IMPROVED FREQUENCY ESTIMATION IN SINUSOIDAL MODELS THROUGH ITERATIVE LINEAR PROGRAMMING SCHEMES

Vighnesh Leonardo Shiv
Catlin Gabel School
shivv@catlin.edu

ABSTRACT

Sinusoidal modeling systems are commonly employed in sound and music processing systems for their ability to decompose a signal to its fundamental spectral information. Sinusoidal modeling is a two-phase process: sinusoidal parameters are estimated in each analysis frame in the first phase, and these parameters are chained into sinusoidal trajectories in the second phase. This paper focuses on the frequency estimation aspect of the first phase. Current methods for estimating parameters rely heavily on the resolution of the Fourier transform and are thus hindered by the Heisenberg uncertainty principle. A novel approach is proposed that can super-resolve frequencies and attain more accurate estimates of sinusoidal parameters than current methods. The proposed algorithm formulates parameter estimation as a linear programming problem, in which the L^1 norm of the residual component of the sinusoidal decomposition is minimized. It achieves 3.5 times the frequency resolution of Fourier-based approaches.

1. INTRODUCTION

Sinusoidal modeling, the problem of representing a signal as a summation of quasi-stationary sinusoids, is a fundamental task in sound and music signal processing. Sinusoidal representations have many important applications. From an analysis standpoint, they transform pulse code modulated signals into meaningful representations for perceptual tasks like auditory scene analysis [1]. From a synthesis viewpoint, they make signals amenable to manipulations including time-scale and pitch-scale modifications, and can be efficiently coded at a low frame rate due to their slowly time-varying nature [2].

State-of-the-art sinusoidal modeling systems adopt a two-phase approach to the problem. The first phase involves extracting the sinusoidal parameters—amplitude, frequency, and phase—in each analysis frame of the signal. The second phase chains these parameters across analysis frames into sinusoidal trajectories, generating sinusoidal “births” and “deaths” as necessary. This paper is concerned with the parameter estimation phase of these systems.

Most recent parameter estimation algorithms employ approaches based on Fourier analysis [3]. First, peaks in the short-time Fourier transform (STFT) of a given analysis frame are detected, either by detecting local maxima greater than a fixed threshold [2] or through more sophisticated algorithms like the cross-correlation method [1]. Then, sinusoidal frequency estimates are refined. Methods for estimating frequencies include interpolation techniques [4], [5], [6], time-frequency reassignment approaches [7], and derivative-based methods [8], [9], [10]. Finally, sinusoidal amplitudes and phases are estimated, either directly from the frame’s frequency spectrum or through more intricate methods like iterative analysis [11], [1].

However, Fourier analysis-based systems are subject to the Heisenberg uncertainty principle: they have limited simultaneous resolution in the time and frequency domains. Thus, long analysis frames are necessary to distinguish nearby frequencies, while short analysis frames are necessary to capture rapid parametric changes and short sound events like musical notes. For an analysis frame of length w and a rectangular window function, the Fourier transform can resolve frequencies $\frac{2}{w}$ apart [4]. Additionally, spectral leakage can cause amplitudes to be overestimated from the frequency spectrum. For these reasons, pure Fourier analysis-based systems may not be optimal for sinusoidal parameter estimation.

In this paper, a novel algorithm for sinusoidal frequency estimation is proposed, that has the ability to super-resolve frequencies with high precision. The algorithm formulates the sinusoidal model as a linear program, in which the objective function is the L^1 norm of the residual component of the signal’s sinusoidal decomposition. The quasi-stationary nature of the sinusoids, in conjunction with properties of the simplex algorithm [12], can be exploited to estimate parameters from analysis frame to analysis frame at a high frame rate efficiently.

Note that the problem of trajectory continuation in sinusoidal modeling is beyond the scope of this paper. This problem has been addressed in such works as [13], etc. Existing trajectory continuation algorithms can be applied to the proposed parameter estimation algorithm in order to extend it to a complete sinusoidal modeling system, although future work would include a trajectory continuation system optimized to the time domain-based nature of the proposed system.

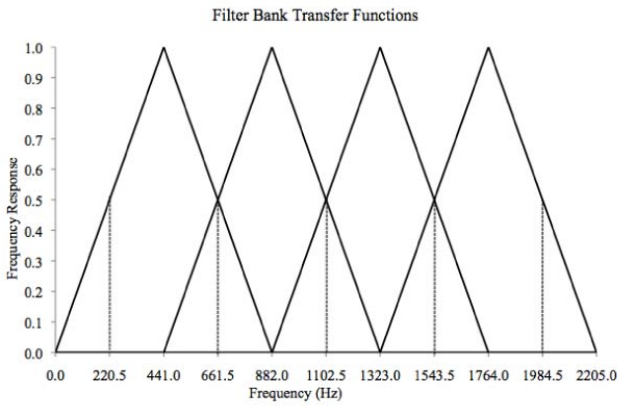

Figure 1: A partial graphical representation of the frequency responses of the filters employed in the system. While only four filters are illustrated in this figure, the filter bank used consists of 25 filters extending beyond 10 kHz. The filters not shown follow the same pattern as the filters displayed in this diagram.

2. SIGNAL PRE-PROCESSING

The input signal is first pre-processed by passing it through a filter bank. Each bandpassed component is more sinusoidally sparse than the raw signal, simplifying sinusoidal analysis on a per-component basis and improving the computational efficiency of parameter estimation. Separating the sinusoids into separate bands also lowers the dependence of sinusoidal analysis quality on signal complexity, increasing robustness.

The proposed system passes the input signal through a filter bank composed of 25 linearly spaced 9 ms-wide Blackman windowed-sinc filters, such that filter r had cutoff frequencies $441r - 220.5$ and $441r + 220.5$ Hz, $r \in [0, 25)$. The filters have approximately triangular responses within 882 Hz-wide bands with 50% overlap. Up to 10 kHz, each frequency has a frequency response of at least 0.5 in some filter band, preventing instability when recombining parameter estimates in bandpassed components.

3. LINEAR PROGRAM

This section describes the formulation of the sinusoidal modeling problem as a linear programming problem. This linear programming setup forms the core of the proposed algorithm. A hypothesis set of K sinusoidal frequencies is assumed; the next section will describe how these K frequencies are generated and how the discussed linear program can be used for sinusoidal parameter estimation.

The traditional sinusoids-plus-noise model takes on either of the following equivalent forms:

$$x(t) = \sum_{k=1}^{K} a_k(t) \sin\left(\phi_k(t)\right) + r(t) \quad (1)$$

$$x(t) = \sum_{k=1}^{K} a_k(t) \sin\left(\int_0^t \omega_k(u) du + \phi_k(0)\right) + r(t) \quad (2)$$

where a_k, ω_k, and ϕ_k are the time-varying amplitude, frequency, and phase respectively of the kth sinusoid, and r is the noise residual. a_k and ω_k are assumed to be locally stable, ϕ_k is assumed to be locally linear, and r is assumed to be a stochastic process. An analysis frame l of length w and a rectangular window function is traditionally defined as

$$x^l(t) = \begin{cases} x(t + hl) & \text{if } t \in \left[-\frac{w}{2}, \frac{w}{2}\right), \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where h is the hop size between analysis frames. Over a short analysis frame, the sinusoidal amplitudes and frequencies can be assumed to be constant. Thus, the sinusoids-plus-noise model of the lth analysis frame can be represented in the form

$$x^l(t) = \sum_{k=1}^{K} a_k^l \sin(\omega_k^l t + \phi_k^l) + r^l(t), \forall t \in \left[-\frac{w}{2}, \frac{w}{2}\right) \quad (4)$$

where a_k^l, ω_k^l, and ϕ_k^l are the amplitude, frequency, and initial phase respectively of the kth sinusoid in the lth analysis frame, and r^l is the residual component in the lth analysis frame. Given the assumption that the hypothesis set of frequencies is known, the goal here is to determine the sinusoidal amplitudes and phases.

The problem is first formulated as an optimization problem in which the L^1 norm of the residual component of the sinusoidal decomposition is minimized, as follows:

$$\begin{gathered} \text{Minimize } \sum_{t=1}^{w} \left|r^l(t)\right| \\ \text{such that, } \forall t \in \left[-\frac{w}{2}, \frac{w}{2}\right), \\ x^l(t) = \sum_{k=1}^{K} a_k^l \sin(\omega_k^l t + \phi_k^l) + r^l(t). \end{gathered} \quad (5)$$

From this form, the problem can then be reformulated into a linear programming problem. First, by defining the variables

$$c_k^l \doteq a_k^l \cos(\phi_k^l), \text{ and} \quad (6)$$

$$s_k^l \doteq a_k^l \sin(\phi_k^l), \quad (7)$$

the sinusoidal component of the sinusoidal decomposition can be represented as a linear expression in terms of c_k^l and s_k^l for each k. If c_k^l and s_k^l can be solved for, a_k^l and ϕ_k^l can be trivially derived. This transforms the optimization problem to:

$$\begin{gathered} \text{Minimize } \sum_{t=1}^{w} \left|r^l(t)\right| \\ \text{such that, } \forall t \in \left[-\frac{w}{2}, \frac{w}{2}\right), \\ x^l(t) = \sum_{k=1}^{K} \left[\sin(\omega_k^l t) c_k^l + \cos(\omega_k^l t) s_k^l\right] + r^l(t). \end{gathered} \quad (8)$$

The constraints are now linear in terms of all the used variables, but the objective function is not. This can be fixed by expressing each residual data point as a difference of nonnegative variables:

$$r^l(t) = r^l_+(t) - r^l_-(t), \tag{9}$$

$$r^l_+(t), r^l_-(t) \geq 0. \tag{10}$$

The constraints are thus still linear expressions in terms of the used variables. As $r^l_+(t)$ and $r^l_-(t)$ are not independent variables, they cannot both be in the basis of the optimal solution found by a linear programming solver. Thus, for each t, one of $r^l_+(t)$ and $r^l_-(t)$ must equal zero, implying that

$$\left|r^l(t)\right| = r^l_+(t) + r^l_-(t). \tag{11}$$

As the objective function is now a linear expression, the final linear programming problem can be written:

$$\begin{gathered}\text{Minimize } \sum_{t=1}^{w} \left[r^l_+(t) + r^l_-(t)\right] \\ \text{such that, } \forall t \in \left[-\frac{w}{2}, \frac{w}{2}\right), \\ x^l(t) = \sum_{k=1}^{K} \left[\sin(\omega^l_k t)c^l_k + \cos(\omega^l_k t)s^l_k\right] + r^l_+(t) - r^l_-(t), \\ r^l_+(t), r^l_-(t) \geq 0. \end{gathered} \tag{12}$$

Given a hypothesis set of K frequencies, this program can be solved. While either the simplex algorithm or interior-point methods can be used to solve the program, the simplex algorithm is preferred for reasons that will be discussed in Section 6.

4. PARAMETER ESTIMATION

This section entails the linear programming-based algorithm for estimating the sinusoidal parameters. The general concept is to use linear programming with a temporal model to iteratively reduce the size of the hypothesis set of frequencies, while also retrieving information about amplitudes, phases, and phase derivatives. The result is that frequencies can be super-resolved and amplitudes and phases can be fit optimally to the input signal.

The input signal, the output of the original signal passed through a single filter, is first partitioned into short overlapping analysis frames of length 46 milliseconds, or $w = 2048$ for CD-quality audio. To estimate the sinusoidal parameters in a given analysis frame l, the analysis frame is first zero-padded with zero-padding factor z. z represents the theoretical increase in frequency resolution of the proposed algorithm over Fourier-based methods, although practical issues set a bound on how high z can be, as discussed in Section 5. A fast Fourier transform of the zero-padded analysis frame is then taken, resulting in spectrum $\hat{x}^l(t)$. The frequencies corresponding to each frequency bin can be considered an inital hypothesis set of sinusoidal frequencies. This hypothesis set can be pruned by eliminating frequency bins of low magnitude, since although an FFT will result in false positives due to spectral leakage, it will rarely result in false negatives. A threshold can be set proportional to the overall energy in the signal; frequencies with magnitudes below this threshold can be eliminated.

The linear program in (12) can be solved using the hypothesis set of frequencies, to determine estimates for the amplitudes and phases of all frequencies in the set. From these values, sinusoids of low amplitudes can be attributed to overfitting noise or modulations in the signal. These sinusoids' frequencies can thus be eliminated from the hypothesis set, and the linear program can be resolved to sharpen the sinusoidal parameters. This iterative pruning procedure of eliminating low-amplitude frequencies from the hypothesis set and resolving the linear program can be repeated until the hypothesis converges.

Fourier-based methods operating on a non-zero-padded rectangular-windowed analysis frame can resolve frequencies two frequency bins apart [4]. The proposed method aims to resolves frequencies two frequency bins apart in the zero-padded spectrum, although it may be able to resolve frequencies even more finely under certain conditions. The reason for this is that the procedure thus far makes an implicit assumption that the frequencies present in the signal will be near enough to a single frequency in E such that the linear program can accurately model the signal. However, this assumption is unsafe to make, as a frequency could very well lie, in the worst case, half way between two frequency bins, undermining the procedure. In such cases where frequencies are close to two adjacent frequencies, a smearing effect will occur making both frequencies appear significant, a phenomenon similar to spectral leakage in Fourier-based methods. If a single sinusoid can be assumed to be responsible for up to two nearby frequencies occurring in the hypothesis set, the proposed algorithm should have a frequency resolution of $\frac{2}{zw}$, z times more sensitive than Fourier analysis.

One remaining problem is that of how to estimate the frequency of a sinusoid "activating" the two frequencies adjacent to it. If the frequency is close enough to a frequency in the original hypothesis set, the smearing effect will not occur; in most cases, however, the smearing effect has be resolved somehow. The proposed approach here begins by estimating the number of resolvable sinusoids given the hypothesis set. The hypothesis set is partitioned into the minimum number of non-overlapping subsets $H_1, H_2, \ldots, H_n$ such that each subset contains only frequency bins adjacent to one another. As a single sinusoid will activate one or two frequencies around it and the algorithm's resolving power is two frequency bins, it thus follows that the number of resolvable sinusoids in H_i is

$$\left\lfloor \frac{|H_i| + 1}{2} \right\rfloor. \tag{13}$$

From here, the sinusoidal frequencies should be estimated. For each H_i, there are two cases to consider: where $|H_i|$ is even and where it is odd. In the former case, H_i can be further partitioned into two-element subsets of adjacent frequency bins. The sum of the two sinusoids in a given subset H_{ij} can be represented as third sinusoid with time-varying amplitude and frequency. At time $t = 0$, that time-varying frequency is [1]:

$$f = \frac{f_2 + f_1}{2} + \frac{(f_2 - f_1)(a_2 - a_1)\left(\sec^2\left(\frac{\phi_2-\phi_1}{2}\right)\right)}{2(a_2 + a_1)\left(1 + \tan^2\left(\frac{\phi_2-\phi_1}{2}\right)\left(\frac{a_2-a_1}{a_2+a_1}\right)^2\right)} \quad (14)$$

where the two amplitudes, frequencies, and phases are those corresponding to the two elements of H_{ij}. Thus, this equation can be used to estimate the true sinusoidal frequency with sufficient accuracy.

In the latter case, H_i contains an odd number of elements. In this case, a new set H_i' can be generated, containing a number of elements equal to $|H_i| + 1$. The distances between adjacent frequencies in H_i' are of the same spacing as in H_i, and the mean frequencies of both sets are equal. Once all sets H_i' have been generated, a linear program can be solved with a hypothesis set that replaces all H_i subsets with H_i' subsets. This allows this case to be reduced to the case of H_i having an even number of elements, for which a method has already been described.

Once the final set of frequencies is determined, a final linear program involving all of these frequencies can be performed to determine the amplitudes and phases of each sinusoid within an analysis frame.

5. FREQUENCY SPACING

One parameter of the algorithm that must be tuned is the frequency spacing between frequency bins of the FFT taken. The larger the frequency spacing is, the lower the frequency resolution of the algorithm. The algorithm may fit parameters poorly to the input signal, and the information determined may not be useful for high-level processing or generalizable to future analysis frames. One inherent advantage to using a temporal model is the ability to super-resolve frequencies, causing a bias toward a smaller frequency spacing. However, if this spacing is made too small, overfitting will occur. Many spurious frequencies may be detected in an attempt to fit a sinusoidal model as closely to the original signal as possible, undermining the purpose of sinusoidal modeling and making analysis of sinusoidal births and deaths difficult during trajectory continuation. A balance between these two extremes must be determined empirically.

6. EFFICIENCY CONSIDERATIONS

The proposed linear programming-based algorithm brings into question the preference for linear programming over a linear least squares approach. Indeed, a similar approach could be developed in which the L^2 norm of the residual component of the sinusoidal decomposition is instead minimized through linear regression, as opposed to the L^1 norm through linear programming. The main reason to opt for the linear programming approach is the efficiency gain it provides. Redundant information from linear program to linear program can be exploited when the simplex algorithm is used. The same cannot be said for linear least squares, however. While efficiency is not the main focus of this paper, as linear programming is still a slow process, these efficiency measures should be included for the sake of both completeness and comparison to least squares methods.

The simplex algorithm is, in essence, a greedy search method that traverses a convex polytope representing the linear program to solve. Conceptually, with each iteration it moves to an adjacent vertex on the polytope such that the objective function becomes closest to optimal. After enough such iterations, the simplex method reaches a global maximum, as all local maxima are global maxima in a linear programming problem. Given the greedy nature of the algorithm, it makes sense that the efficiency measures discussed in this section focus on using redundancy to determine an initial basis close to the optimal solution, thus skipping phase I of the simplex algorithm and dramatically shortening phase II.

Redundant information can be exploited in the proposed system in two ways: from linear program to linear program both within and across analysis frames. Consider the former case first. When analyzing a single analysis frame, the numerous linear programs solved differ only in that some number of variables are removed from the system from linear program to linear program, changing the coefficient matrix. The dual simplex algorithm can be employed to efficiently make use of this redundancy. Rather than removing variables from the system, constraints can be added fixing those variables to zero. The dual simplex can then use the previously computed optimal basis as an initial dual feasible basis. As the variables removed were near zero to begin with, this initial basis should be only a few iterations away from the optimal basis, increasing efficiency. In contrast, linear least squares systems cannot increase their efficiency in this way. If the coefficient matrix changes in a linear least squares system, the entire system must be computed again. As the coefficient matrix slightly changes for each linear system solved for a given analysis frame as frequencies are eliminated, linear least squares becomes an expensive process for this task, unlike linear programming.

Now consider the latter case, where redundancy can be exploited across analysis frames. Given that the sinusoids are assumed to have quasi-stationary parameters, the final parameters found in the previous analysis frame should ideally be reused as the initial parameters for the current analysis frame. Given a high frame rate, it is possible to take advantage of this redundancy using the simplex algorithm. The frame rate is assumed here to be the original sampling rate of the signal to maximize the amount of redundancy between frames, so $h = 1$. The most efficient way to take advantage of this redundancy involves slightly altering the analysis frame system. Rather than defining an analysis frame as a shifted signal that is windowed from its first time sample, it can instead be defined as a non-shifted signal that is multiplied by a shifted window. While these definitions appear to be equivalent, the latter definition allows more efficient transitions between linear programs. A linear programming system requires only a single row change from frame to frame, making matrix inversions more efficient. The linear program under this definition of an analysis frame l is thus

$$\text{Minimize} \sum_{t=1}^{w} \left[r_+^l(t) + r_-^l(t)\right]$$
$$\text{such that, } \forall t \in \left[l - \frac{w}{2}, l + \frac{w}{2}\right),$$
$$x(t) = \sum_{k=1}^{K} \left[\sin(\omega_k^l t) c_k^l + \cos(\omega_k^l t) s_k^l\right] + r_+^l(t) - r_-^l(t),$$
$$r_+^l(t), r_-^l(t) \geq 0. \tag{15}$$

Note that the phases retrieved from c_k^l and s_k^l must be shifted accordingly, such that

$$\phi_k(l) \equiv \arg(c_k^l + i s_k^l) + \frac{\omega_k^l l}{f_s} \mod 2\pi. \tag{16}$$

Now that the analysis frame has been redefined such that a single row of the linear programming system changes from frame to frame, the previous analysis frame's parameters can be reused almost completely as the current analysis frame's initial parameters efficiently. Two situations should be considered here: when one of the two residual variables in the changing row is basic, and when both are non-basic, i.e. zero. Consider the former situation first. Luckily, this situation is both more common and more efficient to work with. Every initial value of c_k^l and s_k^l can simply be set to c_k^{l-1} and s_k^{l-1}, respectively, and every residual variable $r_+^l(t)$ and $r_-^l(t)$, for all $t < l + \frac{w}{2} - 1$, can be set to $r_+^{l-1}(t)$ and $r_-^{l-1}(t)$, respectively. The last variables to determine, $r_+^l(l + \frac{w}{2} - 1)$ and $r_-^l(l + \frac{w}{2} - 1)$, can simply be solved for using the values of the other variables and the restriction that one of the variables must equal zero. This supplies an initial primal feasible basis close to the optimal basis. Considering the second situation, the same steps can be repeated. However, the resulting values do not represent a basic feasible solution, as the basis contains one excess basic variable. An easy way to deal with this issue is to remove the variable that has the value closest to zero, and simply solve for the values of the other variables using a linear system. A single solution of a linear system is generally efficient enough given the sparsity of the matrix. Given an initial basis, the primal simplex method can be employed to traverse to the optimal basis.

7. SYSTEM EVALUATION

The system was evaluated experimentally in two ways. First, the ability of the system to detect the number of resolvable sinusoids was tested, in order to tune the zero-padding factor of the system. The zero-padding factor z and the number of sinusoids n were varied. The test set consisted of synthetic signals of n sinusoids separated in frequency by the expected resolution, where the lowest frequency was randomly generated. The metric used to evaluate the system was the average absolute error in the number of sinusoids detected.

As the data shows in Figure 2, the detection error behaves similarly for different numbers of sinusoids. For zero-padding factors of 3.5 or less, the system consistently detects the correct number of sinusoids, while the behavior of the system becomes more erratic with an upward trend in detection error as the zero-padding factor gets even higher. Taking this information into account, a zero-padding factor of 3.5 was selected for the second experiment.

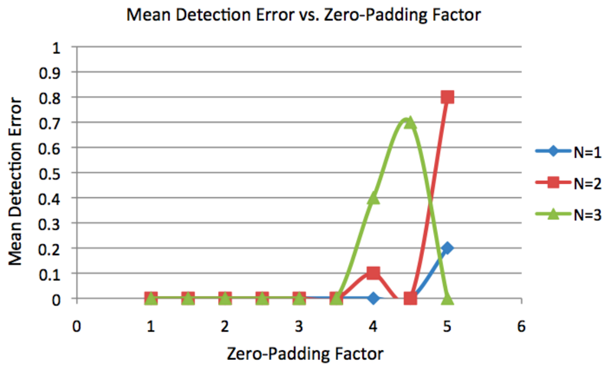

Figure 2: Curves indicating the relationship between the average absolute error in detecting the number of frequencies and the zero-padding factor used. Each curve represents signals of a different number of sinusoids.

The purpose of the second experiment was to directly test the frequency estimation quality of the system. A test set of synthetic signals was constructed in which, in addition to the sinusoidal amplitudes, the exponential decay rates of the sinusoids were varied, so as to test the robustness of the system to common musical modulations. The decay rate of the sinusoids k, where the time-varying amplitude is of the form ae^{-kt}, and the number of sinusoids s were varied. The metric measured here was the average absolute frequency estimation error among sinusoids. As Figure 3a shows, the average frequency estimation error is consistently low and robust against both the number of sinusoids and their exponential decay rates, with the average error never exceeding 2.5 Hz.

Although this paper focuses only on frequency estimation in sinusoidal models, the linear program additionally determines amplitudes and phases of the sinusoids. Thus, it is interesting to analyze the error in determining the amplitudes and phases. When the sinusoids are far apart in frequency, the amplitudes and phases are found with high precision, so it is more interesting to study the effects that sinusoids near each other in frequency have on each other when solving for amplitudes and phases. On the same test set used for measuring frequency estimation error, the average Euclidean distance between target and output amplitude-phase vectors as a percent of the magnitude of the target vector was also measured, a standard metric used in sinusoidal modeling research [1]. The target vector incorporated the amplitude and phase at the center of the analysis frame. The data in Figure 3b indicates that sinusoids near each other do impact each other's computed parameters. There are clear upward trends in amplitude-phase error with respect to both the exponential decay rate and number of sinusoids, which should be studied further in future work.

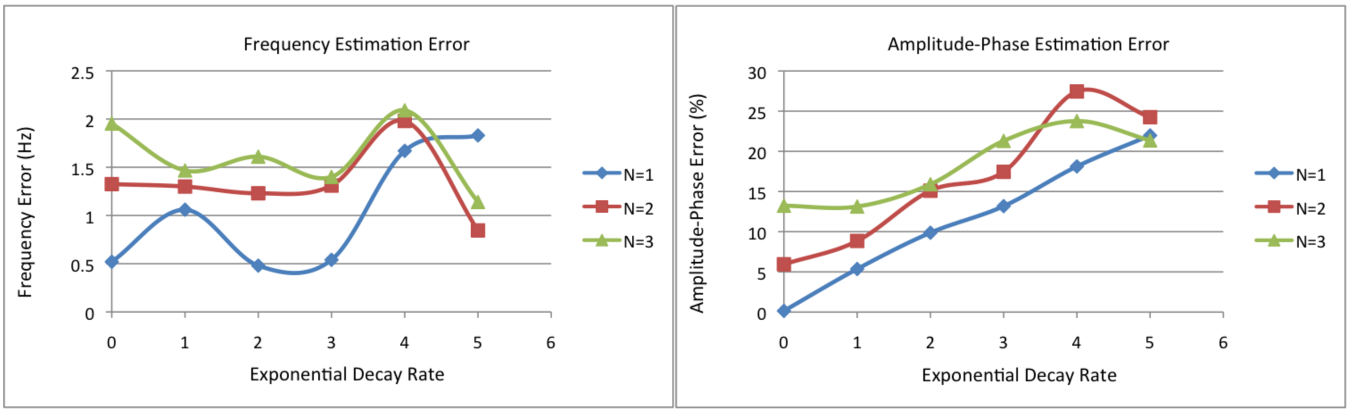

(a) Sinusoidal amplitudes vs. time

(b) Sinusoidal frequencies vs. time

Figure 3: (a) Curves indicating the relationship between the average absolute error in estimating the sinusoidal frequencies and the exponential decay rate of the synthetic sinusoids used. (b) Curves indicating the relationship between the average error in estimating the sinusoidal amplitude and phase, as per the stated metric, and the exponential decay rate of the synthetic sinusoids used.

8. CONCLUSION

In this paper, a novel algorithm for sinusoidal frequency estimation was proposed. The algorithm was based on using linear programming to super-resolve frequencies present in an analysis frame. The frequency estimation system is able to super-resolve frequencies by a factor of 3.5 compared to Fourier-based systems on synthetic signals, and is robust to amplitude modulations. Future work should be directed toward improving amplitude-phase estimation and constructing trajectory continuation systems based on temporal models, as current temporal sinusoidal modeling techniques perform poorly over note transitions in comparison to spectral sinusoidal modeling techniques.

9. REFERENCES

[1] T. Virtanen, "Audio signal modeling with sinusoids plus noise," Master's thesis, Tampere University of Technology, 2000.

[2] R. J. McAulay and T. F. Quatieri, "Speech analysis/synthesis based on a sinusoidal representation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 1986.

[3] M. Betser, P. Collen, B. David, and G. Richard, "Review and discussion on STFT-based frequency estimation methods," in *presented at the Audio Engineering Society 120th Convention*, 2006.

[4] J. O. Smith and X. Serra, "PARSHL: An analysis/synthesis program for non-harmonic sounds based on a sinusoidal representation," in *Proceedings of the International Computer Music Conference*, 1987.

[5] B. G. Quinn, "Estimating frequency by interpolation using Fourier coefficients," *IEEE Transactions on Signal Processing*, 1994.

[6] M. Abe and J. O. Smith III, "Design criteria for simple sinusoidal parameter estimation based on quadratic interpolation of FFT magnitude peaks," in *presented at the Audio Engineering Society 117th Convention*, 2004.

[7] F. Auger and P. Flandrin, "Improving the readability of time-frequency and time-scale representation by the reassignment method," *IEEE Transactions on Signal Processing*, 1995.

[8] M. Desainte-Catherine and S. Marchand, "High-precision Fourier analysis of sounds using signal derivatives," *Journal of Acoustic Engineering Society*, 2000.

[9] S. Marchand and P. Depalle, "Generalization of the derivative analysis method to non-stationary sinusoidal modeling," in *Proceedings of the International Conference on Digital Audio Effects*, 2008.

[10] M. Betser, "Sinusoidal polynomial parameter estimation using the distribution derivative," *IEEE Transactions on Signal Processing*, 2009.

[11] P. Depalle and T. Hslie, "Extraction of spectral peak parameters using a short-time Fourier transform and no sidelobe windows," in *IEEE 1997 Workshop on Applications of Signal Processing to Audio and Acoustics*, 1997.

[12] G. Dantzig, Ed., *Linear Programming and Extensions*. Princeton University Press, 1963.

[13] M. Lagrange, S. Marchand, and J.-B. Rault, "Enhancing the tracking of partials for the sinusoidal modeling of polyphonic sounds," *IEEE Transactions on Audio, Speech, and Language Processing*, 2007.