of my proposed algorithms as well as provide statements of uniqueness I had otherwise been neglecting to establish. Towards the end of the fall semester, Al and I continued making many comparisons between the DTT and Fourier spectral analysis techniques in order to identify both applications for the DTT as well as understand the issues of transient spectral resolution and leakage.

As the final semester began, I developed methods for computing stable approximate transient spectra for large data sets, which led to many of the approximate algorithms and spectral resolution bounds in Chapter 4. Once these techniques had been developed, I began an intense writing period, during which discussions with Al significantly assisted with shaping the thesis into its final form. We had established that the algorithms for dual basis generation also worked for a complex exponential basis, but it wasn't until the editing process, only a couple weeks before the document was submitted, that the proofs and explanations were changed to correspond to a general exponential basis.