Spectral Warping Revisited

Donald Bailey, Warwick Allen
Institute of Information Sciences and Technology
Massey University, Private Bag 11222,
Palmerston North, New Zealand
D.G.Bailey@massey.ac.nz, wallen@iee.org

Serge Demidenko
School of Engineering and Science,
Monash University Malaysia,
Kuala Lumpur, Malaysia
Serge.Demidenko@engsci.monash.edu.my

Abstract

Spectral warping is a time domain to time domain transformation on a signal that effectively warps the frequency content of the original signal. Here we present a matrix formulation of the spectral warping transformation. The transform matrix is decomposed into three steps. The first is a DFT to convert the time signal into the frequency domain. Step two is an interpolation matrix to calculate the signal content at the desired new frequency samples. This effectively provides the frequency warping. The final step is an inverse DFT to transform the signal back into the time domain.

A direct consequence of this matrix representation is a direct FIR implementation of spectral warping, rather than the more commonly used IIR technique. We demonstrate that spectral warping is a generalisation of linear filtering, and show how the conventional all-pass spectral warping transformation can be generalised by using either arbitrary frequency mapping functions or different interpolation schemes. Finally, the conditions for the invertibility of the spectral warping transformation are derived.

1 Introduction

The importance of analogue and mixed-signal testing has grown increasingly during the last decades. Many specialists in different countries have carried out intensive research and development on the topic. This has resulted in many different efficient techniques and approaches. Digital Signal Processing (DSP)-based testing has proved to be one of the most promising among them. It involves employing digital tools and methods to test both digital and analogue components of the Device Under Test (DUT).

In the basic arrangement of DSP-based testing both signal generation and output measurement are realised by means of pure digital circuitry. Test signals that are usually used in this scheme include: digitised sinusoid, digitised multi-tone, pseudorandom, etc., while the following algorithms are traditionally used for response analysis: Discrete Fourier Transform (DFT) (normally computed by means of Fast Fourier Transform (FFT)), sine-wave fitting, cross-correlation, auto-correlation, filtering, etc.

DSP-based testing is an effective way around serious limits of pure analogue instrumentation (cross-talk, non-linearity, noise, drift, aging, improper calibration, long filter settling time, thermal effects and so on). At the same time, DSP-based testing provides benefits inherent to the use of digital components and tools (accuracy, stability, single set-up for multiple types of tests, repeatability of results, etc.). Besides that, in DSP-based testing all test access can be obtained through the same digital I/O ports, while the DUT can be tested for many parameters in one run thus increasing throughput.

Often in analogue and mixed signal devices, the frequency domain characteristics or transfer function of a device under test are of interest. Usually the focus is on a particular region of the frequency spectrum rather than the complete spectrum, so improved resolution is desired in the region of interest. Conventional analysis using an FFT provides equal resolution from DC up to the Nyquist frequency. However, many of the samples produced by the FFT are not required, and more detail is often desired around the frequencies of interest. Improving the resolution requires taking a longer FFT so that the samples are more closely spaced in the frequency domain.

An alternative approach that has been investigated to solve some of these problems is spectral warping [1-5]. The spectral warping transform modifies the signal from the device under test in such a way that the samples provided by an FFT of the warped signal correspond to unequally spaced samples of the Fourier Transform of the original signal. Thus, by warping the signal prior to taking the FFT, unequal frequency resolution is achieved. This allows a particular region of the spectrum to be analysed with higher resolution (more closely spaced samples) without having to increase the size of the FFT.

The spectral warping transform is a time domain transformation of a signal in that both the input and

Proceedings of the Second IEEE International Workshop on Electronic Design, Test and Applications (DELTA 2004)
0-7695-2081-2/04 $20.00 © 2004 **IEEE**

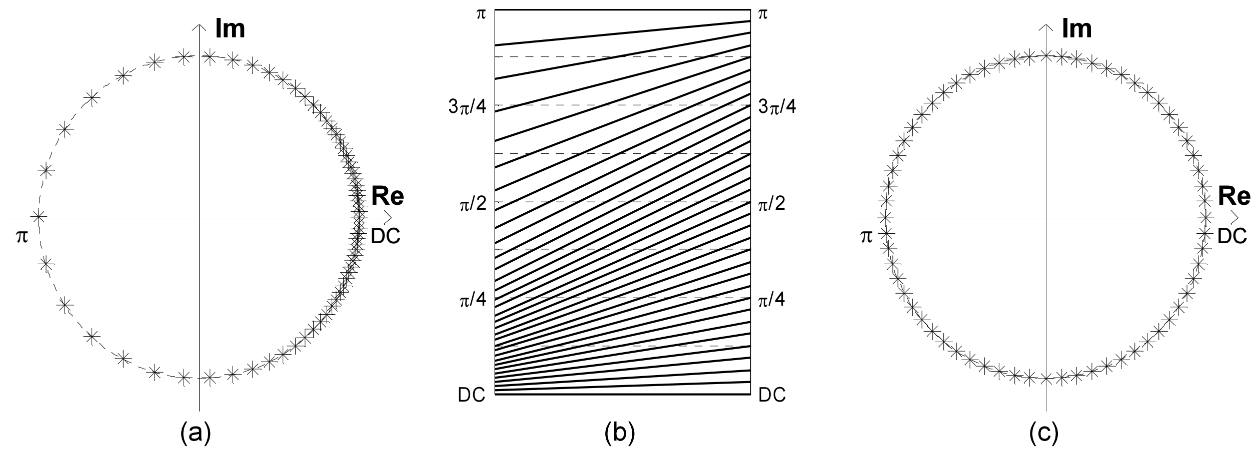

Figure 1. A graphical representation of the steps involved in spectral warping. The unevenly distributed *z*-domain samples (a) are redistributed by the spectral warping function (b) to be evenly spaced around the unit circle (c). The inverse *z*-transform now reduces to an inverse DFT.

output are time domain signals. It has been realised through a cascade of first order IIR filter sections [1,3,5]. The continuous signal to be analysed is sampled, and split into a series of frames of N samples each. The spectral warping transform is applied to each frame by time-reversing the samples within a frame and passing them into the filter network. After final sample has been entered, the outputs of each of the first order filter stages provide the samples of the warped signal. Because of the frequency warping, there will generally be more output samples than input [3,5].

The rest of this paper investigates a matrix formulation of the spectral warping transform. While such a formulation may not necessarily be practical from an implementation point of view, it can aid the understanding of spectral warping, and can provide further insight into its properties.

2 Analysis of the Transform

Conceptually, the spectral warping transform is derived by taking non-uniformly spaced frequency samples, and warping them to make them uniformly spaced [5]. The frequency axis corresponds to the unit circle within the *z*-domain, therefore any warping that maps the unit circle onto itself in the *z*-domain may be used for spectral warping. An all-pass mapping satisfies this relationship [5,6]. If the mapping is one-to-one, then it will also be invertible.

The spectral warping process (illustrated in figure 1) is equivalent to evaluating the *z*-transform of the input signal at the non-uniform sample points around the unit circle, and taking the inverse discrete Fourier transform of the result. The spectral warping takes place when the non-uniform samples of the *z*-transform are treated as uniformly spaced by taking the inverse Fourier transform. This process may be represented mathematically as follows.

2.1 Matrix representation

Let $f[n]$ be the N samples within one frame of the input signal, then its *z*-transform is given by:

$$F(z)=\sum_{n=0}^{N-1} f[n]z^{-n} \tag{1}$$

This is evaluated at M points around the unit circle, $z=e^{j\omega_m}$, where

$$\omega_m=\tan^{-1}\left[\frac{(1-a^2)\sin\frac{2\pi m}{M}}{(1+a^2)\cos\frac{2\pi m}{M}+2a}\right] \tag{2}$$

for a first order all-pass mapping with warp parameter a [5]. The summation of equation (1) may be conveniently represented in matrix form:

$$\begin{bmatrix} G[0] \\ G[1] \\ \vdots \\ G[m-1] \end{bmatrix}=\begin{bmatrix} 1 & e^{-j\omega_0} & \cdots & e^{-(n-1)j\omega_0} \\ 1 & e^{-j\omega_1} & \cdots & e^{-(n-1)j\omega_1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & e^{-j\omega_{m-1}} & \cdots & e^{-(n-1)j\omega_{m-1}} \end{bmatrix}\begin{bmatrix} f[0] \\ f[1] \\ \vdots \\ f[n-1] \end{bmatrix} \tag{3}$$

or $\mathbf{G}=\mathbf{H}_{M,N}\mathbf{f}$, where $\mathbf{f}$ are the N samples arranged as a vector, $\mathbf{H}_{M,N}$ is the MxN *z*-transform matrix representing equation (1), and $\mathbf{G}$ are the samples in the *z*-domain at the points indicated by equation (2). Each row of $\mathbf{H}$ therefore gives one frequency sample in the *z*-domain.

These samples are then warped to make them evenly spaced around the unit circle. The frequency warping is effectively the inverse of equation (2) and is illustrated in figure 1(b). This spectral warping will stretch one portion of the spectrum (low frequencies in figure 1), while compressing another portion (high frequencies in figure 1).

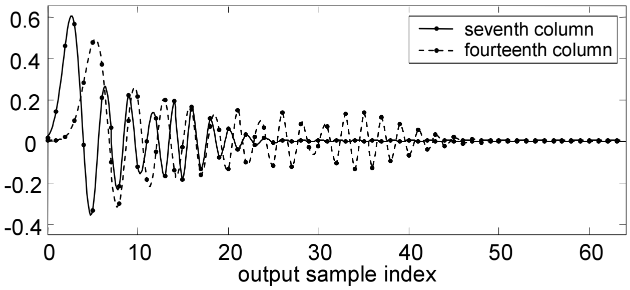

Figure 2. The seventh and fourteenth columns of a 64×16 spectral warping matrix, S, representing the impulse responses of the corresponding input samples, for a warp parameter of 0.5.

Since the samples are evenly spaced around the unit circle, the inverse z-transform corresponds to an inverse discrete Fourier transform:

$$g[n] = \tfrac{1}{M}\sum_{m=0}^{M-1} G[m]W^{-mn} = \tfrac{1}{M}\sum_{m=0}^{M-1} F(e^{j\omega_m})W^{-mn} \tag{4}$$

where $W_M^k = e^{-j2\pi k/M}$. Representing this in matrix form

$$\begin{bmatrix} g[0] \\ g[1] \\ \vdots \\ g[m-1] \end{bmatrix} = \tfrac{1}{M}\begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & W_M^{-1} & \cdots & W_M^{-(m-1)} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & W_M^{-(m-1)} & \cdots & W_M^{-(m-1)(m-1)} \end{bmatrix}\begin{bmatrix} G[0] \\ G[1] \\ \vdots \\ G[n-1] \end{bmatrix} \tag{5}$$

or $\mathbf{g} = \mathbf{D}_M^{-1}\mathbf{G}$, where $\mathbf{D}_M$ is an M-point DFT matrix. Substituting equation (3) into (5) gives:

$$\mathbf{g} = \mathbf{D}_M^{-1}\mathbf{H}_{M,N}\mathbf{f} = \mathbf{S}_{M,N}\mathbf{f} \tag{6}$$

From equation (6) it is clearly seen that the spectral warping transformation, **S**, is linear in that each of the M output samples is a linear combination of the N input samples. Each column of **S** therefore represents the output produced by an impulse at the corresponding input sample. As the columns of the warping transformation are quite different and not related by simple delays, the warping is not time invariant. Therefore a time shift of the input signal will not correspond to a simple shift in the warped output. Figure 2 shows some of the impulse responses associated with the transformation shown in figure 1.

2.2 Frequency interpolation

Matrix **H** can be further decomposed. First consider the N-point DFT of the input sequence. This gives N equally spaced samples around the unit circle in the z-domain. The M points produced by **H** are effectively interpolated between these. That is

$$\mathbf{g} = \mathbf{D}_M^{-1}\mathbf{H}\mathbf{f} = \mathbf{D}_M^{-1}\mathbf{H}\mathbf{D}_N^{-1}\mathbf{D}_N\mathbf{f} = \mathbf{D}_M^{-1}\mathbf{C}\mathbf{D}_N\mathbf{f} \tag{7}$$

where

$$\mathbf{C} = \mathbf{H}\mathbf{D}_N^{-1} \tag{8}$$

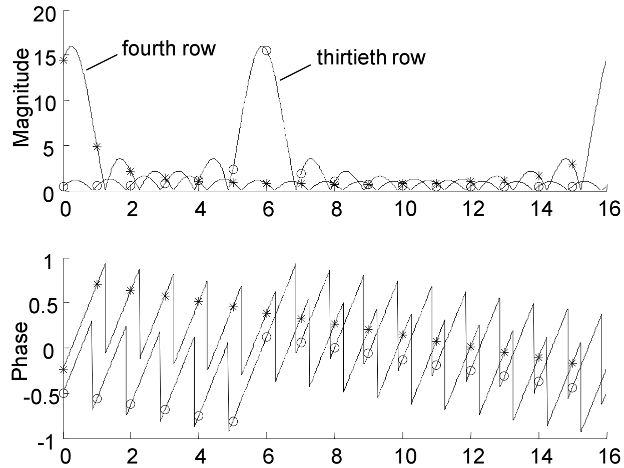

Figure 3. Two example rows of a 64×16 interpolation matrix, C, showing the modified sinc function. The markers show the points at which the function is evaluated.

is effectively an MxN interpolation matrix. **C** therefore directly represents the frequency domain warping of equation (2). Each row of **C** gives one of the warped frequency samples as a linear combination of the uniformly spaced samples of the standard DFT. The interpolation functions are expected to be modified sinc functions – modified because sampled data is periodic in the frequency domain:

$$\left|C(n)\right| = \tfrac{1}{N}\frac{\sin \pi n}{\sin \frac{\pi n}{N}} \tag{9}$$

When these are offset to the desired sample positions and sampled, the interpolation matrix elements are derived from evaluating equation (8):

$$\begin{aligned} C_{MxN}[m,n] &= \tfrac{1}{N}\sum_{k=0}^{N-1} e^{-jk\omega_m} e^{j2\pi kn/N} \\ &= \tfrac{1}{N}\sum_{k=0}^{N-1} e^{jk\left(\frac{2\pi n}{N}-\omega_m\right)} \\ &= \tfrac{1}{N}\frac{1-e^{j(2\pi n - N\omega_m)}}{1-e^{j\left(\frac{2\pi n}{N}-\omega_m\right)}} \\ &= \tfrac{1}{N}\frac{e^{j\left(\pi n - \frac{N\omega_m}{2}\right)}}{e^{j\left(\frac{\pi n}{N}-\frac{\omega_m}{2}\right)}}\left(\frac{e^{-j\left(\pi n - \frac{N\omega_m}{2}\right)} - e^{j\left(\pi n - \frac{N\omega_m}{2}\right)}}{e^{-j\left(\frac{\pi n}{N}-\frac{\omega_m}{2}\right)} - e^{j\left(\frac{\pi n}{N}-\frac{\omega_m}{2}\right)}}\right) \\ &= \tfrac{1}{N}\frac{\sin\left(\pi n - \frac{N\omega_m}{2}\right)}{\sin\left(\frac{\pi n}{N}-\frac{\omega_m}{2}\right)} e^{j(N-1)\left(\frac{\pi n}{N}-\frac{\omega_m}{2}\right)} \end{aligned} \tag{10}$$

Figure 3 shows the magnitude and phase of two rows of the **C** matrix for the above example. The phase term in equation (10) has an interesting effect on the phase of the interpolation samples. The (N-1) factor almost completely compensates for the phase reversal between adjacent

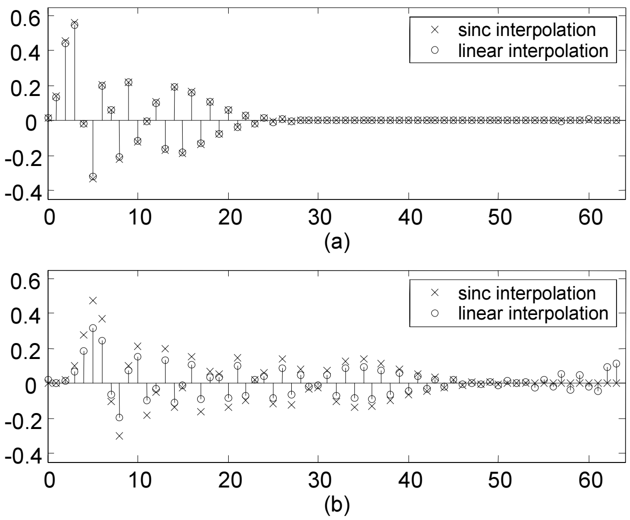

Figure 4. Comparing the impulse responses (columns) between the sinc and linear spectral warping matrices. The seventh (a) and fourteenth (b) columns of the 64×16 S matrix are shown.

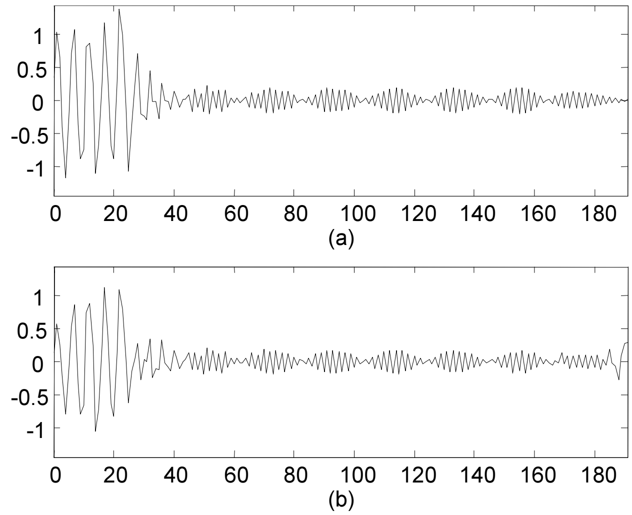

Figure 5. Sinc (a) and linear (b) spectral warping applied to the signal (sin π/7 + sin 6π/7)/2.

nodes of the sinc interpolation, resulting in adjacent coefficients having similar phase.

Although the coefficients of $\mathbf{D}_M$, $\mathbf{D}_N$ and $\mathbf{C}$ are complex in general, the result of their product in **S** is real. This because of the conjugate symmetry involved with each of the steps. Since the sampled data is real, the initial DFT of that sampled data will be conjugate symmetric. If the warping function is symmetric for positive and negative frequencies, then after the interpolation by **C** the warped frequency components will also be conjugate symmetric. Finally the inverse DFT will transform the conjugate symmetric coefficients back into a real sequence.

2.3 FIR implementation

Multiplying the input samples by the matrix **S** effectively represents a bank of real coefficient FIR filters with each row of **S** making up a separate filter in the bank. When the complete frame has been input to the filter (*N* samples), the output of each filter provides a separate sample of the warped waveform. The FIR filter bank therefore provides an alternative to the IIR implementation. Such an FIR implementation wouldn't have the transient components of the IIR implementation from not initialising the memory. Also the FIR filters would be less susceptible to word length effects because they can be implemented in direct form. The IIR filters being recursive will accumulate the error with each filter stage that the signal passes through. The biggest disadvantage of the FIR filters is that all of the coefficients are different, whereas almost all of the filters of the IIR implementation are identical. To change the warp factor, all MxN coefficients must be recalculated, whereas with the IIR filter, only a single parameter is affected.

2.4 Other interpolation functions

The spectral warping provided by equation (2) can be generalised by modifying the interpolation matrix, **C**. One generalisation would be to use something other than an all-pass transformation as the warp function, for example a piecewise linear function. The samples in the frequency domain may be placed arbitrarily around the unit circle, to give an arbitrary warp function. Then, the appropriate interpolation coefficients may be found from equation (8) or (10).

A second generalisation would be to use a different interpolation to that provided by equation (10). For example, a linear interpolation could be used between the nearest two frequency samples (the phase term of equation (10) must also be used for best results). This would result in approximate frequency estimates, and may be adequate in many circumstances. Figure 4 compares the impulse responses generated by using linear interpolation with those for sinc interpolation. For input samples near the centre of the frame, the coefficients are almost identical, whereas those near the start and end of the frame diverge more from ideal. The overall results of this effect may be seen in figure 5. The differences between the two different interpolation functions may be significantly reduced by windowing the input with a Hanning window prior to warping because this reduces the weight of the samples at the start and end of the frame.

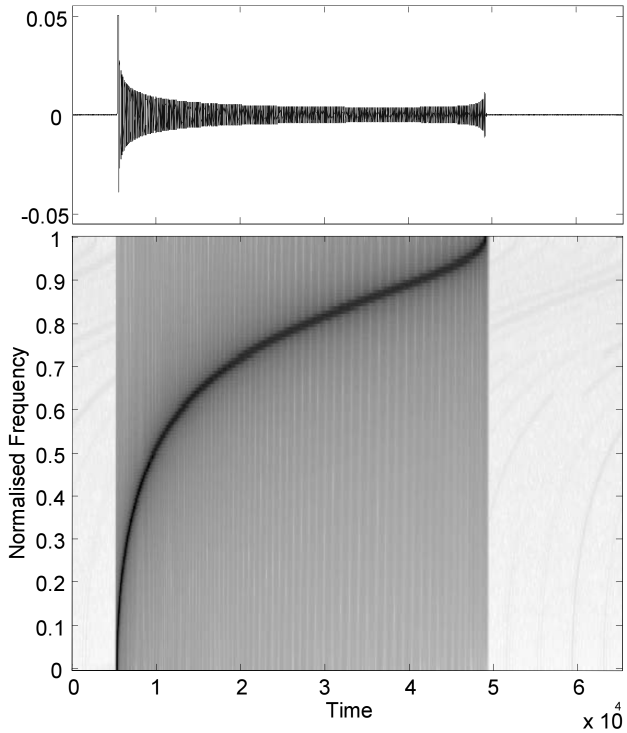

Figure 6. Time domain plot and spectrogram of the result of spectrally warping a unit impulse occurring at the *M*/4th sample (where *M* is the output sequence length). A warp parameter of *a* = 0.5 compresses the low frequencies in time and stretches the high frequencies. All the frequencies are warped towards the Nyquist frequency (as illustrated in figure 1).

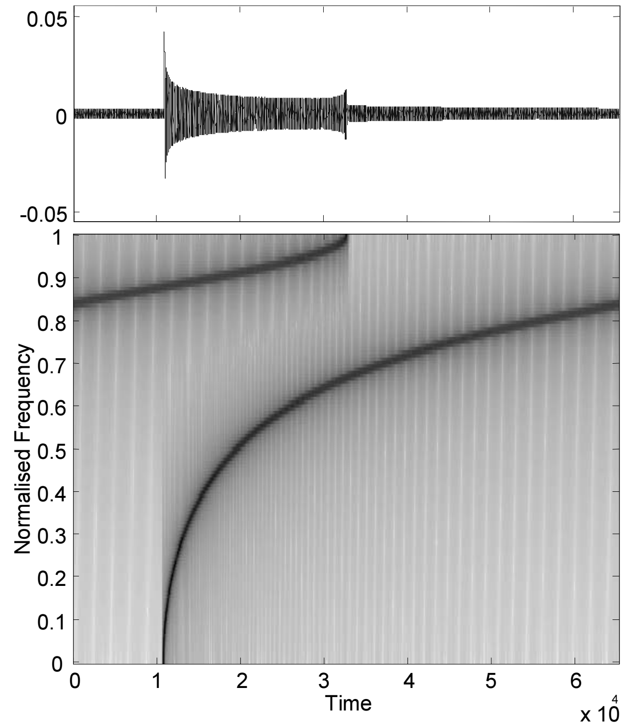

Figure 7. The result of spectrally warping a unit impulse using too few output samples. The plot and spectrogram above show the warped output of a unit impulse occurring at the *M*/2th sample, with the sequence length and warp parameter identical to those of figure 6. The wrap-around effect of temporal aliasing is clearly visible.

One special case interpolation function is worth noting. If **C** is square and diagonal, equation (7) corresponds to a standard linear filter. The input samples are transformed into the frequency domain, where they are weighted by **C** and transformed back into the time domain. The values along the diagonal of **C** would in that case represent the frequency response of the filter. Therefore, spectral warping can be seen as a generalisation of linear filtering.

2.5 Invertibility

For the spectral warping to be invertible, it is necessary for **C**, or equivalently **H**, to be invertible. However, in general **C** (and **H**) are not square, so a pseudo-inverse must exist such that

$$\mathbf{C}^{-1}\mathbf{C} = \mathbf{I}_N \tag{11}$$

where $\mathbf{I}_N$ is the NxN identity matrix. Equation (11) implies that **C** must be of rank N for such an inverse to exist. It can be shown that **H** is invertible provided there are at least N distinct warped frequencies [7]. The implication therefore is that any additional frequency samples are actually redundant. While from a mathematical viewpoint this may be true, the output signal **g** in that case has little physical meaning. It has been shown [3-5] that to completely represent the warped output signal in the time domain, generally more output samples than input samples are required. This should not be a surprise, because a similar effect occurs when implementing convolution via the frequency domain [6]. The input signal must be zero-padded prior to filtering to prevent aliasing in the time domain.

With spectral warping, the warp function stretches or spreads some frequencies in the frequency domain and compresses others. Any such frequency scaling will result in a corresponding inverse time scaling, so where the frequencies are compressed, the corresponding frequency components will become stretched in the output time waveform. This time–frequency relationship can be clearly seen in figure 6, where a spectrogram is taken of a warped unit impulse.

The effect of producing fewer time domain samples is that the signal at those frequencies will become aliased in the time domain (see figure 7). While technically such aliasing can be recovered (since **H** is invertible) the time domain signal has little meaning in this context.

When **C** is not square, the inverse is not unique because only N of the M frequency sample points are linearly independent. When an all-pass mapping is used for the spectral warping transform, the inverse transform is also given by equation (2) except using a warp parameter of $-a$. Since the transformation with a increases the number of samples from N to M, the inverse transformation will also have the same time stretching ratio. However, the time stretching of the inverse transform occurs at frequencies that were shrunk in time by the forward transform. Since there are N input samples, provided that M is sufficiently long to represent the entire transformed sequence without aliasing, then the inverse transform will provide only N non-zero output samples.

3 Summary

This paper has formulated a matrix representation of the spectral warping transform. This effectively decomposes the transformation into 3 components: a DFT from the original time domain samples into the frequency domain; a warping matrix that interpolates the available frequency samples to give a new set of frequency samples; and an inverse DFT to transform these back into the time domain. The transformation matrix effectively represents a bank of FIR filters that can be used to implement the warping. From the matrix representation, it is shown that the spectral warping transform is not time invariant, although the impulse responses are readily available from the transformation matrix.

Spectral warping can be seen as a generalisation of linear filtering, and warping by an all-pass function can also be generalised by using interpolation functions other than sincs, or by applying an arbitrary frequency mapping. For the transformation to be invertible, the original frequency samples must be able to be derived from the new warped samples.

4 References:

[1] A. V. Oppenheim, and D. H. Johnson, "Discrete Representation of Signals," *Proc. IEEE*, vol. 60, June 1972, pp. 681-691.

[2] S. N. Demidenko, and K. Lever, "Analysis of the Frequency Characteristics of a Digital Spectral Warping Network," in *International Workshop on Design Methodologies for Microelectronics and Signal Processing,* Gliwice-Cracow, Poland, 1993, pp. 397-401.

[3] W. Allen, D. Bailey, S. Demidenko, and V. Piuri, "Analysis and Implementation of a Spectral Warping Network," in *Proc. International Symposium on Integrated Circuits,* 2001, pp. 492-496.

[4] W. Allen, D. Bailey, S. Demidenko, and V Piuri, "Test chirp Signal Generation Using Spectral Warping," *International Workshop on Electronic Design, Test, and Applications,* Christchurch, January 29-31, 2002, pp 492-495.

[5] W. Allen, D. Bailey, S. Demidenko, and V Piuri, "Analysis and Application of Digital Spectral Warping in Analog and Mixed-Signal Testing," accepted for publication in *IEEE Transactions on Instrumentation and Measurement*.

[6] S. K. Mitra, Digital Signal Processing: a Computer-Based Approach. McGraw-Hill, Singapore, 1998.

[7] S. Bagchi, and S. K. Mitra, The Nonuniform Discrete Fourier Transform and its Application in Signal Processing, Kluwer Academic Publishers, Boston, 1999.