Solving Demodulation as an Optimization Problem

Gregory Sell and Malcolm Slaney, *Fellow, IEEE*

Abstract—We introduce two new methods for the demodulation of acoustic signals by posing the problem in a convex optimization framework. This allows the parameters of the modulator and carrier to be explicitly defined as constraints in an optimization problem. We first show the theory used to define the demodulation relationship within the rules of convex programming. Then, for the two approaches introduced, we derive specific cost functions and constraints to solve for modulators specifically motivated by perceptual rules. The methods described here perform well with simple, harmonic, and stochastic carriers, and also in the presence of noise.

Index Terms—Modulation, optimization methods.

I. INTRODUCTION

DEMODULATION divides a signal into its modulator $m(t)$ and carrier $c(t)$. In this context, the original signal is the product of the two components

$$s(t) = m(t)c(t). \tag{1}$$

The modulator is typically defined as a lower frequency signal, and the carrier is a higher frequency signal.

Demodulation, originally just used in radio communications, has become a more interesting problem because of a number of uses in speech analysis and processing. Several researches have used demodulation, albeit with limited success due to the shortcomings of the Hilbert approach [7], [16], to separate the spectral and fine time-structure in a speech signal [17], [19], [10]. They recombine these components to create so-called chimaeric sounds. Other uses of demodulation include speech enhancement [15], speech recognition [12], source separation [1], [11], and hearing devices [15].

The main hurdle for any demodulation algorithm is the mathematical fact that there are an infinite number of legitimate modulator–carrier pairs. It is the same problem as solving $xy = 8$ for x and y, because without further constraints, there is no way to limit the solution.

The ambiguity of demodulation can be illustrated with a simple example—the product of 5-, 50-, and 150-Hz sinusoids

$$s(t) = \cos(2\pi 5t)\cos(2\pi 20t)\cos(2\pi 150t).$$

One very reasonable decomposition defines the modulator as the 5-Hz sinusoid, and the rest as the carrier

$$m(t) = \cos(2\pi 5t)$$
$$c(t) = \cos(2\pi 20t)\cos(2\pi 150t).$$

An equally reasonable decomposition defines the carrier as the 150-Hz sinusoid, and the rest as the modulator

$$m(t) = \cos(2\pi 5t)\cos(2\pi 20t)$$
$$c(t) = \cos(2\pi 150t).$$

Both of these decompositions are demodulations, and, mathematically speaking, they are equally correct. There are infinite more possible decompositions that are equally right. So, a demodulation algorithm needs to include some means to select the desired solution. Without this, there is no way to distinguish between the infinite set of valid modulator–carrier pairs.

The goal of this paper is to introduce optimization theory as a vehicle for solving the demodulation problem. It is our belief that the decision process necessary to select the ideal modulator–carrier pair for a signal aligns perfectly with constrained cost function optimization, especially within a convex framework for easier computation. Furthermore, to demonstrate the power of this new approach, we develop two specific demodulators and apply them to many types of signals, including those with harmonic, stochastic and time-varying carriers. No other current demodulation algorithm can successfully deconstruct this collection of signals as well.

This paper describes past demodulation methods and their limitations, as well as our goals for a general demodulator, in Section II. Section III describes the theoretical and optimization issues, using both a linear and a logarithmic framework. Section IV describes the specific cost functions necessary to solve the optimization problem. Finally, Section V demonstrates the power of the methods on a wide range of signals, including those with harmonic, stochastic, and time-varying carriers.

II. BACKGROUND

There are a number of approaches to solve the demodulation problem. This section describes the successes and limitations of conventional (Section II-A) approaches to demodulation, and describes the goals for an ideal demodulation (Section II-B). We defer to Section V-G the role of dividing the input signal into subbands using a filterbank.

A. Past Work

A classic method for demodulation is Hilbert envelope detection. This process simply assumes the modulator is the magnitude of the analytic signal. This method certainly returns a valid decomposition from a purely mathematical perspective. However, there is no opportunity to control which solution the algorithm returns. As a result, the solution does not adhere to the definition of acoustic demodulation, to be given in Section II-B, when the carrier $c(t)$ is harmonic or noisy.

Recent work has also shown that, because the modulator magnitude $|m(t)|$ is not necessarily bandlimited, filtering the modulator or carrier can lead to artifacts upon remodulation [7], [16]. In the case where these artifacts are caused by sign errors as a result of the non-negativity of the modulator magnitude, the sign restoration method, discussed in Section IV-A4, restores the correct modulator. However, even with sign restoration, bandlimiting cannot be guaranteed, and the issues with noisy or harmonic carriers remain.

A spectrogram is a type of demodulation because the magnitude coefficient of each channel of the filterbank gives a down-sampled energy estimate over time. This method is familiar, easy to implement, and it allows for a great deal of versatility. By intelligently choosing the parameters for the spectrogram (i.e., narrowband versus wide-band), a wide range of decompositions are possible. However, this method is subject to the same time–frequency tradeoffs that any spectrogram encounters, where increasing resolution in one dimension decreases resolution in the other. Furthermore, if a filterbank channel has more than one component, there is no way to distinguish between them in the representation.

Additionally, a modulation spectrogram has been proposed [9]. The method begins with a filterbank decomposition, and then low-pass filters each channel. The modulation spectrogram is then built from the 4-Hz coefficient of each channel's short-time Fourier transform. The extension of the low-pass filter is a valuable addition, in that it adds control over the modulator, but the issues with the spectrogram listed above still persist.

Building off the work on homomorphic demultiplication [14], Atlas *et al.* [2] implemented a homomorphic demodulator after a subband decomposition. In the logarithmic domain, the multiplicative relationship in (1) becomes a summation. The logarithmic-domain modulator is extracted as the real part of the logarithmic signal, and then filtered before it is linearized with an exponential. The method is based on the same modulator definition as the Hilbert envelope, and so the same difficulties with harmonic or noisy carriers apply here as well.

A biologically motivated demodulation algorithm is presented by Elihilali, Chi, and Shamma [6]. The method first applies an auditory model to the input signal to yield an internal neural representation called the auditory spectrogram. The spectrogram is then analyzed to estimate the spectral and temporal modulations using filters modeled after ones found in the mammalian primary auditory cortex. This is the only method to date that views modulation as both a temporal and spectral process.

Turner and Sahani introduced a method for probabilistic demodulation based on a cascade process connecting demodulations at different time scales [18]. The demodulations extract increasingly fine time structures for sentences, phonemes, and formants. This paper is important in that it allows the definition of the modulator to be tuned, and also suggests the process can be solved as an optimization. The need for modulation criteria in order to differentiate between the infinite possible solutions in a demodulation is also highlighted. We extend this work by deriving optimization criteria that guarantee a global optimum solution.

Recent work proposed subband carrier demodulation as a method for coherent modulation detection [16]. The algorithm estimates the carrier phase of each channel of a complex filterbank. Each channel is then demodulated with a complex signal with an inversion of the phase estimate, essentially shifting the spectrum by the estimated carrier frequency. This is the first approach to attempt demodulation by first estimating the carrier. One of the primary advantages of this method is that the modulator and carrier are guaranteed to be bandlimited. Also, if the carrier phase estimate is accurate, the subband demodulation is perfect. However, if the phase estimate is inaccurate (due to noise, multiple carriers in the channel, or a nonstationary carrier), then the demodulation will improperly shift the spectrum. Even in the errant cases, though, the modulator can be guaranteed to be bandlimited.

These methods for demodulation decompose a signal with different levels of complexity and accuracy. This paper introduces a method that can return any possible decomposition depending on the objective function in the optimization, and that can improve performance on non-stationary and harmonic carriers.

B. Goals for A Demodulator

In addition to extracting a valid modulator and carrier from a signal, a demodulation algorithm should meet a few additional criteria. We believe that an acoustic demodulator should distinguish pitch from modulation consistently and based on a transparent and clearly understandable metric, it should act as an identity operator on modulators, and it should satisfy the projection property.

- **Distinguishing Pitch and Modulation:** Several demodulation algorithms are unable to explicitly define the characteristics that comprise a modulator or a carrier. The two components are determined on a case-by-case basis instead of operating under a higher level definition of the modulator or carrier class. We argue that an effective demodulation algorithm should explicitly define the characteristics of a modulator and a carrier and then obey those characteristics.

  Generally, we define a modulator as a lower frequency signal, and a carrier as a higher frequency signal. For the purposes of this paper, we will expand this definition to account for the perceptual experience. A human listener will interpret low-frequency modulation (below approximately 25 Hz) as amplitude variation, while higher frequency modulation is interpreted as multiple carrier frequencies. An acoustic modulator should therefore only in-

clude the low-frequency content that is interpreted as amplitude, but exclude the higher frequency content.

- **Modulator Identity:** A modulator is defined as a signal that cannot be further demodulated [5]. A demodulation algorithm should obey this property. In other words, demodulating a modulator should yield the trivial carrier, $c(t) = 1$.
- **Projection Property:** Ghitza developed a test for modulator distortion in modulation filtering [7], sometimes referred to as the projection property [5]. This test was later refined to the following two conditions.
  1) Remodulated subbands are bandlimited within the channel filter bounds.
  2) Demodulation of a remodulated carrier with a filtered modulator $(\tilde{s}(t) = \mathcal{H}\{m(t)\}c(t))$ yields the original carrier, $c(t)$.

Because we do not require a subband decomposition, we do not address the first condition. However, a demodulation algorithm should adhere to the second condition so that it can satisfy the requirements for modulation filtering.

We can achieve these goals by formulating demodulation as an optimization problem and designing appropriate cost functions and constraints. We show how to do this in the next section.

III. DEMODULATION AS AN OPTIMIZATION PROBLEM—THE THEORY

We describe general optimization problem as a cost (or objective) function minimization with any number of equality and inequality constraints [3]

$$
\begin{aligned}
\text{minimize} \quad & f(x) \\
\text{subject to} \quad & g_i(x) \le 0, \forall i \\
& h_i(x) = 0, \forall i.
\end{aligned}
$$

If the cost function $f(x)$ and inequality constraints $g_i(x)$ are all convex and the equality constraints $h_i(x)$ are all affine, then the problem is said to be a convex optimization problem. This is desired, because convex optimization problems can be iteratively solved using gradient descent algorithms, allowing them to be solved dependably and efficiently. For more on the definition of convexity, convex problems, and numerical solvers, see the textbook by Boyd and Vandenberghe [3].

However, properly constraining modulation within the convex restrictions is not trivial. The modulation relationship in (1) is not convex, because the known variable $s(t)$ is defined as the product of the two unknowns, $m(t)$ and $c(t)$. Therefore, a convex optimization problem cannot directly utilize (1) as an equality constraint. Instead, we must find a way to represent the relationship while preserving convexity.

In this paper, we write expressions for the modulation and carrier as a function of time. To implement these equations, we sample the signals above the Nyquist rate. We sum and norm the loss functions over all samples. We implement the continuous-time derivatives with finite differences. The minimizations

we describe here are high-dimensional, with one free parameter per equation per time sample. Fortunately, modern convex-optimization codes make it easy to solve for hundreds of free variables in fractions of a second on a laptop computer.

The next two sections show two different frameworks for converting demodulation into a convex optimization problem. Both approaches produce viable demodulations, but they differ in how they fit the demodulation relationship into a convex constraint. Sections III-A and III-B describe these constraints for linear- and logarithmic-domain demodulation and then we show how to satisfy modulator identity in each of these cases in Section III-C. Then, in Section IV, we will describe specific examples of cost functions and constraints for the two frameworks.

A. Logarithmic-Domain Convex Demodulation

One solution to the convexity issue is to work in the logarithmic domain. There, the optimization variables can be defined simply as the logarithm of the squared linear optimization variables $m(t)$ and $c(t)$, or

$$
\begin{aligned}
l_m(t) &= \log(m(t)^2) \\
l_c(t) &= \log(c(t)^2).
\end{aligned}
$$

The variables are squared to avoid taking the logarithm of negative numbers. The magnitude could also be used, but squaring is preferred because, unlike the magnitude, it preserves the bandlimiting on the modulator, and it simplifies working with the derivatives, which will be important in the cost functions introduced in Section IV.

Now, the non-convex relationship in (1) is represented with a simple sum of the logarithmic-domain optimization variables so that

$$
\log(s(t)^2) = l_m(t) + l_c(t).
$$

Using this constraint, framing a legal convex optimization problem is simple:

$$
\begin{aligned}
\text{minimize} \quad & C_m(l_m(t)) + C_c(l_c(t)) \\
\text{subject to} \quad & \log(s(t)^2) - l_m(t) - l_c(t) = 0, \forall t.
\end{aligned} \tag{2}
$$

Here, C_m and C_c are the cost functions for the modulator and carrier variables, respectively. These functions can be any convex function, including the l_1-, l_2-, and l_∞-norms. The cost functions dictate which of the infinite valid decompositions will be selected, and they are the aspect of the optimization framework that allows it to be customizable to a specific application.

Additional constraints and cost functions can also be added to the problem, but any additional cost functions or inequality constraints must be convex in the optimization variables, and any additional equality constraints must be affine. We will design specific cost function in Section IV-A, after describing the restrictions imposed in order to achieve modulator identity in Section III-C.

B. Linear-Domain Convex Demodulation

An optimization framework can also be derived in the linear domain. First, note that the constraint defined in (1) implies that, given the modulator $m(t)$, there is only one carrier $c(t)$

$$c(t) = m(t)^{-1}s(t). \qquad (3)$$

Using this relationship, an optimization problem can be posed that eliminates the optimization variable $c(t)$ and instead operates only on $m(t)$. First, the problem is framed with the non-convex constraint

$$\begin{aligned}\text{minimize} \quad & C_m(m(t)) + C_c(c(t))\\ \text{subject to} \quad & s(t) - m(t)c(t) = 0, \forall t.\end{aligned}$$

By substituting in (3), the constraint is reduced to the trivial $0 = 0$ and eliminated, and the problem is legally convex:

$$\text{minimize} \quad C_m(m(t)) + C_c(m(t)^{-1}s(t)).$$

The modulator cost function C_m can be any convex function. However, because the carrier cost function C_c is operating on the inverse of the optimization variable $m(t)$, it must be both convex and non-decreasing (non-negative first derivative) for the problem to be convex.

Another issue is that utilizing the inverse of the modulator $m(t)$ means that there is a discontinuity in the search space at $m(t) = 0$. Additionally, $m(t)^{-1}$ is convex only for $m(t) > 0$ (since its second derivative $m(t)^{-3}$ is positive for $m(t) > 0$) and concave for $m(t) < 0$. So, to eliminate these complications, the modulator is restricted to be non-negative.

$$\begin{aligned}\text{minimize} \quad & C_m(m(t)) + C_c(m(t)^{-1}s(t))\\ \text{subject to} \quad & -m(t) \le 0, \forall t.\end{aligned} \qquad (4)$$

We will design specific loss functions for the linear-domain demodulator in Section IV-B

C. Satisfying Modulator Identity

As described in Section II-B, modulator identity requires that demodulating a modulator leaves the trivial carrier, $c(t) = 1$. In the optimization context, this means that, with a modulator as input, the cost of the modulator and the trivial carrier is less than the cost of any other possible decomposition. For both optimization frameworks introduced above, we can derive a set of criteria that ensures modulator identity is satisfied in all cases.

1) Logarithmic Domain: If a signal is optimally decomposed in the logarithmic domain into the modulator–carrier pair $l_m(t)$ and $l_c(t)$, then the cost of these signals is necessarily less than the cost of any other decomposition so that

$$\begin{aligned}C_m(l_m(t)) &+ C_c(l_c(t))\\ &\le C_m(l_m(t) - x(t)) + C_c(l_c(t) + x(t)) \quad (5)\end{aligned}$$

for any feasible signal $x(t)$. The above equation must be true, because the optimization process minimizes the cost function,

and therefore the solution will have the minimal cost of all feasible modulator–carrier pairs.

Now let us say that we use the optimal logarithmic modulator $l_m(t)$ as the input signal for the algorithm. In order to have modulator identity, the cost of no further decomposition must be the minimal cost, since, as stated above, the modulator–carrier pair that minimizes the cost function is the solution of the optimization problem or

$$C_m(l_m(t)) + C_c(\mathbf{0}) \le C_m(l_m(t) - x(t)) + C_c(x(t)) \quad (6)$$

where, again, $x(t)$ is any feasible signal.

Equation (6) provides a condition for modulator identity for the optimization framework in (2). We will show that the specific logarithmic-domain cost functions to be derived in Section IV satisfy this condition.

2) Linear Domain: In the linear domain, if a signal $s(t)$ is found to have the optimal modulator $m(t)$, then the cost for that modulator is optimally minimal so that

$$\begin{aligned}C_m(m(t)) &+ C_c(m(t)^{-1}s(t))\\ &\le C_m(x(t)^{-1}m(t)) + C_c(m(t)^{-1}x(t)s(t)) \quad (7)\end{aligned}$$

for any feasible signal $x(t)$. If we use the optimal linear modulator $m(t)$ as the input signal, then the cost of no further decomposition should be the minimal cost, according to modulator identity:

$$C_m(m(t)) + C_c(\mathbf{1}) \le C_m(x(t)^{-1}m(t)) + C_c(x(t)). \quad (8)$$

Note that a unit carrier in the linear domain is $\mathbf{1}$, as opposed to $\mathbf{0}$ in the logarithmic domain.

Equation (8) provides a condition for satisfaction of modulator identity for the optimization framework in (4). We will show that the specific linear-domain cost functions derived in Section IV satisfy this condition.

IV. DEMODULATION AS AN OPTIMIZATION PROBLEM—THE DESIGN

In (2) and (4), we outlined an optimization framework for logarithmic-domain and linear-domain demodulation, respectively. We will now derive two sets of cost functions and constraints, one for each case, designed to perform the desired demodulation on an arbitrary input signal.

There is an art required in fitting a problem into a convex-optimization framework. Simple rules one might want to use to specify the optimal solution are often not convex. In some cases, a similar and convex rule exists to accomplish the same goal, but, in many cases, the requirements of convexity prevent using seemingly intuitive approaches. We have already seen this to a small degree, with the adjustments necessary to legalize (1) for inclusion in the optimization. Similar problems were found with designing specific cost functions and constraints, and the work that follows represents challenges that could be solved with creativity and compromise, and also challenges that led only to dead ends.

A. *Designing Logarithmic-Domain Cost Functions and Constraints*

In this section, we derive a cost function for the logarithmic-domain demodulation framework in (2). The cost function penalizes high frequencies in the modulator by minimizing the derivative of $l_m(t)$ and encourages spectral sparsity in the logarithmic carrier variable $l_c(t)$. Additionally, we minimize the carrier energy. We describe these two solutions for the loss function in the next two subsections.

1) Logarithmic-Domain Modulator Cost Function: The definition of an acoustic modulator suggested in Section II-B is given in terms of frequency bandlimiting. The most direct way to select a modulator based on this definition is to weight the spectrum appropriately. However, these spectral weights are in terms of the linear variable $m(t)$, and the logarithmic-domain demodulation framework from (2) optimizes the variable $l_m(t)$. The spectrum of $l_m(t)$ is related to the spectrum of $m(t)$, but has extra harmonics as a result of the nonlinear logarithmic transformation. The frequencies in the new spectrum are still related to the original fundamental frequency (as the periodicity of the signal is unchanged), but the sharp drops to large negative numbers as the original signal approaches 0 result in high-frequency harmonics that were not present in the original signal. As a result, bandlimiting constraints on the logarithmic variable do not cleanly translate to the linear domain.

However, we can impose a spectral weighting by working with the derivatives, which accentuate the higher frequencies of the signal by imposing a highpass filter. Therefore, if we minimize the derivative of the modulator, we will be essentially penalizing high frequencies. Higher-order derivatives penalize against high frequencies. Weighting the derivatives of the logarithmic variable $l_m(t)$ still does not directly translate to the linear modulator, but, unlike with spectral weights, we can derive equations that do translate, resulting in a modulator cost function that removes high frequencies from $m(t)$ by minimizing the derivatives of $l_m(t)$.

For the first derivative, this derivation is simple. The only difference between the first derivative of $l_m(t)$ and $m(t)$ is a scaling by the amplitude of the linear modulator $m(t)$:

$$\frac{d}{dt}l_m(t) = 2\frac{m'(t)}{m(t)}.$$

For the second derivative, the relationship is slightly more complex:

$$\frac{d^2}{dt^2}l_m(t) = 2\frac{m''(t)}{m(t)} - 2\frac{m'(t)^2}{m(t)^2}.$$

However, optimizing simultaneously over the second derivative and the square of the first derivative of $l_m(t)$ results in the optimization of $m''(t)$, though it is still scaled by the amplitude of $m(t)$:

$$\frac{1}{2}\left(\frac{d^2}{dt^2}l_m(t)\right) + \frac{1}{4}\left(\frac{d}{dt}l_m(t)\right)^2 = \frac{m''(t)}{m(t)} - \frac{m'(t)^2}{m(t)^2} + \frac{m'(t)^2}{m(t)^2}$$
$$= \frac{m''(t)}{m(t)}.$$

Higher order derivatives beyond the second derivative become far too complex for useful implementation. We found through experimentation that minimization of the second derivative is sufficient to impose the desired high-frequency penalty. Minimizing only the first derivative is generally not as effective.

So, we use the second derivative relationship in this implementation. Ideally, we simply minimize the l_2-norm

$$C_m(l_m(t)) = \left\| \frac{1}{2}\left(\frac{d^2}{dt^2}l_m(t)\right) + \frac{1}{4}\left(\frac{d}{dt}l_m(t)\right)^2 \right\|_2. \quad (9)$$

However, because of the rules of convex programming, we cannot do this. The squared first derivative is a convex function applied to the variable $l_m(t)$. Any function applied to that value must be convex as well, and also non-decreasing (non-negative slope). Norms do not satisfy the non-decreasing requirement.

Instead, we can minimize the summation of the second derivative. However, this will encourage large negative values that will drive down the total. To counter this, we also include a minimization on the norm of the second derivative of $l_m(t)$. Thus, the loss C_m for the proposed modulator l_m is the sum of the second derivative approximation in (9) and a term that prevents the second derivative from getting too large, or

$$C_m(l_m(t)) = \sum \left[\frac{1}{2}\left(\frac{d^2}{dt^2}l_m(t)\right) + \frac{1}{4}\left(\frac{d}{dt}l_m(t)\right)^2 \right]$$
$$+ \left\| \frac{d^2}{dt^2}l_m(t) \right\|_2. \quad (10)$$

2) Logarithmic-Domain Carrier Cost Function: For the cost function on the carrier variable $l_c(t)$, we use the l_1-norm of the spectrum of $l_c(t)$ to encourage sparsity in the frequency domain. This is used on the assumption that the desired carrier is a harmonic signal, and therefore has a sparse spectrum. Note, the harmonics added in the nonlinear logarithmic transformation do affect this metric by crowding the spectrum a bit more, but a carrier with a sparse spectrum will still have a sparse spectrum in the logarithmic domain:

$$C_c(l_c(t)) = \|\mathcal{F}\{l_c(t)\}\|_1.$$

By including the l_2-norm of $l_c(t)$ weighted by a constant W_E, the method produces better results by allowing more energy to shift to the modulator, giving us

$$C_c(l_c(t)) = \|\mathcal{F}\{l_c(t)\}\|_1 + W_E \|l_c(t)\|_2.$$

Also, we found that using the analytic signal $\tilde{s}(t)$ and then optimizing over the real part of $l_c(t)$, similar to the process by Atlas *et al.* [2], produced better results than the real signal $s(t)$. Essentially, this means that we are demodulating the Hilbert envelope with the optimal criteria. Thus the complete loss function for the carrier is written as

$$C_c(l_c(t)) = \|\mathcal{F}\{\mathcal{R}\{l_c(t)\}\}\|_1 + W_E\|\mathcal{R}\{l_c(t)\}\|_2. \quad (11)$$

Note that the use of norms here falls within the constraints of convex optimization because, unlike the modulator cost function in (9) and (10), the norms are performed on linear operations of the optimization variables.

3) Logarithmic-Domain Optimization Problem: So, substituting (10) and (11) into (2), we are able to define the convex problem as a minimization of

$$W_{l_m} \sum \left[\left(\frac{1}{2} \left(\frac{d^2}{dt^2} l_m(t) \right) + \frac{1}{4} \left(\frac{d}{dt} l_m(t) \right)^2 \right) \right.$$
$$\left. + \left\| \frac{d^2}{dt^2} l_m(t) \right\|_2 \right)$$
$$+ \| \mathcal{F}\{\mathcal{R}\{l_c(t)\}\} \|_1 + W_E \| \mathcal{R}\{l_c(t)\} \|_2 \qquad (12)$$

subject to

$$\log \left(s(t)^2 \right) - l_m(t) - l_c(t) = 0, \forall t. \qquad (13)$$

Note that we included an additional parameter weight W_{l_m} in the modulator cost function. This is set to $W_{l_m} = 1$ for signals with a single-component carrier, but must be increased to improve results for harmonic carriers, which will be shown later. Based on experiments, we set $W_E = 1/10$ for all cases.

4) Sign Restoration: This optimization problem solves for $l_m(t)$, from which the magnitude of the linear modulator can be extracted:

$$|m(t)| = \sqrt{e^{l_m(t)}}. $$

However, this results in a similar problem to that discussed in Section II-A for Hilbert envelope detector. The modulator magnitude is non-negative and therefore will not be able to track sign changes in the original modulator. These errors occur at zero crossings, where the magnitude stays positive while the original modulator is negative, as discontinuities in the first derivative of the modulator magnitude. So, in order to restore the original modulator $m(t)$ from the magnitude $|m(t)|$, we iteratively solve for the combination of the positive and negative magnitude that minimizes jumps in the first derivative, or the magnitude of the second derivative. To restore the sign we

$$\begin{array}{ll} \text{minimize} & \left| \frac{d^2}{dt^2} \tilde{m}(t) \right| \\ \text{subject to} & \tilde{m}(t) \in \{|m(t)|, -|m(t)|\}. \end{array}$$

5) Modulator Identity: We can show that this algorithm acts as an identity operator for the modulator by proving that (6) is true. It is not difficult to show that one set of criteria for this is

$$C_c(\mathbf{0}) \leq 0 \quad \text{and}$$
$$C_c(l_c(t) + x(t)) \leq C_c(l_c(t)) + C_c(x(t)). \qquad (14)$$

The first inequality must be true, because the cost function is simply a sum of norms, and any norm of $\mathbf{0}$ is 0. The second inequality can also be proven to be true. First, we begin with Jensen's inequality, which is true for any convex function (C_c has already been restricted to be convex). Here, we also use the trivial substitutions $l_c(t) = \alpha(1/\alpha l_c(t))$ and $x(t) = (1 - \alpha)(1/1 - \alpha x(t))$ and rewrite (14) to find

$$C_c(l_c(t) + x(t)) \leq \alpha C_c \left(\frac{1}{\alpha} l_c(t) \right) + (1 - \alpha)$$
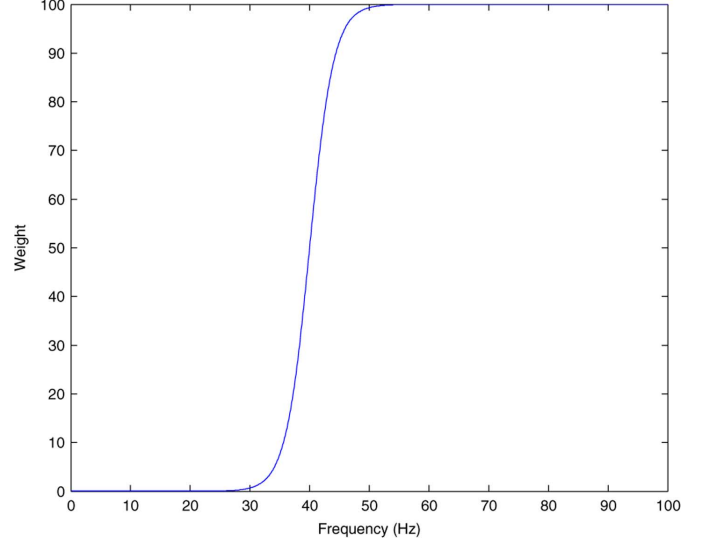$$C_c \left(\frac{1}{1\alpha} x(t) \right), \quad \alpha \in [0, 1]. \quad (15)$$

Fig. 1. Spectral weight function $W_m(f)$ used to penalize high frequencies in the modulator $m(t)$. Frequencies in the modulator below 25 Hz cost very little, and the cost increases until it flattens near 55 Hz and above. A symmetric weighting is applied to the negative frequencies.

If C_c also satisfies the scaling property

$$\gamma C_c(x(t)) \leq C_c(\gamma x(t)), \quad \gamma \in [0, 1] \qquad (16)$$

then we can show that the desired inequality in (14) must be true, because

$$C_c(l_c + x(t)) \leq \alpha C_c \left(\frac{1}{\alpha} l_c(t) \right)$$
$$+ (1 - \alpha) C_c \left(\frac{1}{1 - \alpha} x(t) \right)$$
$$\leq C_c(l_c(t)) + C_c(x(t)). \qquad (17)$$

For this implementation, C_c does fulfill (16) because C_c contains only norms, and norms by definition satisfy the scaling property with equality. Therefore, this set of cost functions satisfies modulator identity.

6) Projection Property: Once again, the logarithmic transformation complicates the process. Assessing the affects of filtering in the linear domain on a logarithmic signal while preserving filter generality is extremely difficult. So, we will instead leave proof of the projection property for logarithmic domain for future work.

B. Designing Linear-Domain Cost Functions and Constraints

1) Linear-Domain Modulator Cost Function: Unlike the logarithmic-domain case, in the linear-domain framework of (4), the cost functions can be written directly in terms of the spectrum of the modulator $m(t)$. So, we use high weights in the frequency domain to penalize high frequencies. In our work, the weighting function is a sigmoid function, displayed in Fig. 1. The weights are designed so that the bandlimiting is smooth, and so that, at some point between 25 and 50 Hz, the cost of including the spectral content grows too high, imposing the desired cutoff point from Section II-B.

The cost function used for this modulator is then the l_2-norm of the weighted spectrum (weighted at frequency f with weight $W_m(f)$) plus the l_2-norm of the modulator itself or

$$C_m(m(t)) = \|W_m(f)\mathcal{F}\{m(t)\}\|_2 + \|m(t)\|_2. \tag{18}$$

2) Linear-Domain Carrier Cost Function: Because the carrier is unique for any given modulator $m(t)$, finding the desired modulator means that the desired carrier has been found as well. The modulator cost function is sufficient for determining the desired modulator, so no cost on the carrier is needed for the optimization. We ignore the carrier by setting its cost function to zero or

$$C_c\left(m(t)^{-1}s(t)\right) = 0. \tag{19}$$

3) Linear-Domain Constraints: Because the carrier is left unpenalized, it is necessary to add another constraint. We observe that, without loss of generality, the magnitudes of the signal, carrier, and therefore modulator are all less than 1 at all times (values greater than 1 are clipped) so

$$|s(t)| \le 1, \quad |c(t)| \le 1, \quad |m(t)| \le 1$$

Also, using (3) and the above inequalities, it must be true that the magnitude of the modulator is greater than or equal to the magnitude of the signal at all times so

$$|c(t)| = \left|m(t)^{-1}s(t)\right| = \left|m(t)^{-1}\right|\left|s(t)\right| \le 1$$

therefore

$$|s(t)| \le |m(t)|.$$

Since we have already constrained the modulator to be non-negative, we know that the modulator and its magnitude are equal at all times, and the relationship can be simplified to

$$|s(t)| \le m(t). \tag{20}$$

4) Linear-Domain Optimization Problem: With the constraint in (20) and the cost functions in (18) and (19), we have a complete optimization problem within the linear-domain framework:

$$
\begin{aligned}
\text{minimize} \quad & \|W(f)\mathcal{F}\{m(t)\}\|_2^2 + \|m(t)\|_2^2 \\
\text{subject to} \quad & m(t) - 1 \le 0, \forall t \\
& |s(t)| - m(t) \le 0, \forall t.
\end{aligned}
$$

Note that the norms in the cost function are now squared, so that the problem becomes a quadratic optimization problem (a subclass of convex optimization) and thus is easier to solve.

5) Sign Restoration: Because this algorithm only considers non-negative modulators it is preferable to include a means for sign restoration. However, because of the bandlimiting of the modulator, sign errors will not manifest themselves as discontinuities in the first derivative, but instead will be smooth, making it very difficult to restore the sign changes from the original modulator. Squaring the signal prior to demodulation solves this problem mathematically but, in practice, still does not produce sharp enough discontinuities to allow the method

in Section IV-A4 to restore the sign changes. Development of an extension to this algorithm that would allow it to track sign changes in the modulator is a direction for future work.

6) Modulator Identity: Fulfilling modulator identity in the linear domain means that (8) must be true. It is easy to show that this is the case if

$$C_c(\mathbf{1}) \le C_c(m(t)^{-1}s(t))$$

and

$$C_c(m(t)^{-1}x(t)s(t)) \le C_c(x(t)).$$

This is trivially true, because $C_c(c(t)) = 0$ for all $c(t)$. Therefore, modulator identity is satisfied.

7) Projection Property: We can show that the projection property holds for these cost functions. First, we note that the carrier $c(t)$ has already been optimally demodulated. Therefore, any components in the original signal $s(t)$ that are sufficiently low cost in the modulator cost function C_m have already been removed as modulation.

The unfiltered optimal modulator $m(t)$, on the other hand, consists exclusively of components that are of sufficiently low cost in C_m (defined in Fig. 1 as low frequencies). Furthermore, if we normalize the filter $\mathcal{H}$ in the frequency domain

$$H(\omega) \le 1, \forall \omega$$

then it must be true that the filtered modulator $\mathcal{H}\{m(t)\}$ also consists of only low-cost frequency components. Here, it is critical that the modulator cost function C_m is defined exclusively in terms of frequency weights and energy norms, which makes it possible to directly define the effect of modulation filtering on the cost in the optimization.

Based on these conclusions, in the remodulated signal $\tilde{s}(t) = \mathcal{H}\{m(t)\}c(t)$, $\mathcal{H}\{m(t)\}$ contributes only modulator content (provided that $\mathcal{H}\{m(t)\}$ satisfies all other modulator constraints, which, in this case, means non-negativity), and $c(t)$ contributes only carrier content, as defined by the cost functions C_m and C_c. As a result, the demodulated components will be the filtered modulator $\mathcal{H}\{m(t)\}$ and the original carrier $c(t)$, satisfying the projection property.

V. RESULTS

In this section, we demonstrate the utility of the optimal demodulation approaches on several different kinds of signals. The simplest approaches to demodulation, such as an envelope detector, do a good job of finding the envelope for a single modulated sinusoid with no noise. To better illustrate the difference between these algorithms we demonstrate optimal demodulation using frequency chirps, harmonic carriers and a noisy carrier.

We often display the demodulation results using signal plots, but since the results are so good it is hard to judge their performance. Thus, we also characterize the performance of an algorithm by adding noise to stress the demodulation approach. Ideally, the performance should degrade gradually as noise is added, until the model is no longer valid and the modulation and carrier are no longer recoverable. Noise is a difficult component for demodulation algorithms to deal with because noise is not part of the model $s(t) = m(t)c(t)$.
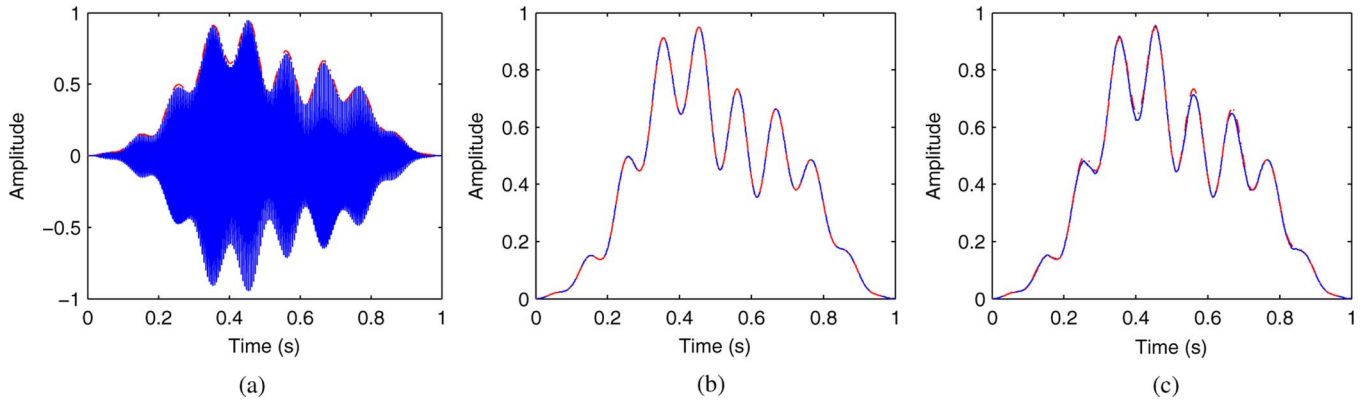
Fig. 2. Low frequency modulator applied to a single sinusoidal carrier of 200.7 Hz (a) and the modulators extracted by logarithmic-domain demodulation (b) and linear-domain demodulation (c) plotted with the original modulator.
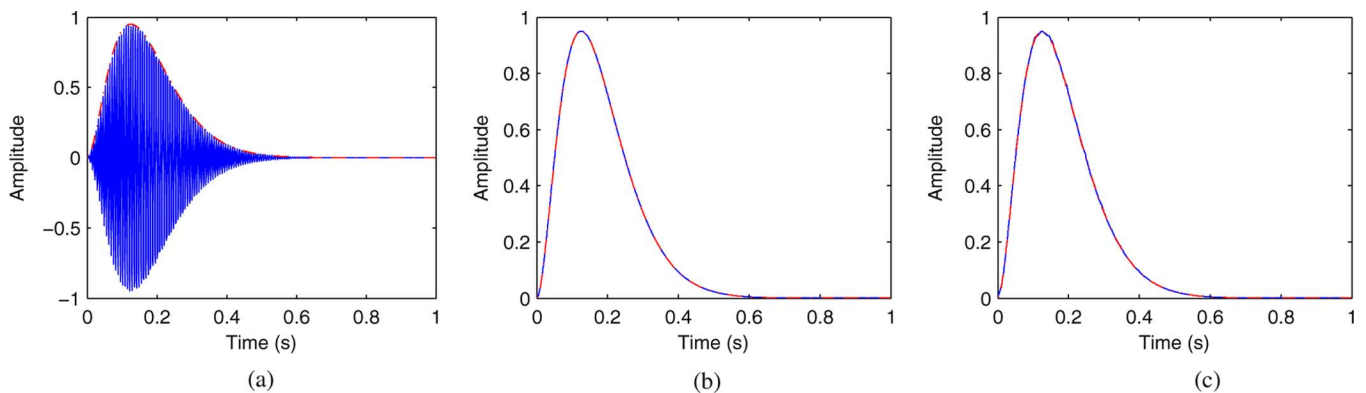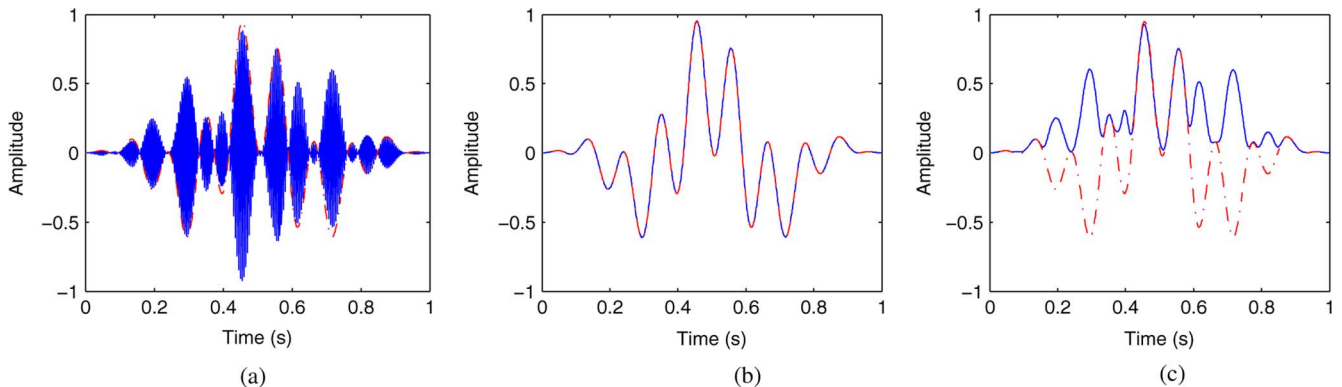

Fig. 3. Exponentially derived modulator applied a single sinusoidal carrier of 150.7 Hz (a) and the modulators extracted by logarithmic-domain demodulation (b) and linear-domain demodulation (c) plotted with the original modulator.

Fig. 4. Phase-inverting sinusoidal modulator applied to a single sinusoidal carrier of 200.7 Hz (a) and the modulators extracted by logarithmic-domain demodulation (b) and linear-domain demodulation (c). The original modulator is shown as a dot-dashed line. Linear-domain demodulation fails to track the sign changes because it restricts the modulator to be non-negative.

For all examples, input signals were broken into 100-ms clips with a 50% overlap and windowed with a Hann window. This allows computation time to scale linearly with increasing signal length instead of exponentially at the cost of limited time scope for the optimization. However, we have found no significant difference in the demodulation as a result of this decision. Optimizations were performed with the Matlab toolbox cvx [8].

All synthetic examples are also windowed with a Hann window of length equal to the entire signal, which is reflected in the plots. This windowing is to smooth the onsets and offsets. For real speech examples, though, this window was not needed.

A. Stationary Sinusoidal Carrier

1) Sinusoidal Modulator: The simplest example is a high-frequency carrier with a low-frequency modulator. An example is shown in Fig. 2(a), with a 200.7-Hz carrier. The extracted modulators are in Fig. 2(a) and (b) for logarithmic-domain demodulation and linear-domain demodulation, respectively. Both methods estimate the modulator very accurately.
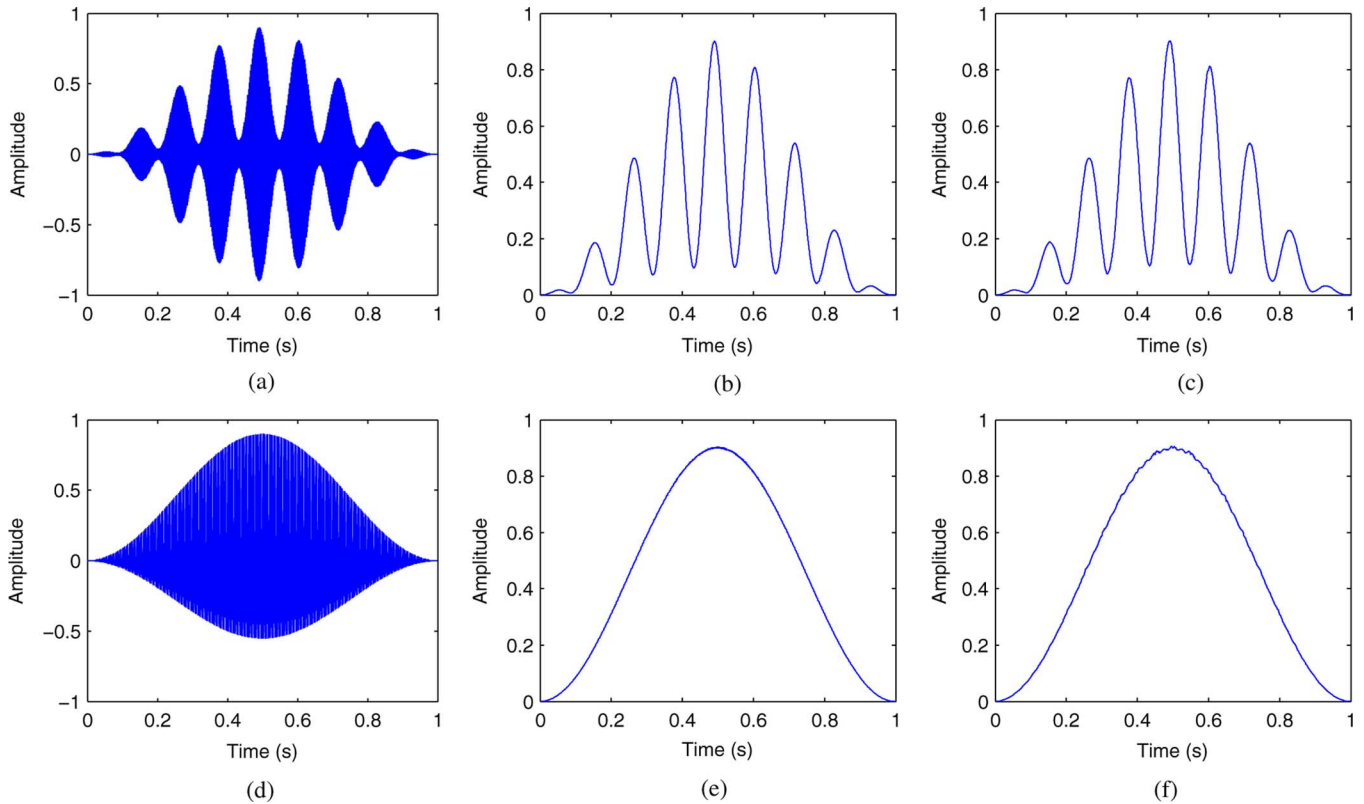
Fig. 5. A 500-Hz sinusoid with an 8.7-Hz modulation and its extracted modulators are shown in the top row. The same carrier with a 250-Hz modulation is shown with its extracted modulators in the bottom row. Both demodulators extract the 8.7-Hz modulation and the signal's overall window shape, but then ignore the 250-Hz pitch modulation. (a) Original, 8.7-Hz Modulation. (b) Logarithmic. (c) Linear. (d) Original, 250-Hz Modulation. (e) Logarithmic. (f) Linear.

2) Exponential Modulator: An exponentially derived modulator applied to a single sinusoidal carrier of 150.7 Hz is shown in Fig. 3(a) with its extracted modulators in Fig. 3(b) and (c). Again, both methods are very accurate.

3) Suppressed-Carrier Sinusoidal Modulator: Fig. 4(a) shows a single sinusoidal carrer of 200.7 Hz with a suppressed-carrier sinusoidal modulator. Logarithmic-domain demodulation estimates the modulator very well, as shown in Fig. 4(b). Linear-domain demodulation fails to track the sign changes [Fig. 4(c)], because the modulator is restricted to non-negativity. However, the estimated modulator does track the magnitude of the original modulator well.

4) Low-Frequency Versus High-Frequency Modulator: We suggested in Section II-B that an effective demodulator should differentiate between low-frequency modulation and high-frequency modulation in a manner similar to a human listener. To demonstrate this, Fig. 5(a) shows a single sinusoidal carrier of 500 Hz with a low-frequency modulator of 8.7 Hz, and Fig. 5(d) shows the same single sinusoidal carrier with a high-frequency modulator of 250 Hz. The only difference between these two signals is the modulation frequency.

The low-frequency estimated modulators are shown in Fig. 5(b) and (c) for the logarithmic-domain modulator and linear-domain modulator, respectively. The high-frequency estimated modulators are shown in Fig. 5(e) and (f).

Both demodulation methods successfully differentiate the two cases. The low-frequency modulator is extracted as modulation, while the high-frequency modulation is seen as pitch,
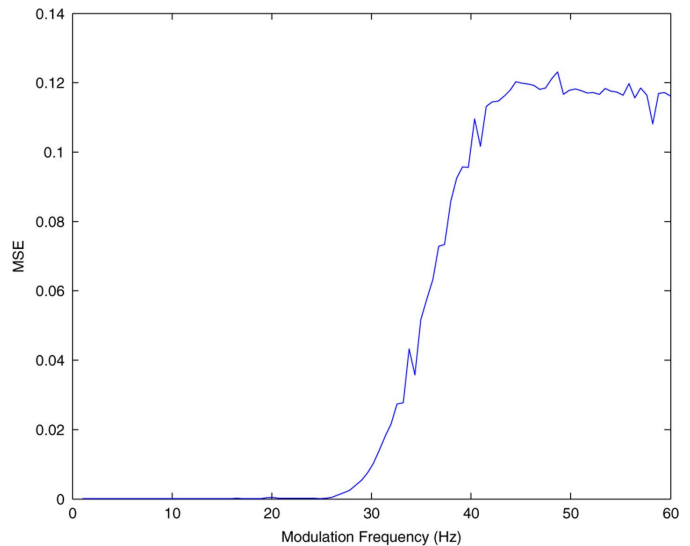

Fig. 6. Mean squared error of the estimated modulator plotted against the frequency of the original modulator for a single carrier at 353.8 Hz. At low frequencies, the error is small. At high frequencies, the error is large, because the estimated modulator excludes high frequencies. This plot is for linear-domain demodulation. Note the similarity in shape to the spectral weights $W_m(f)$ in Fig. 1.

and so the estimated modulators are only the window that was applied to the original signal.

Logarithmic-domain demodulation makes this distinction because, as the frequency of the modulator increases, the penalty
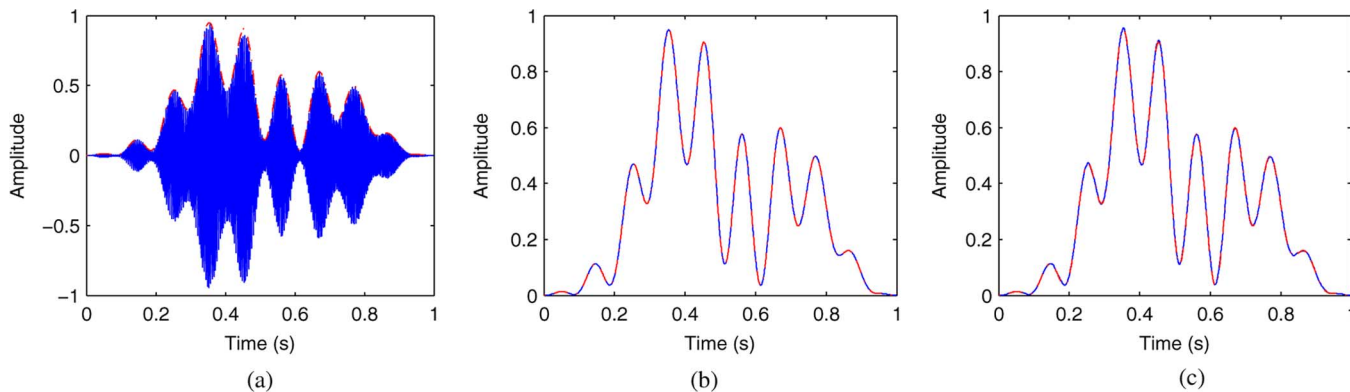
Fig. 7. Low-frequency modulator applied to a chirp carrier (a) and the modulators extracted by logarithmic-domain demodulation (b) and linear-domain demodulation (c). In both cases the errors are negligible.

for its high second derivative increases as well. Eventually, the cost becomes too high, and so the high-frequency content is excluded from the estimated modulator.

Linear-domain demodulation penalizes higher frequencies in the modulator more directly, using the frequency-dependant penalty function shown in Fig. 1. As the frequency increases, so does the penalty, and, as was the case with logarithmic-domain demodulation, the cost eventually becomes too high and the higher frequency content is excluded from the modulator.

To illustrate this transition from amplitude to pitch as a function of modulation frequency more clearly, Fig. 6 shows the mean squared error of the estimated modulator for different modulation frequencies, using linear-domain demodulation. The plot shows that modulation under 25 Hz is extracted accurately. Likewise, it is clear in Fig. 1 that the cost for the modulator frequency at these levels is very low. As the frequency increases above 25 Hz, the cost on the modulator increases as well, and so the demodulation algorithm begins to push the energy out of the estimated modulator. This is reflected as increased error in Fig. 6, since the estimated modulator is deviating from the original modulator (counterintuitively, in this case, error is desired as modulation frequency increases). Finally, above 45 Hz, the cost of the modulation frequency has become sufficiently high, and the error plateaus, indicating that all of the higher frequency modulation energy has been pushed from the estimated modulator. At this point, estimated modulators are similar to Fig. 5(f).

B. Swept Sinusoidal Carrier

Fig. 7(a) shows a sweeping sinusoidal carrier (from 200 to 250 Hz in 1 second) modulated by a low-frequency modulator. The estimated modulators are shown in Fig. 7(b) for logarithmic-domain demodulation and Fig. 7(c) for linear-domain demodulation. Both methods are able to accurately estimate the original modulator.

C. Stochastic Carrier

All previous examples have used single sinusoid carriers (except in the high-frequency modulation cases, where the carrier can be seen as multiple sinusoids with a unit modulator). Fig. 8(a) shows a low-frequency modulator instead applied to a stochastic carrier of uniformly distributed noise. Fig. 8(b)
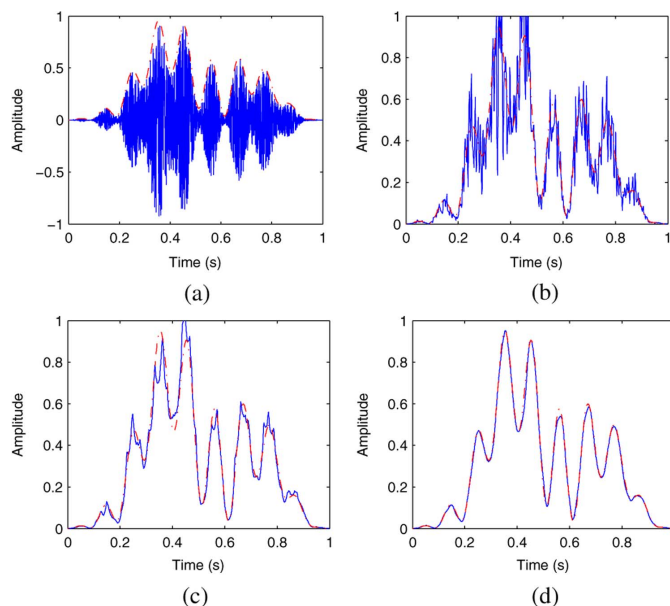

Fig. 8. Low-frequency modulator applied to a stochastic carrier (a) and the modulators extracted by logarithmic-domain demodulation with $W_{l_m} = 1$ (b) and $W_{l_m} = 10$ (c), and linear-domain demodulation (d). All are plotted with the original modulator. Additional weight is needed for the modulator in the logarithmic-domain case because the added spectral content in the harmonic carrier disrupts the balance of the cost function. Setting $W_{l_m} = 10$ restores the balance.

shows the modulator estimated with logarithmic-domain demodulation. The estimate is not very good, because the carrier is not spectrally sparse. The estimate can be greatly improved if extra weight is added to the modulator cost function by setting $W_{l_m} = 10$, resulting in the modulator shown in Fig. 8(c). This is a much better estimate, though it is still not as accurate as previous examples. Linear-domain demodulation is still able to estimate the modulator very accurately, as shown in Fig. 8(d).

Linear-domain demodulation outperforms logarithmic-domain demodulation in this case because, in the linear case, the dense carrier is not penalized in any way. So, increasing the complexity of the carrier, either with added harmonics or noise, will have a very minimal effect. Logarithmic-domain demodulation, on the other hand, penalizes non-sparsity in the carrier spectrum, so added harmonics or noise can throw off the
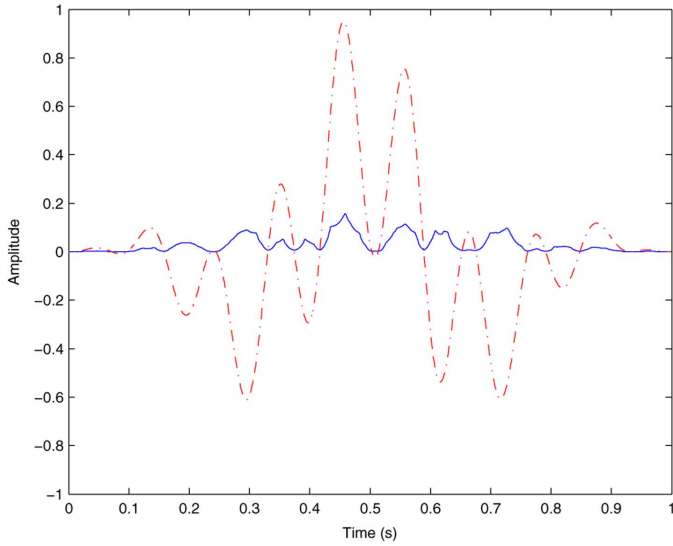
Fig. 9. Logarithmic-domain demodulation result for the same modulated single sinusoidal carrier from Fig. 4 with $W_{l_m} = 10$, which means the balance will favor a modulator with less cost on its slope. This shift results in an overly jagged modulator. In this case, the weight $W_{l_m} = 10$ is excessive for the simple carrier.
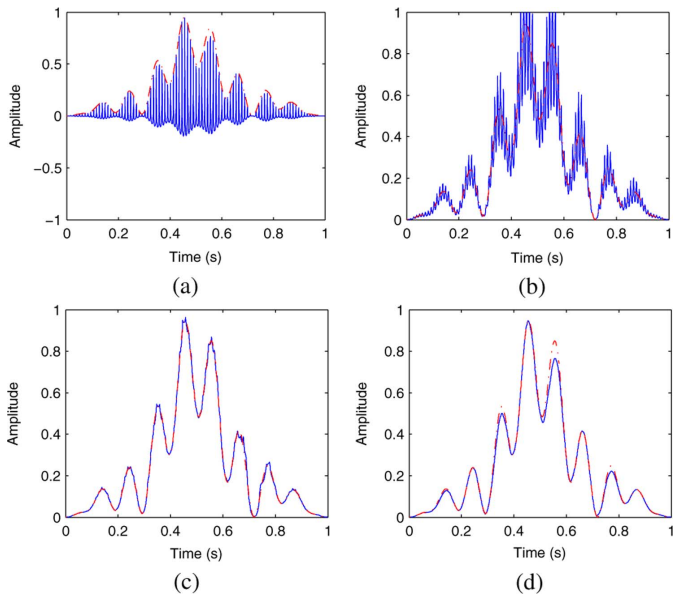

Fig. 10. Low-frequency modulator applied to a harmonic carrier of 90.7 Hz and its first ten harmonics (a) and the modulators extracted by logarithmic-domain demodulation with $W_{l_m} = 1$ (b) and $W_{l_m} = 10$ (c), and linear-domain demodulation (d). All show the original modulator as a dot-dashed line.

Fig. 11. Mean squared error of the modulator estimated with Hilbert envelope detection (solid line), linear-domain demodulation (dotted line), and logarithmic-domain demodulation (dashed line) for a single 200.7-Hz sinusoid modulated with a 4.3-Hz sinusoid with added noise. All three methods are robust to noise and then deteriorate similarly below 20-dB SNR. However, logarithmic-domain demodulation does show some inconsistent deterioration as a result of the higher level noise.

Fig. 12. Mean squared error of the modulator estimated with Hilbert envelope detection (solid line), linear-domain demodulation (dotted line), and logarithmic-domain demodulation (dashed line) for a 90.7-Hz harmonic tone modulated with a low-frequency sinusoid with added noise. Unlike Fig. 11, the inability of Hilbert envelope detection to extract modulators for harmonic signals causes a consistently high error for that method. Both optimization methods still perform similarly well compared to the single sinusoidal carrier case.

balance between the cost of the modulator and the cost of the carrier, pushing carrier energy into the modulator.

At this point, it is sensible to wonder why the weight W_{l_m} is not set to 10 in all cases. The problem with this change is that simpler carriers, such as single sinusoidal carriers, are not strongly penalized. So, increasing the cost on the modulator leads to a jagged and overly minimized modulator. For example, in the case of the single sinusoidal carrier shown in Fig. 4, the estimated modulator from logarithmic-domain demodulation with $W_{l_m} = 10$ is shown in Fig. 9. This result is clearly suboptimal.
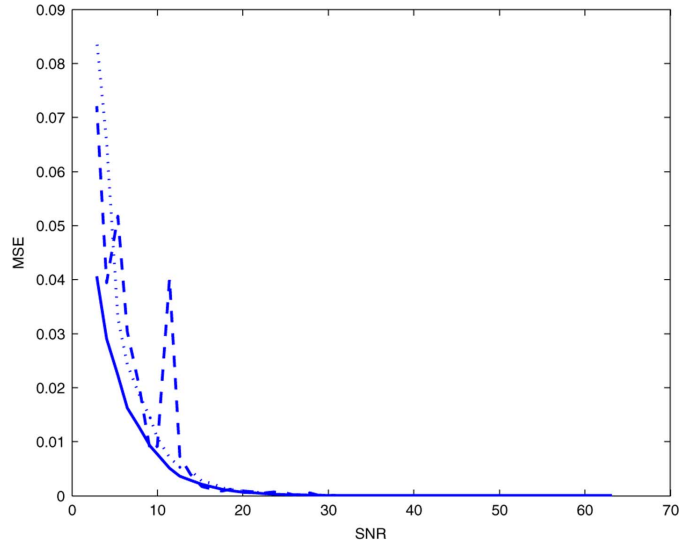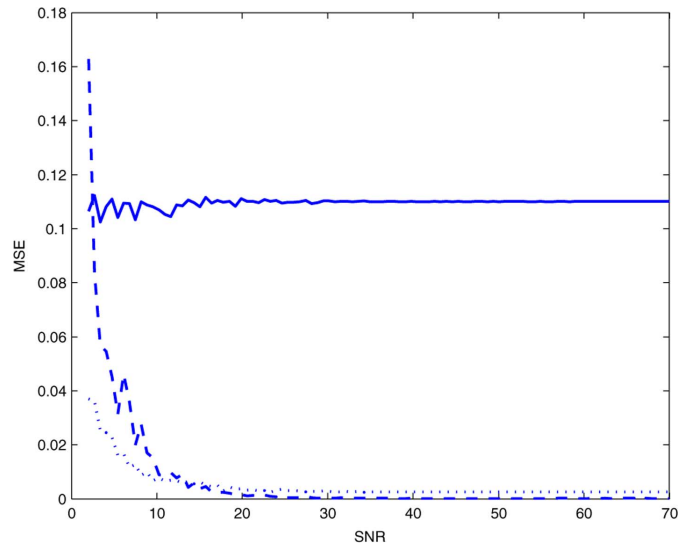
So, tuning the weight W_{l_m} is required for optimal results with carriers of varying cost in the optimization functions.

D. Harmonic Carrier

Fig. 10(a) shows a harmonic carrier (90.7 Hz fundamental with ten harmonics) modulated with a low-frequency modulator. Fig. 10(b) shows the modulator estimated with logarithmic-domain demodulation. As was the case with a stochastic carrier, this estimate is not very good. Here, though
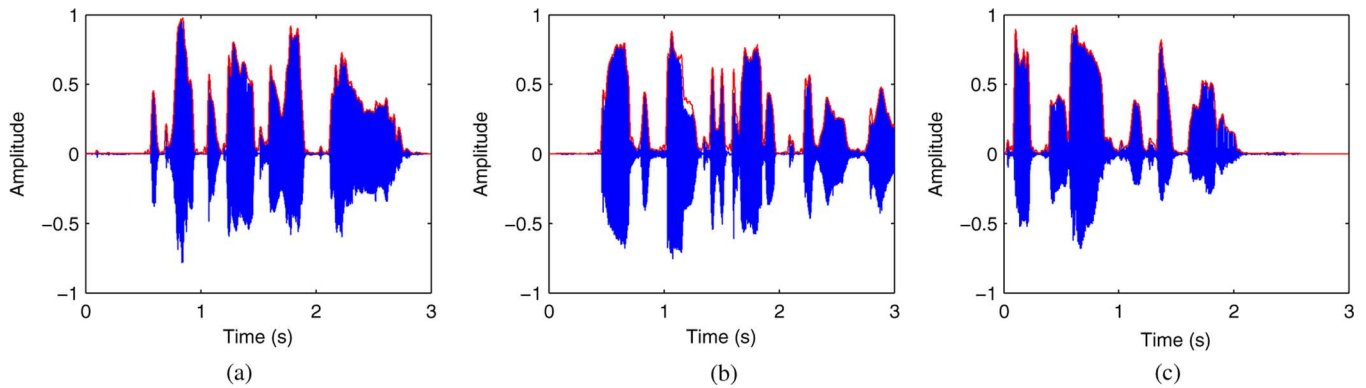
Fig. 13. Several speech signals and the modulators estimated with linear domain demodulation. (a) Female speaker, "The pipe began to rust while new." (b) Female speaker, "Add the sum to the product of these three." (c) Male speaker, "Cats and dogs each hate the other."

the carrier is still sparse, the added harmonics increase exponentially in the logarithmic domain, and so the balance of the optimization is skewed. Once again, the estimate is improved if extra weight is added to the modulator cost function by setting $W_m = 10$, resulting in the modulator shown in Fig. 10(c). Also like the stochastic carrier example, linear-domain demodulation is still able to estimate the modulator very accurately, as shown in Fig. 10(d).

E. Noisy Signals

Linear-domain and logarithmic-domain demodulation were tested in the presence of noise and compared to Hilbert envelope detection. The mean squared error (MSE) of the estimated signal for different signal-to-noise ratios (SNRs) is shown in Figs. 11 and 12, with Hilbert envelope detection shown as a solid line, linear-domain demodulation as a dotted line, and logarithmic-domain demodulation as a dashed line.

In the single sinusoidal carrier case (Fig. 11), all of the methods show very accurate demodulation above 20-dB SNR, resulting in very low mean squared error of the estimated modulator. Deterioration is seen below 20 dB, though it is smooth and comparable for Hilbert envelope detection and linear-domain demodulation. Logarithmic-domain demodulation shows a slightly more erratic behavior with increasing noise, though the overall level of deterioration is similar to the other methods.

In the case of the harmonic carrier (Fig. 12), both optimization algorithms perform similarly to the single sinusoidal carrier case, and degrade similarly as well. However, Hilbert envelope detection fails at all SNR levels, because it is unable to properly demodulate harmonic signals.

It is worth noting that, unlike Hilbert envelope detection and logarithmic-domain demodulation, the reason for the degradation of the linear-domain algorithm is not that it fails as a result of the noise. Instead, it is because of the constraint $|s(t)| \leq m(t)$, in which the added noise is included in the signal $s(t)$. So, the algorithm is restricted to demodulating the sum of the original modulated signal and the noise. As the noise becomes more significant, its effect on the waveform $s(t)$ increases as well, and so the signal being demodulated, $s(t)$, resembles the original modulated signal $m(t)c(t)$ less and less.

F. Speech Examples

We also applied linear domain demodulation to several real speech signals. We chose to only utilize linear-domain demodulation in this case because of the harmonic and noisy nature of speech, which logarithmic-domain demodulation struggles to demodulate as accurately, as shown in the Sections V-C and V-D. Also, because linear-domain demodulation is a quadratic optimization problem, it can be solved much faster than logarithmic-domain demodulation, which began to get cumbersome with the real audio signals.

Fig. 13 shows several speech signals and their estimated modulators. The modulators track the signals very well, and look to be very reasonable and accurate estimations. Since there is not a known true modulator, no numerical comparison can be made. However, these examples do confirm that this method is capable of demodulating real audio consisting of harmonic and noisy components.

G. Subband Decomposition

Most demodulation methods use a front-end subband decomposition, either to separate source components or to simulate critical bands in the auditory system. Our methods can also include a subband decomposition prior to demodulation, as this section will show. The extension from (1) is simple, where instead of being viewed as a single product model, the problem is viewed as a sum-of-products model [14]:

$$s(t) = \sum_i s_i(t) = \sum_i m_i(t)c_i(t).$$

In this model, $s_i(t)$ is a subband channel. If each channel demodulation is viewed as its own problem $s_i(t) = m_i(t)c_i(t)$, then each can be solved individually with convex demodulation.

However, there are several issues that the inclusion of a decomposition raises. Namely, demodulation after a subband decomposition eliminates the ability to extract wideband carriers and modulators, compromises the symmetry of the spectrum potentially resulting in the need for complex modulators, and changes the very definition of modulation to include spectral filtering effects.
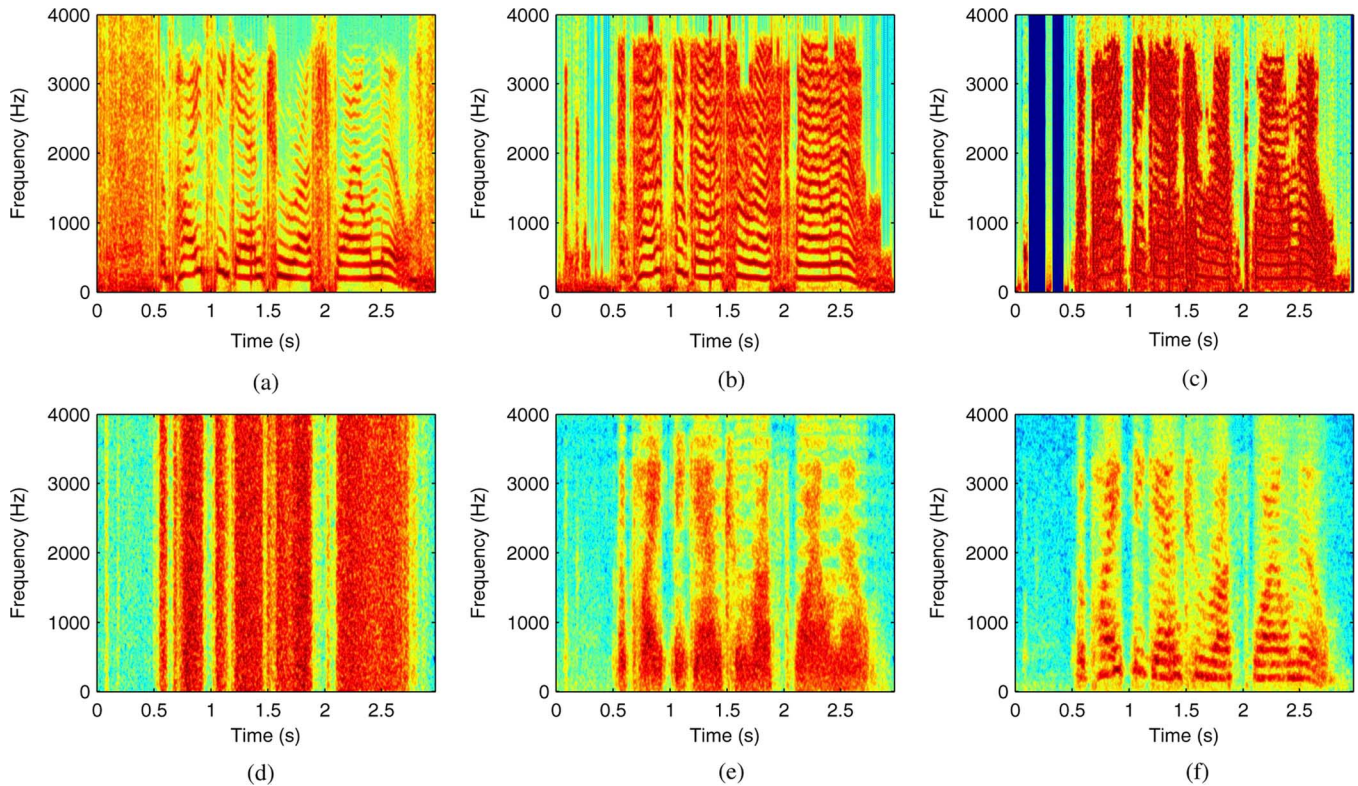
Fig. 14. Several plots of the linear-domain demodulated components of the signal "The pipe began to rust while new" as spoken by a female speaker with and without subband decomposition. The first row shows the carriers after (a) no decomposition, (b) 11-channel, 200-Hz Hamming filterbank decomposition. (c) 101-channel, 20-Hz Hamming filterbank decomposition. The second row shows white noise modulated with the extracted envelopes after (d) no decomposition, (e) 11-channel decomposition, (f) 101-channel decomposition. Higher decomposition channel resolution results in spectral filtering and pitch information moving from the carrier to the modulator.

If a speech signal is demodulated in its entirety, without any decomposition, then the filtering effects are preserved in the carrier [as shown in Fig. 14(a), demodulated with the linear-domain method], while the modulator [shown with a white noise carrier in Fig. 14(d)] represents only the changes in amplitude. Because the demodulation is performed on the wideband signal, it will not affect the relative harmonic amplitudes. However, in the case where decomposition is included, the spectral effects (both from filtering present in the original signal and from the subband filtering itself) are removed from the carrier in the demodulation. Fig. 14(b) shows the rebuilt carrier (the sum of the channel carriers) after demodulation on each channel of an 11 channel, 200 Hz Hamming filterbank decomposition. The relative harmonic amplitudes are diminished. In the corresponding modulator [applied to subband filtered white noise and shown in Fig. 14(e)], it is clear that the speech filtering data and formants are considered modulation rather than filtering, even with only 11 subband channels. With a 101 channel, 20-Hz Hamming filterbank decomposition [Fig. 14(c)], the filtering effects are even further removed in demodulation. In the 101 subband modulator [Fig. 14(f)], because the channel resolution is high enough to fit between the harmonics, even pitch data is present, in addition to speech filtering data. These figures demonstrate that, in the case of demodulating after decomposition, the filtering information is represented in the array of modulators instead of remaining in the carrier. It is unclear which of these approaches, no subbands or a higher number of subbands, is preferable.
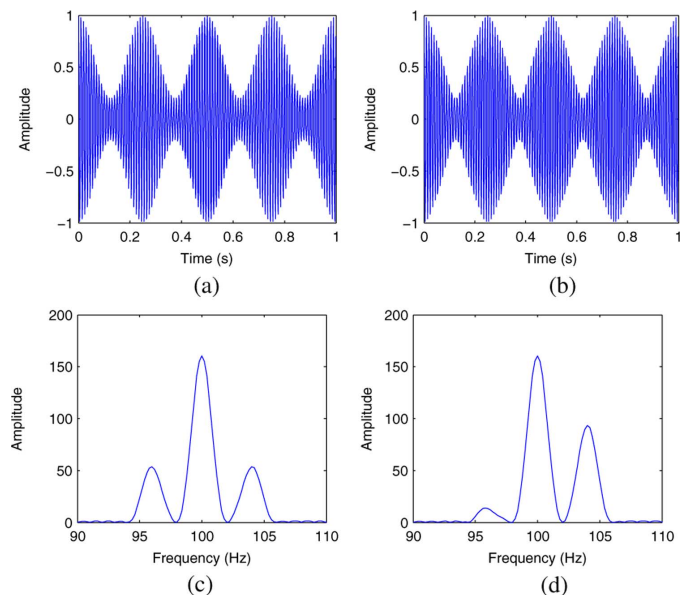

Fig. 15. A modulated signal is shown in (a). Note that its spectrum in (c) is symmetric as a result. If a signal is passed through a subband filter and results in unequal sidebands, shown in (b), an asymmetric spectrum results, as in (d). See Fig. 16 for estimated modulators from these signals.

A problem raised by subband decomposition is that it is possible that the channel filters will unequally distort the modulation components. In this situation, the modulation spectra is not
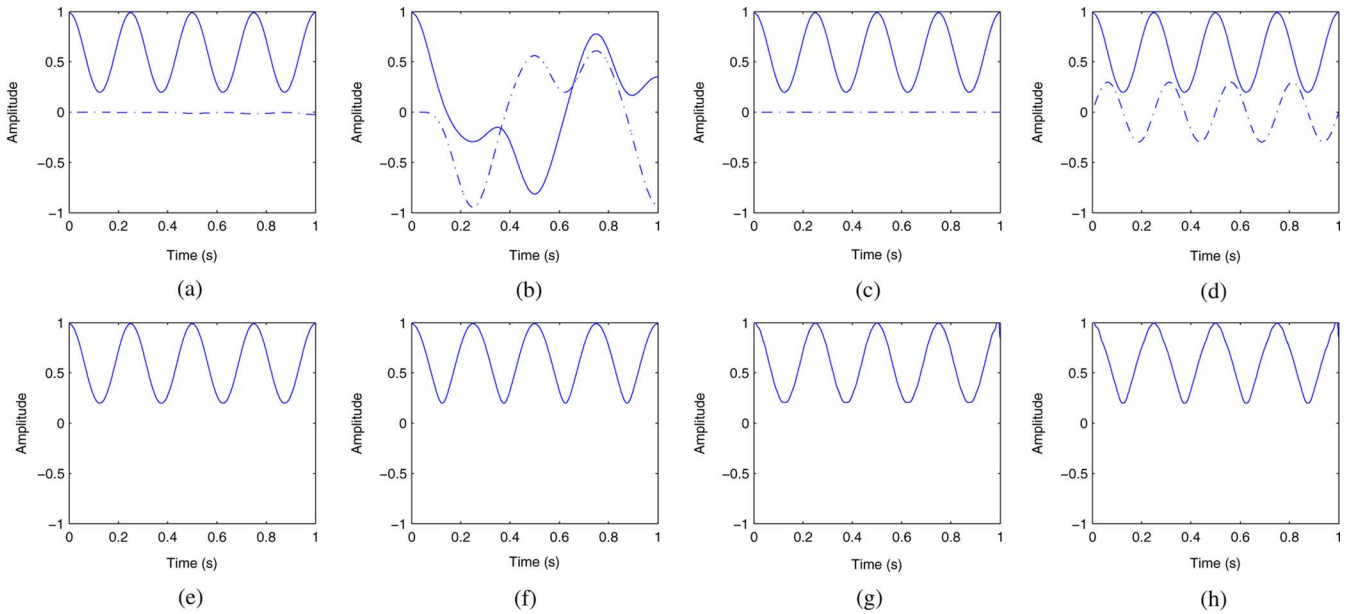
Fig. 16. Using subband filters can lead to asymmetric spectra, an example of which shown in Fig. 15. The first row shows the complex modulator (real part in solid, imaginary in dot-dash) extracted with subband carrier demodulation, with both a spectral center-of-mass estimated carrier and an exact carrier. The bottom row shows real modulators extracted with Hilbert and linear domain optimization (introduced in Section IV-B). The complex modulator is better mathematically, but a real modulator has more perceptual relevance. (a) Symmetric signal, Subband carrier demodulation with estimated carrier. (b) Asymmetric signal, Subband carrier demodulation with estimated carrer. (c) Symmetric signal, Subband carrier demodulation with exact carrier. (d) Asymmetric signal, Subband carrier demodulation with exact carrier. (e) Symmetric signal, Hilbert. (f) Asymmetric signal, Hilbert. (g) Symmetric signal, Linear Domain Optimization. (h) Asymmetric signal, Linear Domain Optimization.
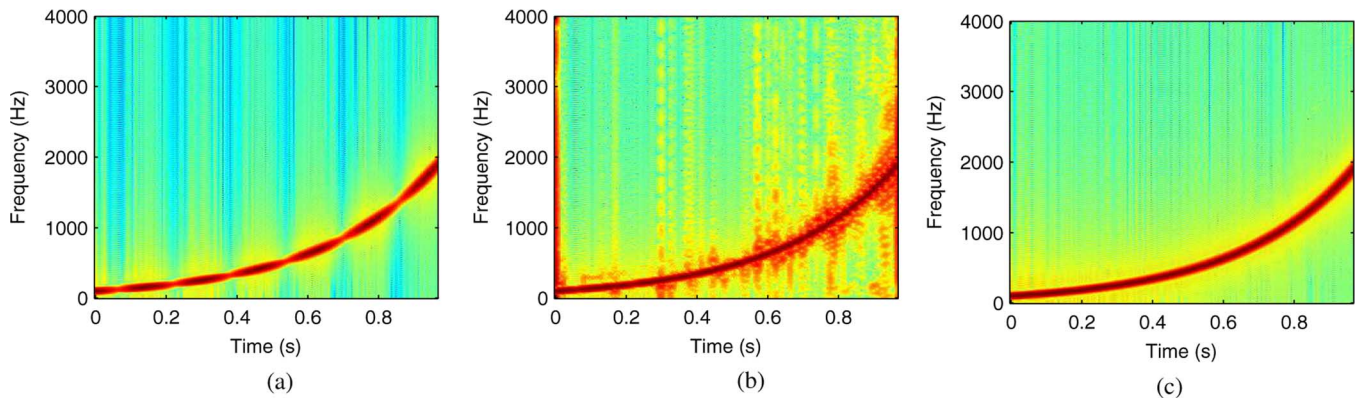

Fig. 17. Subband decomposition is often necessary for harmonic signals, but introduces its own problems as a carrier transitions through each passband. (a) shows the spectrogram of a 6-Hz modulated chirp. The reconstructed carrier after subband demodulation with a 51-channel, 80-Hz Hamming filterbank decomposition is shown in (b) and contains high levels of spectral spatter, while the extracted wideband Hilbert carrier in (c) is clean. (a) Original. (b) With subbands. (c) Without subbands.

symmetric within the channel. Atlas *et al.* suggest that a complex modulator is necessary to properly demodulate a channel with distortion from subband filtering [2]. Mathematically, this argument is sound, but it is conceptually difficult to justify the need for a complex modulator to describe the amplitude envelope of real signals.

To demonstrate this issue, Fig. 15 shows a modulated signal with a symmetric spectrum, and also a corresponding signal with the same frequency components but an asymmetric spectrum because one sideband is close to the band-edge. Fig. 16 shows a series of modulators, both complex and real, extracted from these signals. There are two details to note in these plots. First, there is a significant difference in results between the complex modulators (all extracted with subband carrier demodu-

lation) derived from an estimated carrier [Fig. 16(b)] and an exact carrier [Fig. 16(d)] in the case of the spectrally asymmetric signal. Because of the asymmetry, the spectral center-of-mass estimate will lead to a slightly skewed carrier frequency estimate, resulting in the differences. In the symmetric case, the estimate is much more accurate, and so this issue does not arise (and the complex modulator is not even necessary).

A second aspect of these plots to note is that the real modulators do effectively track the real amplitude of the signal. So, while the complex modulator mathematically solves the asymmetric spectrum, the real modulators do contribute valuable information that is potentially even more relevant perceptually. Which of these modulators is the desirable outcome is unclear, and would likely be application specific.

Furthermore, recombining the post-analysis subbands raises troubling issues with subband decomposition. Because a demodulation process is, by definition, nonlinear, the wideband modulator and carrier cannot easily be restored from the channel modulators and carriers. Fig. 17 shows a modulated chirp (a) and the demodulated carriers that result from Hilbert demodulation with (b) and without (c) subband decomposition. The distortions added as a result of the subband decomposition are clear in (b). The wideband demodulation (c) performs well in this case because the modulator used is non-negative and the carrier has only one sinusoidal component, conditions that Hilbert demodulation can easily solve.

Recent work [13] has suggested implementing pitch-tracking filters in order to eliminate inconsistent alignments between the filters and the components. While this approach improves some of the issues that arise with a subband decomposition, such as non-stationary components, the results of this analysis appear more like those found with sinusoidal peak tracking or a phase vocoder.

Including the subband decomposition also signals a movement away from modeling the physical sound. By demodulating channels, spectral effects from filtering are also classified as amplitude modulation. For example, demodulating speech on a subband level will classify the spectral variations from the shifting vowel formants as modulation when in fact these resonances are created by filtering in the vocal tract. Demodulating without the subband decomposition leaves these filtering effects in the carrier, while the modulator consists of only relative amplitude variations.

We wonder if a subband decomposition should be included in the demodulation process. We believe the ability of our methods to demodulate wideband signals without a subband decomposition is an asset that will allow for further insight into the issue.

VI. DISCUSSION

We have shown a demodulation approach that allows a large range of signals to be analyzed. The examples presented above show the strengths and weaknesses of the two demodulation methods described in this paper.

Both the linear-domain and the logarithmic-domain methods are able to differentiate between low-frequency and high-frequency modulators. This is an important distinction in a perceptual context, and it is a feature that has not been demonstrated in past demodulation work.

The presented implementation of logarithmic-domain demodulation performs very well with single sinusoidal carriers. For stochastic and harmonic carriers, the cost on the carrier grows, and so the balance between the modulator and carrier costs is skewed, leading to carrier components appearing in the modulator [Figs. 8(b) and 10(b)]. This can be managed by adding extra weight to the modulator, as seen in Figs. 8(c) and 10(c). However, finding the right weight to select the desired modulator can be challenging and require several iterations. This process could be posed as an additional (non-convex) optimization problem that iterates over several weights and compares the modulators to find the optimum, which may be a good direction for future work. However, it seems that the ideal solution is one that does not require changing the weights at all. Deriving an implementation that handles stochastic and harmonic carriers as well as this implementation demodulates single sinusoidal carriers would be a significant improvement.

The implementation of linear-domain demodulation performs very well, regardless of the carrier. The examples show that it accurately demodulates harmonic and stochastic carriers and even real speech, as well as simpler cases. The method does have two potential weaknesses, though. First, a linear-domain demodulation is unable to track the sign changes in suppressed-carrier modulation [Fig. 4(c)]. However, it is unclear if this ability is necessary for acoustic demodulation. Additionally, the spectral weighting in the cost function bandlimits the modulator, so the issues with the artifacts in modulation filtering of the Hilbert envelope are avoided. The second potential weakness is that the method does not infer any structure of the carrier, but rather only considers the samples. This can lead to an extracted modulator that is slightly less than the original modulator [as seen, for example, in the peak at 0.55 seconds of Fig. 10(d)] because of slight phase differences between the carrier frequency and the sampling frequency. It also causes inaccurate demodulation for very noisy signals. Again, though, it is unclear if this is a relevant distinction in acoustic demodulation. However, future work to remove these potential issues would be an improvement.

VII. CONCLUSION

We have introduced here two new methods for demodulation of audio signals using convex optimization. Each method uses a different approach to fit the demodulation relationship within the rules of disciplined convex programming, and, as a result, each method has its drawbacks. However, the new methods do contribute many novel additions to modulation signal processing. They allow for the definition of a modulator to be designed and clearly outlined within the cost function, while previous methods have only offered minimal control.

By solving demodulation as an optimization problem, we obtained clean, low-frequency modulators for harmonic and noisy signals, also a task that previous methods could not perform. Because of this improved performance without the necessity for narrowband decompositions, it is possible to examine the relationship between subband channel width, such as that seen in Fig. 14. The methods shown here are the only to date that allow for this sort of study.

We believe that posing demodulation as an optimization problem is a very promising direction for this field of research. The decision process between pitch and modulation corresponds well with the minimization of a summation of cost functions, and the quality of results presented here for a wide array of signals demonstrates the exciting potential of the approach.

References

[1] L. Atlas and C. Janssen, "Coherent modulation spectral filtering for single-channel music source separation," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, 2005, pp. 461–464.

[2] L. Atlas, Q. Li, and J. Thompson, "Homomorphic modulation spectra," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, 2004, pp. 761–764.

[3] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004 [Online]. Available: http://www.stanford.edu/~boyd/cvxbook

[4] P. Clark and L. Atlas, "A sum-of-products model for effective coherence modulation filtering," in *Proc. ICASSP*, 2009.

[5] P. Clark and L. Atlas, "Time-frequency coherent modulation filtering of nonstationary signals," *IEEE Trans. Signal Process.*, vol. 57, no. 11, pp. 4323–4332, Nov. 2009.

[6] M. Elihilali, T. Chi, and S. A. Shamma, "A spectro-temporal modulation index (STMI) for assessment of speech intelligibility," *Speech Commun.*, vol. 41, pp. 331–348, 2003.

[7] O. Ghitza, "On the upper cutoff frequency of the auditory critical-band envelope detectors in the context of speech perception," *J. Acoust. Soc. Amer.*, vol. 110, no. 3, pp. 1628–1640, Sept. 2001.

[8] M. Grant, S. Boyd, and Y. Ye, "cvx Users' Guide," Sep. 2008, available at [Online]. Available: http://www.stanford.edu/~boyd/cvx

[9] S. Greenberg and B. E. D. Kingsbury, "The modulation spectrogram: In pursuit of an invariant representation of speech," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, 1997, pp. 1647–1650.

[10] M. G. Heinz and J. Swaminathan, "Quantifying envelope and fine-structure coding in auditory-nerve responses to chimaeric speech," *JARO*, vol. 10, no. 3, pp. 407–423, Sep. 2009.

[11] G. Hu and D. Wang, "Monaural speech segregation based on pitch tracking and amplitude modulation," *IEEE Trans. Neural Netw.*, vol. 15, no. 5, pp. 1135–1150, Sep. 2004.

[12] B. E. D. Kingsbury, N. Morgan, and S. Greenberg, "Robust speech recognition using the modulation spectrogram," *Speech Commun.*, vol. 25, pp. 117–132, 1998.

[13] Q. Li and L. Atlas, "Coherent modulation filtering for speech," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, 2008, pp. 4481–4484.

[14] A. V. Oppenheim, R. W. Schafer, and T. G. Stockham, "Nonlinear filtering of multiplied and convolved signals," *Proc. IEEE*, vol. 56, no. 8, pp. 1264–1291, Aug. 1968.

[15] S. Schimmel, "Theory of modulation frequency analysis and modulation filtering, with applications to hearing devices," Ph.D. dissertation, Univ. of Washington, Seattle, 2007.

[16] S. Schimmel and L. Atlas, "Coherent envelope detection for modulation filtering of speech," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, 2005, pp. 221–224.

[17] Z. M. Smith, B. Delgutte, and A. J. Oxenham, "Chimaeric sounds reveal dichotomies in auditory perception," *Nature*, vol. 416, no. 6876, pp. 87–90, Mar. 2002.

[18] R. E. Turner and M. Sahani, "Modeling Natural Sounds with Modulation Cascade Processes," in *Advances in Neural Information Processing Systems 20*, J. C. Platt, D. Koller, Y. Singer, and S. Rowels, Eds.   Cambridge, MA: MIT Press, 2008, pp. 1545–1552.

[19] L. Xu and B. E. Pfingst, "Relative importance of temporal envelope and fine structure in lexical-tone perception," *J. Acoust. Soc. Amer.*, vol. 114, no. 6, pp. 3024–3027, Dec. 2003.

Gregory Sell received the B.A. and M.A. degrees in music in 2005 and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 2007, where he is currently pursuing the Ph.D. degree in the Center for Computer Research in Music and Acoustics, Stanford University.

His research interests vary from modulation signal processing and auditory perception to music information retrieval and musical database search and organization.

Malcolm Slaney (F'10) received the Ph.D. degree in electrical engineering from Purdue University, West Lafayette, IN.

He is a Principle Researcher with Yahoo! Research, Sunnyvale, CA, and a Consulting Professor at Stanford University, Stanford, CA. He is a coauthor of the book *Principles of Computerized Tomographic Imaging (a Classic in Applied Mathematics)* (SIAM, 2001) and coeditor of the book *Computational Models of Hearing* (IOS Press, 2001).