IEEE SP Paper # 5820
Accepted for publication in IEEE Trans. on Sig. Proc. (First draft submitted 1991)

The Chirplet Transform: Physical Considerations*

Steve Mann† and Simon Haykin‡

Abstract

We consider a multidimensional parameter space formed by inner products of a parameterizable family of *chirp* functions with a signal under analysis. We propose the use of quadratic chirp functions (which we will call q-chirps for short), giving rise to a parameter space that includes both the time-frequency plane and the time-scale plane as two-dimensional subspaces. The parameter space contains a "time-frequency-scale volume", and thus encompasses both the short-time Fourier transform (as a slice along the time and frequency axes), and the wavelet transform (as a slice along the time and scale axes).

In addition to time, frequency, and scale, there are two other coordinate axes within this transform space: shear-in-time (obtained through convolution with a q-chirp) and shear-in-frequency (obtained through multiplication by a q-chirp). Signals in this multidimensional space can be obtained by a new transform which we call the "q-chirplet transform", or simply the "chiplet transform".

The proposed chirplets are generalizations of wavelets, related to each other by two-dimensional affine coordinate transformations (translations, dilations, rotations, and shears) in the time-frequency plane, as opposed to wavelets which are related to each other by one-dimensional affine coordinate transformations (translations and dilations) in the time-domain only.

1 INTRODUCTION

Underlying a great deal of traditional signal processing theory is the notion of a *sinusoidal wave*. With the advent of modern computing, and the Fast Fourier transform, the use of and interest in frequency-domain signal processing has increased dramatically. More recently, however, researchers are becoming aware of the limitations of frequency-domain methods. Although the Fourier transform yields perfect reconstruction of a broad class of signals, it does not necessarily provide a meaningful interpretation when the signals lack global stationarity. For example, consider the time series formed by a typical passage of music. An estimate of its power spectrum tells us which musical notes are present (how much energy there is around each of the frequencies), but fails to tell us when each of those notes was sounded.

Much of the recent focus of signal processing is on the so-called *time-frequency* (TF) methods, which allow us to observe how a spectral estimate evolves over time. One of these TF methods, the short-time Fourier transform (STFT), has been used extensively for analyzing speech, music, and other non-stationary signals.

Suppose we want to perform a STFT analysis, but are uncertain what the window size should be. We could perform the STFT of a signal, $s(t)$, using a window of relatively short duration, then stretch the window out a small amount and compute another STFT, and so on, gradually increasing the window size and computing another STFT for each value of window size. Stacking uncountably many of these STFTs on top of one another results in a continuous volumetric representation of s that is a function of time, frequency, and the size of the window (Fig. 1(a)). We will refer to this volumetric representation as the *time-frequency-scale* (TFS) transform[1].

Another time-frequency representation (which might more appropriately be called a time-scale representation) is the well-known wavelet transform [2], [3], [4], [5]. The wavelet transform can be expressed as an inner product of the signal under analysis with a family of translates and dilates of one basic primitive. This primitive is known as the *mother wavelet*. A member of the wavelet family is produced by a particular one-dimensional affine coordinate transformation acting on the time axis of the mother wavelet; this geometric transformation is parameterized by two numbers (corresponding to the amounts of translation and dilation). The continuous wavelet transform is formed by taking inner products of the signal with the uncount-

*This work was supported by the Natural Sciences and Engineering research Council of Canada

†Massachusetts Institute of Technology Room E15-389, 20 Ames Street, Cambridge, MA 02139 e-mail steve@media.mit.edu

‡Communications Research Laboratory, McMaster University Hamilton, Ontario, Canada L8S 4K1. e-mail haykin@synapse.crl.mcmaster.ca

[1] When using a multidimensional parameter space, it is often impossible to establish *frame bounds* [1] on the energy in the parameter space. With only one parameter, we cannot always reconstruct the signal. With two *effective* parameters, we can reconstruct the signal, and also bound the energy of the representation. With three or more parameters, the energy in the transform space will be infinite. To the extent that multi-dimensional parameter spaces are still useful, we will not let this infinite energy hinder our progress.

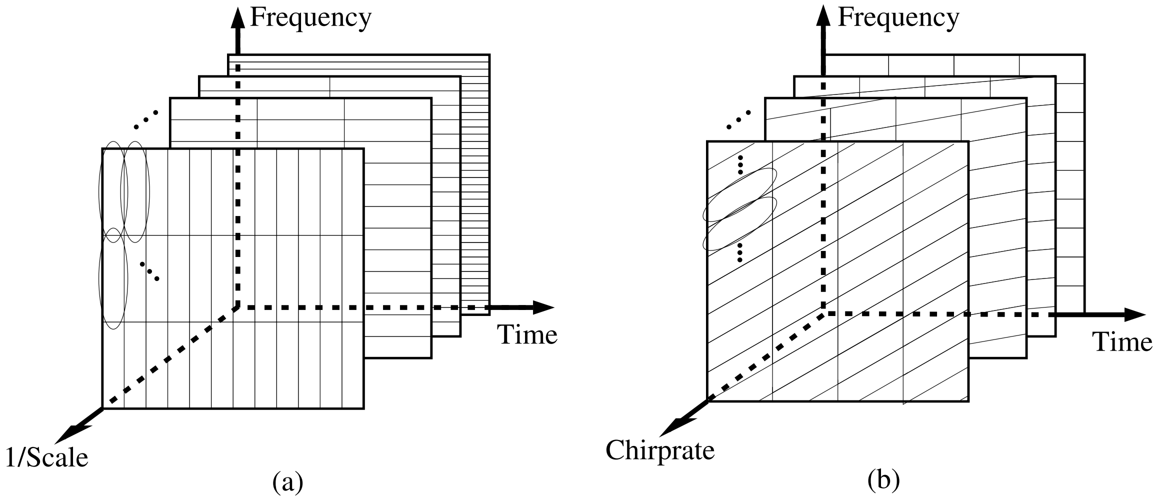

Figure 1: Volumetric family of short-time Fourier transforms. (a) A family of uncountably many STFTs where the window is allowed to dilate continuously gives us a "time-frequency-scale" (TFS) transform. The bottom plane, $f_c = 0$, is the time-scale plane which is a continuous wavelet transform if $g \in L^2(\mathbb{R})$ is a suitably chosen *mother wavelet*. Here we only show one octant of the volume. Note also that the plane $1/s = 0$ is not defined, for it would correspond to infinite scale. (b) Sheared STFTs with a variety of assumed chirprates. Shearing of the TF plane is performed through multiplication of the signal by a chirp, with chirprate c. If we stack up uncountably many such TF planes, allowing c to vary continuously, the result is a "time-frequency-chirprate" (TFC) transform.

ably many members of the two-parameter wavelet family. The continuous wavelet transform is, with an appropriate choice of window/mother-wavelet, simply the time-scale (TS) plane of the TFS volume (Fig. 1(a)).

We begin to see that, even if it is not practical from a computational or data-storage point of view, the *time-frequency-scale* space is useful from a conceptual point of view. In particular, if we only desire the magnitude TFS volume, we can easily extract this information from the Wigner distribution, by the appropriate coordinate transformations and uniform smoothing of the coordinate-transformed Wigner distributions. A continuous transition from the magnitude TF plane (spectrogram) to the magnitude TS plane (scalogram) is possible through appropriate smoothing of the Wigner distribution [6].

Now suppose we were to multiply the signal, $s(t)$, by a linear FM (chirp) signal $\exp(j2\pi\frac{c}{2}t^2)$ and then compute its STFT. If we vary the chirp rate, c, continuously, and repeat the process uncountably many times, stacking the resulting STFTs one above the other, we obtain a different three-dimensional volume (Fig. 1(b)). This time we have a function of time, frequency and chirprate.

Of course there is no reason to limit ourselves to a choice between these two parameter spaces; to motivate what follows, it will prove helpful to keep in mind a continuous four-dimensional "time-frequency-scale-chirprate[2]" (TFSC) parameter space.

[2] Traditionally, the term *chirp-rate* (with a hyphen) is used, but in this paper, we use the single word "chirprate", to avoid confusion arising out of hyphens in compounded parameter lists.

1.1 Historical Notes

In 1946, in his seminal paper on communication theory [7], Gabor (who later won the Nobel prize for his work on holography), provided a new interpretation of the one-dimensional Gaussian-windowed STFT and examined the time-frequency plane in terms of a two-dimensional tiling. Although Gabor's development was not completely rigorous (and, in fact his representation was later shown to be unstable [1]), his notion of a time-frequency tiling was a very significant contribution. Gabor referred to the elements of his tiling as *logons*.

Beginning around 1956, Siebert began to formulate a radar detection philosophy with some particularly useful insights in terms of time-frequency [8][9]. Much of his insight was obtained through the use of *Woodward's uncertainty function* [10], also known as the *radar ambiguity function* [11] or the *Fourier-Wigner transform* [12]. Siebert also considered chirping functions for pulse compression radar, and studied these in detail, observing that chirping in the time domain gives rise to a shearing in the time-frequency plane (or equivalently, a shearing in the 2D Fourier transform of the time-frequency plane).

In 1985, Grossman and Paul [13] rigorously formulated some of these important ideas in terms of affine cannonical coordinate transformations to a coherent space representation. They also considered two-parameter subgroups of these affine coordinate transformations.

Papoulis, in his book[14] described the use of a linear frequency-modulated (chirped) signal as the basis of an ordinary Fourier analyzer, and also presented the chirped signals as shearing operators in the time-frequency plane, foreshadowing the development of the chirplet transform.

In 1987, Jones and Parks [15] formulated the problem of window selection in terms of time-frequency leakage. They made an important connection between the work of Szu and Blodgett [16] who showed that frequency shearing is accomplished through multiplication by a chirp, and the work of Janssen [17] who proved that any area-preserving affine coordinate transformation of the time-frequency plane yields a valid time-frequency plane of some other signal, though they were unaware of Siebert's earlier unpublished work. In a simple and insightful example, Jones and Parks showed the time-frequency distribution of both a Hamming window and a chirped Hamming window, one being a sheared version of the other.

Berthon [18] proposed a generalization of the radar ambiguity function based on the semi-direct product of two important groups:

- the special linear group, SL(2,$\mathbb{R}$) that embodies shear in the time-frequency plane, and
- the Heisenberg group that involves both time and frequency shifts.

In 1989 and early 1990, we formulated the chirplet transform, a multidimensional parameter space whose coordinate axes correspond to the pure parameters of planar affine coordinate transformations in the time-frequency plane. (This formulation was motivated by a discovery made by the senior author and his research associates, namely, that the Doppler radar return from a small piece of ice floating in an ocean environment is chirp-like[19].) We also formulated a variety of new and useful transforms that were two-dimensional subspaces of this multidimensional parameter space. Furthermore, we suggested using the work of Landau [20][21][22][23][24][25] who introduced prolate spheroidal functions, and we noted their significance in the context of the shearing phenomenon in the time-frequency plane, as they form idealized parallelogram tilings of this plane.

Later, we applied the chirplet transform and some of the new two-dimensional subspace transforms to problems in marine radar and obtained results that were better than previous methods, so we published these findings [26]. Independently, at around the same time (ironically, only a few days later) Mihovilovic and Bracewell also presented a related idea [27] (ironically, using the same name, "chirplets"), though not in the same level of generality of the multidimensional parameter space. Later they also presented a practical application of chirplets [28].

A point that needs to be emphasized here is that there is more to the chirplet transform than just the shear phenomenon. In particular, time-shear and frequency-shear are examples of *affine* coordinate transformations – mappings from the TF-plane to the TF-plane – while the chirplet transform is a mapping from a continuous function of one real variable to a continuous function of five (or six) real variables.

In 1991, Torresani [29] examined some relations that were intermediate between the affine and the Weyl-Heisenberg groups. The work of Segman and Schempp [30] incorporates scale into the Heisenberg group, and the work of Wilson *et al.* [31][32] examines the use of a TFS representation that they call the multiresolution Fourier transform.

Baraniuk and Jones studied several "chirplet transform subspaces" and made precise some of the mathematical details of the two-dimensional chirplet transform subspaces [33]. They also provided an alternative derivation [33] of the chirplet transform, based on the Wigner distribution. This derivation involved noting, as we did, that each point in the analysis space of the chirplet transform corresponds to a particular operator in the time domain. This time-domain operator acting on the analysis primitive ('mother chirplet') also has, associated with it, a 2-D area-preserving affine coordinate transformation in the TF plane. Baraniuk and Jones also addressed discretization issues [33][34].

Recently, researchers have considered fractional Fourier domains and their relation to chirp and wavelet transforms[35].

1.2 RELATED WORK

Early on, our interest in chirping analysis functions was motivated by a different kind of chirping phenomenon: chirping due to perspective. Our urban or indoor world contains a plethora of periodicity, repeating rows of bricks, tiles, windows, or the like abound, yet pictures of these structures fail to capture the true essence of this periodicity. When photographed at an oblique angle (where the film plane is not necessarily parallel to the planar surface), these surfaces give rise to an image whose spatial frequency changes as we move across the image plane. The distant bricks will appear smaller and smaller as we move toward the *vanishing point* which may be defined to be the point of infinite spatial frequency. Our first generalization of the wavelet transform was to take the "zooming-in" property of wavelets and extend it to *panning* and *tilting*, to model the movements of a camera. Our interest in radar, however, drew us toward processes that are more accurately analyzed by linear-FM chirplets. We realized that, listening to radar sounds from marine radar, automobile traffic radar, and the like, that in many cases there was a strong "chirping", and so the usual Fourier Doppler methods seemed inappropriate in these cases. In particular, the warbling sound of small iceberg fragments suggested that we should consider alternatives to windowed harmonic oscillations and the like (e.g. alternatives to waves and wavelets).

Of the many different kinds of chirping analysis primitives possible, we may distinguish two families of analysis primitives that are of particular interest in practice: the "projective chirplet" (p-chirplet), and the "quadratic chirplet" (q-chirplet), the latter being the one described in this paper. These two forms have been presented in a combined fashion with the "time-frequency perspectives" [36], which is a more general chirplet that has eight parameters. The resulting eight-parameter signal representation includes the "projective chirplet transform" as one five-parameter subspace, and the "quadratic chirplet transform" (e.g. the one presented in this paper) as another five-parameter subspace with the time, frequency, and dilation axes being common to both of these two subspaces. Computational issues have yet to be addressed, although special-purpose hardware has been proposed [37] with an emphasis on use of FFT-based hardware.

We have also constructed other chirplet transforms, such as a three-parameter Doppler chirplet representation that models a source producing a sinusoidal wave, while moving along a straight line (e.g. a train whistle). The three parameters are center-frequency, maximum rate of change of frequency, and frequency swing. Also, a log-frequency chirplet has been formulated where the underlying chirps appear as straight lines in the time-scale plane.

Generalizations of the STFT and wavelet transform, that make use of chirping analyzing functions, have been previously suggested [26], [27], [38], [39], [40], [36], [41], [28]. Comparisons between traditional TF methods and chirplets have also been made, in the context of practical applications in both radar [26], [42], and geophysics [28].

1.3 Overview

This paper is devoted to physical (intuitive) considerations of the chirplet transform. It is organized as follows:

- We first introduce chirping analysis functions which may be thought of as generalized wavelets ("chirplets").
- We then generalize Gabor's use of the Gaussian window for his tiling of the time-frequency plane. This generalization gives rise to the four-dimensional time-frequency-scale-chirprate (TFSC) parameter space.
- We next consider non-Gaussian analysis functions, giving rise to a five-dimensional parameter space.
- We then consider the use of multiple analyzing wavelets/windows, first to generalize Thomson's method of spectral estimation to the TF plane, and then to further generalize this result to the chirplet transform. The multiple analyzing wavelets/windows (which we call "multiple mother chirplets" when they are used in the latter context) collectively act to define a single "tile" in the TF plane, corresponding to each point in the chirplet transform parameter space. Such a tile has a true parallelogram-shaped TF distribution whose shape is governed by the six 2-D affine parameters.
- We generalize autocorrelation and cross-correlation by using the signal itself (or another signal) as a "mother chirplet". In other words, we analyze the signal against chirped versions of itself (or against chirped versions of another signal).
- Finally, we consider chirplet transform subspaces, leading to a variety of new transforms.

2 THE CHIRPLET

The STFT consists of a correlation of the signal with constant-size portions of a wave, while the wavelet transform consists of correlations with a constant-Q family of functions. The two transforms, however, are in some ways similar. Although the former is generally thought of as a *time-frequency* method, and the latter, a *time-scale method*, both attempt to localize the signal in the time-frequency plane. In a rather loose sense, both the modulated window of the STFT, and the wavelet[3] of the wavelet transform, may be regarded as "portions of waves". Chirplets, in a similar manner, may be regarded as "portions of chirps". We generally use complex-valued chirplets to avoid the mirroring in the $f = 0$ axis that results from using only real-valued chirplets.

Figure 2 provides a comparison in terms of real and imaginary components as well as time-frequency distributions, between a wave, wavelet, chirp, and chirplet. In Fig. 3, we provide the same comparison with a 3-D *particle-rendering*, where the three coordinate axes are the function's real value, imaginary value, and time. Discrete samplings of four chirplets are shown: the top two have chirprate set to zero, and the leftmost two have an arbitrarily large window.

2.1 Gaussian Chirplet

The chirplets in Fig. 2 and Fig. 3 were derived from a single Gaussian window by applying simple mathematical operations to that window. The window may be thought of as the primitive that generates a family of chirplets, much like the *mother wavelet* of wavelet theory. We will, therefore refer to this primitive (whether Gaussian or otherwise[4]) as the "mother chirplet", and will denote it by the letter g.

A *Gaussian wave packet* (also known to physicists as simply a *wave packet*), is a wave with a Gaussian envelope. Mathematically, a wave packet, g, may be represented:

$$g'_{t_c,f_c,\sigma}(t) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}(\frac{t-t_c}{\sigma})^2} e^{j2\pi f_c(t-t_c)+j\phi} \qquad (1)$$

where $j = \sqrt{-1}$, $t_c \in \mathbb{R}$ is the center of the energy concentration in time, $f_c \in \mathbb{R}$ is the center frequency, $\sigma \in \mathbb{R} > 0$ is the spread of the pulse, and $\phi \in \mathbb{R}$ is the phase shift of the wave, which we will not consider as one of the parameters. The subscripts of g represent the degrees of freedom, which comprise the *parameter list*.

We like the wave packet to have unit energy. Hence we reformulate the definition of the Gaussian envelope (taking advantage of the fact that a Gaussian function raised to any exponent, in our case, $1/2$, is still a Gaussian function if multiplied by the appropriate normalization constant):

$$\begin{aligned} g'_{t_c,f_c,\sigma}(t) &= \sqrt{\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}(\frac{t-t_c}{\sigma})^2\right)} \exp(j2\pi f_c(t-t_c)) \\ &= \frac{1}{\sqrt{\sqrt{\pi}\Delta_t}} \exp\left(-\frac{1}{2}(\frac{t-t_c}{\Delta_t})^2\right) \exp(j2\pi f_c(t-t_c)) \end{aligned} \qquad (2)$$

where $\Delta_t = \sqrt{2}\sigma$.

Theoretically bandlimited signals have infinite duration, but it is customary, in electrical engineering, to use the *3dB bandwidth* which is defined as the difference in frequencies, on either side of the peak, where the energy or power falls to half the peak value. This definition, however, is not theoretically motivated, nor particularly useful in our context. Therefore, in the case of the wave packet, we simply define the duration to be equal to Δ_t in (2). By

[3] The term "wavelet" will appear in quotes when it is used in this less restrictive sense. In particular, a "wavelet" will be permitted to have a non-zero DC component.

[4] In general, $g(t)$ is a complex-valued function of a real variable, and has finite energy: $g \in L^2(\mathbb{R})$.

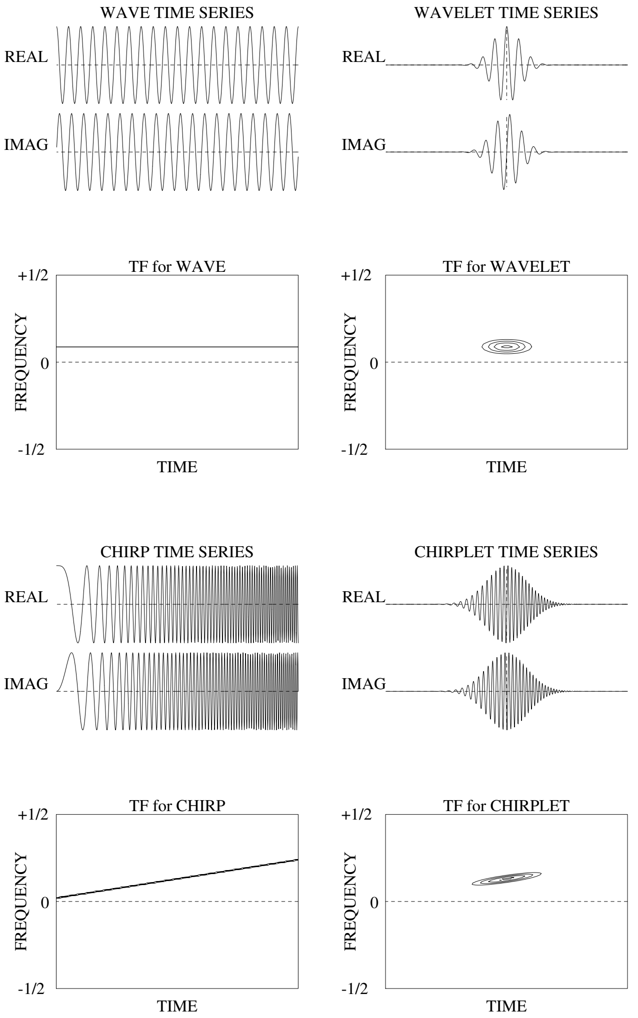

Figure 2: Relationship between wave, "wavelet", chirp and chirplet, in terms of time series and magnitude time-frequency (TF) distributions. The "wavelet" provides a tiling the TF plane with tiles that are lined up with the time and frequency axes, while the chirplet permits us to construct a more general tiling of the TF plane because the tiles may rotate or shear. More generally, each of these four functions is actually a chirplet. For example, the wave is a special case of a chirplet where the chirprate is zero and the window size is arbitrarily large. Note the use of a bipolar frequency axis, since we often wish to distinguish between positive and negative frequency components. Figure reproduced from [26]; used with permission.

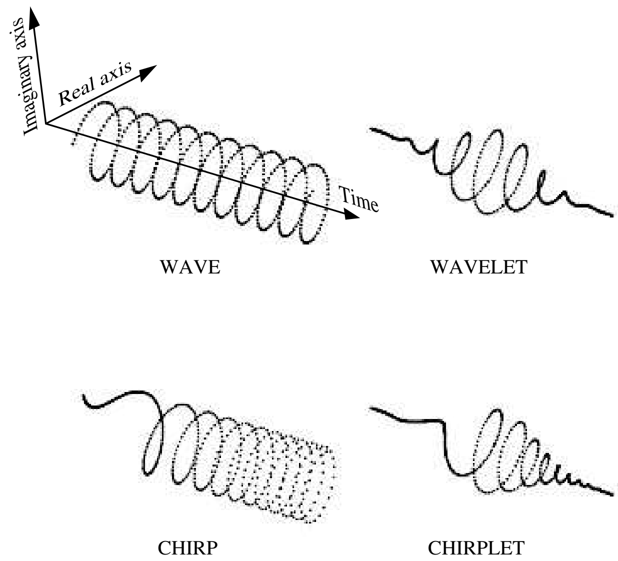

Figure 3: Wave, "wavelet", chirp, and chirplet re-visited. The x-axis corresponds to the real value of the function, and the y-axis to the imaginary value. Though the functions are continuous, a coarse sampling is used to enhance the 3-D appearance. Each sample is rendered as a *particle* in (x, y, t) space. (WAVE) The wave appears as a 3-D helix. The angle of rotation between each sample and the next is constant, hence the frequency, which is the rate of change of phase with respect to time, is constant. (WAVELET) The "wavelet" is a windowed wave, where the reduction in amplitude is observed as a decay toward the t axis. The angle of rotation between each sample and the next is still constant. (CHIRP) The chirp is characterized by a linearly increasing angle of rotation between one sample and the next. Note the increased *particle* density at the origin. (CHIRPLET) The chirplet is characterized by the same linearly increasing angle of rotation, but with first a growing and then a decaying amplitude.

the reciprocal nature of Δ_t and Δ_f, we are also implicitly specifying the bandwidth.

In (2), we can identify the Gaussian part as an envelope, which is modulated by a harmonic oscillation. The family of Gaussian chirplets is given by replacing the harmonic oscillation (*wave*) with a linear FM *chirp*:

$$g_{t_c,f_c,\log(\Delta_t),c}(t) = \frac{1}{\sqrt{\sqrt{\pi}\Delta_t}} e^{-\frac{1}{2}(\frac{t}{\Delta_t})^2} e^{j2\pi(c(t-t_c)^2+f_c(t-t_c))} \tag{3}$$

where we have used a logarithmic scale for the duration so that the unit width (default) is represented by a parameter of zero. Whenever a parameter is missing from the parameter list, we will assume it to be zero. For example, if only three parameters are present, we assume zero chirprate; if only two are present, we also assume that the log-duration is zero ($\log(\Delta_t) = 0$). Summarizing, the Gaussian chirplet (3) has four parameters: time-center, t_c; frequency-center, f_c; log-duration, $\log(\Delta_t)$; and chirprate, c.

2.2 Notation

The family of chirplets is generated from the mother chirplet by applying simple parameterized mathematical operations to it. The parameters of these operations form an index into the chirplet family.

The operations corresponding to the coordinate axes of the chirplet transform parameter space are presented in Table 1. The operators will be explained as they are used. The general notion to keep in mind is that any combination of these operators results in a 2-D affine coordinate transformation in the TF plane, which may be represented using the homogeneous coordinates often used in computer graphics [43].

The continuous STFT may be formulated as an inner product of the signal with the family of functions given in (2):

$$\mathbf{S}_{t_c,f_c} = \langle g_{t_c,f_c,\log(\Delta_t)} | s(t) \rangle \tag{4}$$

where Δ_t is a suitably-chosen (fixed) window size, and $s(t)$ is the original signal. We use the Dirac innner product notation, defined by:

$$\langle g|s \rangle = \int_{-\infty}^{\infty} g^*(t) s(t)\, dt \tag{5}$$

where g^* denotes the complex conjugate of g. We use the vertical bar between the arguments and absorb the conjugation into the first element so that we can write $\langle g|$ by itself, as an operator that acts on whatever follows, in this case the signal, $|s\rangle$.

Suppose we take the Gaussian window, centered at $t = 0$, with unit pulse duration, as given by:

$$g(t) = \frac{1}{\sqrt{\sqrt{\pi}}} \exp(-\frac{1}{2}t^2) \tag{6}$$

We denote a time shift to the position t_c, with an operator that has a multiplicative *law of composition*: $\boxed{\rightrightarrows}_{t_c}$ (Table 1). A frequency shift to the position f_c consists of multiplying the window by $\exp(j2\pi f_c t)$, which we will denote $\boxed{\upuparrows}_{f_c}$. The single-operator notation (Table 1, second column) consists of a pictorial icon depicting the effect each operator has on the TF plane, even when the operator is acting in the time-domain. For example the symbol with the two up-arrows indicates a uniform upward shift along the frequency axis, of the time-frequency plane, for positive values of the parameter. These pictorial icons are consistent with our observation that each of these operators acts in the time domain to perform an area-preserving *affine*[5] coordinate transformation in the time-frequency plane.

Using the new notation, we can re-write (4) as:

$$\mathbf{S}_{t_c,f_c} = \left\langle \boxed{\rightrightarrows}_{t_c} \boxed{\upuparrows}_{f_c} g \middle| s \right\rangle \tag{7}$$

where we have also eliminated the time coordinate, recognizing that for any operator in the time-domain, there is an equivalent operator in the frequency domain, or in the TF plane, or in whatever other reasonable coordinate space one might wish to work in. The *multiplicative law of composition* of the *operators* is applied in the order that they appear from right to left (e.g. $\boxed{\upuparrows}_{f_c}$ is applied first, and then $\boxed{\rightrightarrows}_{t_c}$ is applied to that result). Note that these two operators do not commute. Adopting the convention of applying the frequency-shift first, and then the time-shift, results in the term $t - t_c$ appearing in the second exponent of (2). Applying the operators in the reverse order would result in a different phase-shift. In order to form a true group, we need a third parameter, ϕ, to indicate the degree to which the two operators do not commute. Such a group structure is known as the Heisenberg group [12]. If we are only interested in the magnitude of the TF plane (e.g. the spectrogram) then we can simply consider the two-dimensional (two-parameter) translational group, and describe the operations in terms of this simpler group. Both (4) and (7) are equivalent, providing us with some measure of the signal energy *around* coordinates (t_c, f_c), but (7) emphasizes the fact that the STFT is a correlation between members of a two-parameter family of time-shifted and frequency-shifted versions of the same primitive, g.

Using the simplified law of compostion, we may compose a time shift by t_c with a frequency shift by f_c, as follows:

$$\boxed{\rightrightarrows}_{t_c} \boxed{\upuparrows}_{f_c} = \mathcal{C}_{t_c,0,0,0,0}\mathcal{C}_{0,f_c,0,0,0} = \mathcal{C}_{t_c,f_c} \tag{8}$$

where omissions from the parameter list of C indicate values of zero.

Equation 7) may be re-written, using the "Composite notation" (Table 1, third column):

$$\mathbf{S}_{t_c,f_c} = \left\langle \mathcal{C}_{t_c,f_c} g(t) \middle| s(t) \right\rangle \tag{9}$$

[5]Segal [44] and others sometimes refer to these coordinate transformations as *symplectomorphisms*. It is well-known [12], [45] that the actual geometry of phase space is symplectic geometry, and that it is a coincidence that SP_2 corresponds to area-preserving affine geometry. Therefore, we must keep in mind, that if we desire to extend our thinking to the analysis of signals of dimension $n > 1$, then we must consider the symplectic geometry of SP_{2n}.

Table 1: Operators corresponding to the coordinate axes of the chirplet transform parameter space

Description	1-parameter notation	Composite notation	Time-domain $g(t)$
Time-translation	$\boxed{\rightrightarrows}_{t_c} g(t)$	$= \mathcal{C}_{t_c,0,0,0,0}\ g(t)$	$= g(t - t_c)$
Frequency-translation	$\boxed{\uparrow\uparrow}_{f_c} g(t)$	$= \mathcal{C}_{0,f_c,0,0,0}\ g(t)$	$= e^{j2\pi f_c t} g(t)$
Time-dilation/Freq.contraction	$\boxed{\begin{smallmatrix}\leftrightarrow\\ \leftrightarrow\end{smallmatrix}}_{\Delta_t} g(t)$	$=\mathcal{C}_{0,0,\log(\Delta_t),0,0}\ g(t)$	$=\frac{1}{\sqrt{\lvert\Delta_t\rvert}} g(\frac{t}{\Delta_t})$
Frequency-shear	$\boxed{\downarrow\uparrow}_{c} g(t)$	$= \mathcal{C}_{0,0,0,c,0}\ g(t)$	$= e^{j2\pi\frac{c}{2}t^2} g(t)$
Time-shear	$\boxed{\rightleftarrows}_{d} g(t)$	$= \mathcal{C}_{0,0,0,0,d}\ g(t)$	$= e^{j2\pi\frac{d}{2}t^2} \star g(t)$

2.3 Time-Frequency-Scale Volume

The STFT is a mapping from a one-dimensional function (the *domain*, which is a function of time) to a two-dimensional function (the *range*, which is a function of time and frequency). Now suppose, rather than holding Δ_t constant (4), we also allow it to be a parameter. The new mapping we so obtain is a mapping from the one-dimensional *domain* (time) to a three-dimensional *range* (time, frequency, and log-scale) that we previously referred to as the TFS parameter space[6] (Fig. 1(a)).

2.4 Gaussian Chirplet Transform (GCT)

We can further extend the multidimensional parameter space. Suppose we also allow the chirprate, c, in (3) to be one of the coordinates of the parameter space. The resulting transform is given by:

$$\mathbf{S}_{t_c,f_c,\log(\Delta_t),c} = \left\langle \mathcal{C}_{t_c,f_c,\Delta_t,c}\, g(t) \middle| s(t) \right\rangle \tag{10}$$

We refer to (10) as the "Gaussian chirplet transform" (GCT).

One characteristic of the one-dimensional Gaussian window is that its TF energy distribution is a bi-variate Gaussian function. Therefore its TF energy countours are elliptical, so shearing the TF distribution along the time axis provides no new degrees of freedom that cannot be obtained by combinations of shearing along the frequency axis together with dilation. If we consider other windows, however, we do not, in general, have this degenerate property.

[6] Note that, if we were interested in exploiting the phase of this representation, we would need to add a fourth parameter, to account for the extent to which the operators do not commute.

2.5 Continuous Chirplet Transform (CCT)

We have been using the frequency shear operator, which we obtained through multiplication by a linear FM chirp. In a dual manner, we may introduce the time shear operator (Table 1, last row) which we obtain by convolving with a linear FM chirp.

Fourier transformation of a chirp, with chirprate, d, produces another chirp which has chirprate $-1/d$. Thus convolution of a signal, $s(t)$, with a chirp, having chirprate d is equivalent to multiplying $S(f)$ with a chirp of rate $-1/d$, and taking the inverse Fourier transform of the product. In short, we have rotated the TF plane 90°, sheared it left-right, and rotated it back. This three-step process has the net effect of shearing the TF plane top-bottom.

The full continuous chirplet transform (CCT) is defined in the same manner as (10):

$$\mathbf{S}_{t_c,f_c,\log(\Delta_t),c,d} = \left\langle \mathcal{C}_{t_c,f_c,\Delta_t,c,d}\, g(t) \middle| s(t) \right\rangle \tag{11}$$

except that we have one new operator, time-shear, that is composed with the other four operators.

Again, the *law of composition* [46] of any two chirplet operators (multiplicatively) follows by virtue of the fact that both represent affine coordinate transformations of the TF plane.

The intuition behind (11) is that entries in the first column of Table 1 simply represent the coordinate axes of the multidimensional parameter space, and their subscripts represent the distances along these axes.

Segal exploited various coordinate transformations in the TF plane in the development of his theory of dynamical systems of infinitely many degrees of freedom [47]. His *harmonic map* or *oscillator map*, as he called it (the *Segal Shale Weil representation* [48]), is indeed related to the chirplet transform.

2.6 Multiple Mother Chirplets: The Prolate Chirplets

2.6.1 Thomson's method of spectral estimation

Thomson's multiple window method of spectral estimation [49] provides a very good estimate of the power spectrum by measuring the energy contained within a collection of rectangular[7] shaped frequency intervals. The spectral estimate is formulated by averaging together, with appropriately chosen weightings (the eigenvalues), multiple power spectral estimates, each computed with a different window.

The windows comprise a family of discrete prolate spheroidal sequences (DPSS), have been studied extensively (Landau, Pollack, Slepian [20], [21], [22], [23], [24]) and are commonly referred to as *prolates* or *Slepians*.

The remarkable property of this family of windows is that their energy contributions add up in a very special way that collectively defines an ideal (ideal in the sense of the total in-bin versus out-of-bin energy concentration) rectangular frequency bin. Furthermore, for a time series of a given length, the power spectrum may be estimated at various resolutions (e.g. we can choose the frequency bin size). While it might at first seem unclear why one would want anything other than the highest resolution, the Thomson method allows us to trade resolution for improved statistical properties (reduced variance of the spectral estimate). Often, much of the fine structure of a spectral estimate is due to noise. It should be stressed that while other methods of spectral estimation (such as the Welch [50] method) exist, the Thomson method is particularly noteworthy for its precisely defined rectangular frequency bins.

Generally, the Thomson method is thought of as a *multiple window method*, but another way of thinking of the Thomson method is by the way that the energy in each frequency bin is calculated. To determine the quantity of energy inside the bin centered at f_c, we frequency-shift each of the windows to f_c, and sum the energy contributions from each of the frequency-shifted windows:

$$S(f_c) = \sum_i \left| \left\langle \mathcal{C}_{0,f_c,0,0,0}\, g_i \middle| s \right\rangle \right|^2 \qquad (12)$$

Writing the Thomson method in this way, we can generalize it further by replacing the one-parameter operator, $\mathcal{C}_{0,f_c,0,0,0}$, with multi-parameter operators.

2.6.2 True rectangular tiling of the TF plane

While many researchers depict certain tilings of the TF plane (such as given by the STFT), schematically, using rectangular grids [7], and even refer to them as rectangular tilings, it is important to note that the actual shape of the individual tiles is better described as a tesselation of overlapping "blobs", perhaps Gaussian, as was the case with the Gaussian-windowed STFT.

However, the same family of discrete prolate spheroidal sequences (DPSS) used in the Thomson method synthesizes a concentration of energy in the TF plane where the energy is uniformly distributed throughout one small rectangular region, and minimized elsewhere[8].

Observing this fact (others have also observed this fact [48]), we now extend the Thomson method to operate in the TF plane. In practice, we calculate a discrete version from the discrete-time signal, simply by partitioning the signal into short segments and applying the Thomson method to each segment. This amounts to a sliding-window spectral estimate where the entire family of windows slides together. As in (12), however, we may write the proposed time-frequency distribution, pointwise. That is, to calculate the energy within a rectangle centered at (t_c, f_c), we sum over the set of windows that have all been moved to the point (t_c, f_c):

$$S(t_c, f_c) = \sum_i |\langle \mathcal{C}_{t_c,f_c} g_i \mid s \rangle|^2 \qquad (13)$$

2.6.3 Pyramidal (multiresolution) true-rectangular TF tiling

The area occupied by a particular family of DPSS is just the *time-bandwidth product* and is denoted by the letters NW (the notation used by Thomson and others). The quantity N denotes the number of samples (duration) of a window, and W denotes the bandwidth collectively defined by a plurality of such windows of equal length. We can compute the TF plane, of a particular signal, at any desired value of NW, by using the discrete prolate sheroidal sequences (subject to the constraints that NW can only be adjusted in integer increments, and that it also has a lower bound dictated by the uncertainty relation [51]). If we compute the TF plane at each possible value of NW, and stack these one above the other (Fig. 4), we obtain a 3-parameter space, where the axes are time, frequency, and resolution ($1/NW$).

Hierarchical or *pyramidal* [52] representations have been previously formulated, in the context of image processing, using multiple scales in the physical domain (e.g. spatial scale). The proposed multi-resolution TF representation, however, is new. In particular, here the scale axis is NW – the area of the rectangular tiles at each level of the pyramid. Here the scale is in the time-frequency plane, not the physical (time or space) domain.

At this point, a reasonable question to ask might be: Why vary the area; do we not always desire maximum resolution or maximum concentration in the TF plane? The same answer we gave earlier, regarding smoothed spectral estimates, applies here.

Smoothing is well-known in time-frequency analysis, particularly with the Wigner distribution where we wish to reduce or eliminate cross terms. Many smoothing kernels have been proposed [53][54]. Each of these smoothing kernels has a particular shape and many of these are optimum in one sense or another. The use of the DPSSs, however, has been shown to be equivalent to a rectangular

[7]The term "rectangular" is used here in the context of "rectangular window", meaning a 1-D function that is unity in a certain frequency interval and zero outside that interval, not to be confused with our later use of "rectangular" which will be more consistent with its everyday usage to specify a 2-D shape.

[8]In actual fact, there is a small amount of frequency smearing, but zero time smearing, as the energy is entirely contained in the time interval under consideration.

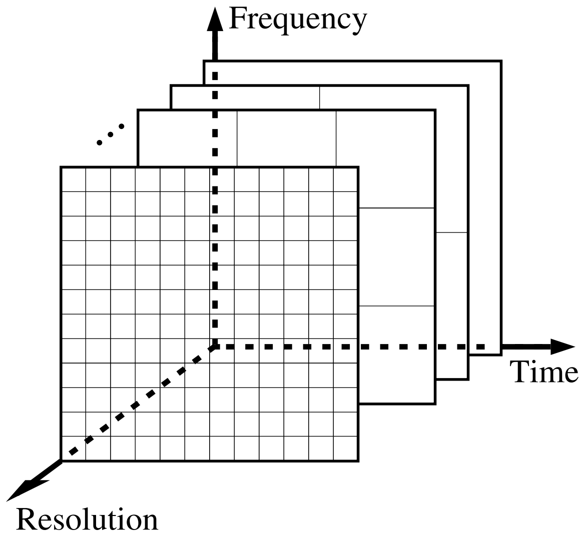

Figure 4: Three dimensional parameter space based on the use of multiple windows. This pyramidal representation may either be computed by applying the appropraite set of DPSSs to compute the "true-rectangular TF tiling" at each level, or alternatively, by computing a pyramid from the TF (Wigner) distribution of the signal. In the latter case, the TF pyramid is computed in much the same way that Gaussian pyramid of an image is computed, except using a rectangular filter rather than a Gaussian filter.

smoothing of the Wigner distribution [48], and therefore deserves special attention, particularly when we wish to describe a tiling of the TF plane in a very simple way.

We may use the result of Shenoy and Parks [48] to generalize the pyramidal true-rectangular TF tiling further, by smoothing the TF distribution with a continuously variable rectangle size. When uncountably many of these rectangulary smoothed TF planes are stacked, one above the other, a continuous volumetric parameter space results, having parameters time, frequency, and resolution.

2.6.4 Parallelogram-shaped tilings of the TF plane

The method of multiple windows may be extended further to the chirplet framework.

This further extension makes use of the same families of multiple windows that are used in the Thomson method, and that we first extended to the true rectangular tiling of the TF plane, but instead they will now be used within the context of the operators of Table 1. In the same way that the Thomson method consists of computing power spectra with a plurality of windows, and averaging the power spectra together, we compute the power CCTs with a plurality of windows, and average the results together. To compute an appropriately-smoothed version the chirplet transform, we compute a CCT (11) using each one of the multiple windows as the mother chirplet. We then average the squared-magnitudes of the resulting CCTs together, weighting by the eigenvalues, just as with the Thomson method. This gives us the CCT at a particular value of NW.

Alternatively, we may consider a given point in the five-dimensional CCT parameter space, say, $(t_c, f_c, \log(\Delta_t), c, d)$, is given by applying the operator $\mathcal{C}_{t_c,f_c,\log(\Delta_t),c,d}$ to the set of multiple windows and then computing the sum of absolute squared energy:

$$S(t_c, f_c, \log(\Delta_t), c, d) = \sum_i |\langle \mathcal{C}_{t_c,f_c,\log(\Delta_t),c,d}\, g_i(t)\, | s(t) \rangle|^2 \tag{14}$$

We now refer to the multiple windows as "multiple mother chirplets" as they have collectively taken the role of the single mother chirplet. They act collectively to produce an idealized parallelogram-shaped smoothing of the TF (Wigner) distribution (Fig. 5), where the area of the parallelogram is NW.

For example, if we apply a frequency shear, with parameter $c = 0.85$, to each of the mother chirplets, the new set of functions will collectively occupy the parallelogram-shaped region of the TF plane indicated in the Fig 5, lower right. This energy concentration represents a single point located at coordinates $(0, 0, 0, 0.85, 0)$ in the averaged squared-magnitude CCT.

2.6.5 The pyramidal (multi-resolution) CCT

Suppose we compute the above CCT (Section 2.6.4) at a few different tile-sizes, and combine these CCTs into a single six-parameter representation. The value of tile-size, NW, may be thought of as a sixth coordinate axis in the chirplet transform parameter space – TF-area. Including this sixth coordinate axis provides us with a hierarchical (multi-resolution) CCT.

To compute the proposed hierarchical CCT, we repeat the computation of the CCT (14) for each of the desired tile sizes and place them in a six-dimensional space, equally spaced along the sixth coordinate axis. Part of the computation involves re-synthesizing a new set of multiple mother chirplets for each value of NW.

Various 2-D slices through the multiresolution CCT may correspond to useful tilings of the TF plane with true parallelograms (true to the extent that the DPSS define a truly rectangular region in the TF plane). For example, the time-scale slice of the multiresolution CCT, taken at a particular resolution, is a wavelet transform based on multiple mother wavelets.

Again, we may use the result of Shenoy and Parks [48] to generalize the multiresolution CCT by smoothing the TF distribution with a continuously variable parallelogram size. When uncountably many of these parallelogram-smoothed TF planes are "stacked", a continuous six-dimensional parameter space results, having parameters time, frequency, scale, chirp, dispersion, and resolution.

Others have done work to further generalize energy concentration to arbitrarily-shaped regions of the TF plane [55], rather than just parallelograms. It would therefore be possible to use these results to define more general parameterizable transforms, based on families of multiple analysis primitives acting collectively in the TF plane.

2.7 Autochirplet and Cross Chirplet Transforms

If, in (11) we choose the mother chirplet to be the signal itself:

$$\mathbf{S}_{t_c,f_c,\log(\Delta_t),c,d} = \langle \mathcal{C}_{t_c,f_c,\log(\Delta_t),c,d}\, s(t) | s(t) \rangle \qquad (15)$$

then we have a generalization of the autocorrelation function, where, instead of only analyzing time-lags we analyze self-correlation with time-shift, frequency-shift, and chirprate. We call this generalization of autocorrelation the 'autochirplet ambiguity function'. If, for example, the signal contains time-shifted versions of itself, modulated versions of itself, dilated versions of itself, time-dependent frequency-shifted versions of itself, or frequency-dependent time-shifted versions of itself, then this structure will become evident when examining the 'autochirplet ambiguity function'. The 'autochirplet ambiguity function' is not new, but, rather, was proposed by Berthon [18] as a generalization of the *radar ambiguity function*. Note that the radar ambiguity function [56], [11] is a special case of (15).

It is well known that the power spectrum is the Fourier transform of the autocorrelation function, and that the Wigner distribution is the two-dimensional (rotated) Fourier transform of the radar ambiguity function. Recent work has also shown that there is a connection between the wideband ambiguity function and an appropriately coordinate-transformed (to a logarithmic frequency axis) version of the Wigner distribution [57], where the connection is based on the Mellin transform. This connection gives us a link between the three-parameter "time-shift—frequency-shift—scale-shift" subspace of (15) and the time-frequency-scale subspace of the chirplet transform. Extending this relation to the entire five-parameter CCT would give us the autochirplet transform. This extension is one of our current research areas in the continued development of the chirplet theory.

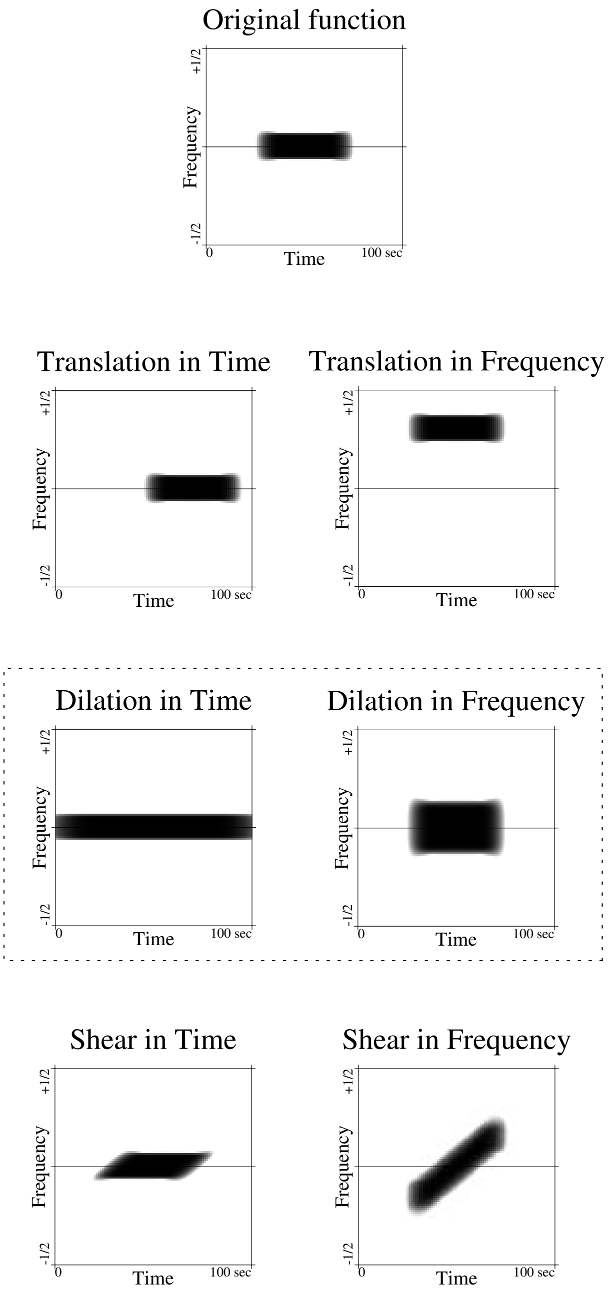

Figure 5: Illustrating the six affine transformations of the TF plane, using multiple "mother chirplets". In this example, the mother chirplets consist of a a set of 24 Discrete Prolate Spheroidal Sequences (DPSS) that collectivey define a rectangular energy concentration in the TF plane with an area $NW = 12$. Members of this chirplet family each comprise $2NW$ functions that collectively define some parallelogram-shaped region of the TF plane. When considering the tile-size as an additional parameter, there are six dimensions in the chirplet transform parameter space. Figure reproduced from [39]; used with permission.

3 CHIRPLET TRANSFORM SUBSPACES

In practice, from a computational, data storage, and display point of view, the chirplet transform is unwieldy. Therefore, we consider subspaces of the entire parameter space. Planes are particularly attractive choices in this regard both because of the ease with which they may be printed or displayed on a computer screen, and the fact that they lend themselves to finite-energy parameter spaces.

Well-known examples are the TF and TS planes discussed previously. Other subspaces, however, correspond to entirely new transforms. For example, consider the chirprate-frequency (CF) plane, computed with a Gaussian window (Gaussian so that chirprate and dispersionrate do not need to be dealt with separately). It turns out to be useful in two cases: (1) when we have only a short segment of data we wish to analyze (and therefore do not wish to partition it into even smaller time segments by the STFT), or (2) when we have a longer time-series, but are not interested in the time axis. In the latter case, the CF plane lets us average out time, and observe long-term slowly-varying frequency trends.

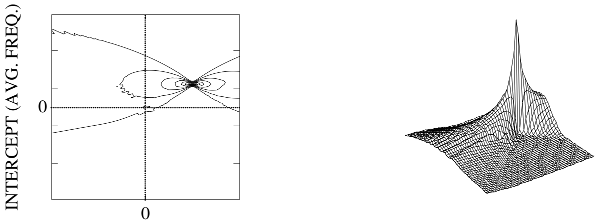

Figure 6: Chirprate-frequency (CF) plane calculated for a signal which is itself a chirp. The bowtie-shaped spread around the peak is due to the finite length of the analysis interval. Figure reproduced from [26]; used with permission.

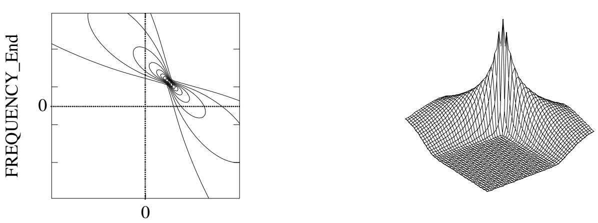

Figure 7: Frequency-Frequency (FF) plane of chirplet transform computed from a pure tone. Here we parameterize the chirplets by a change of coordinates (rotation of the plane by 45 degrees), using f_{beg} and f_{end} rather than c and f_c. Figure reproduced from [26]; used with permission.

3.1 The Frequency-Frequency (FF) Plane

We begin by discussing the CF plane, and then present an argument for re-parameterizing this plane in terms of two frequency indices, leading to what we will be calling "frequency-frequency" (FF) analysis.

Consider a two-dimensional slice through the five-dimensional CCT parameter space that we defined in (11):

$$\mathbf{S}_{c,f_c} = \langle C_{0,f_c,0,c,0}\, g(t) \,|\, s(t) \rangle \tag{16}$$

where $s(t)$ is an arbitrary time series, and the two dimensions of the transform space are the slope of the frequency rise, c, and the center frequency f_c. This transform is known [26] as the "bowtie ($\bowtie$) subspace" since the CF plane of a chirp is a sharp peak surrouned by faint bowtie-shaped contours (Fig. 6). Computing the CF plane of a signal, $s(t)$, is equivalent to correlating the signal with a family of chirps that are parameterized by chirprate, c and center frequency, f_c. Calculating the CF plane from a signal that contains pure tones results in peaks on the slope = 0 axis. Downchirps in the signal result in peaks to the left of this line, and upchirps result in peaks to the right.

For a discrete function[9], we would have periodicity in the CF plane, and the Nyquist boundary is diamond ($\diamond$) shaped. The Nyquist limit dictates that the chirps with the highest (lowest) c values begin with a fractional frequency of $-1/2$ ($+1/2$) and end with a frequency of $+1/2$ ($-1/2$). These chirps will both lie on the $f_c = 0$ axis of the CF plane. Consider a chirp that begins with a frequency $1/4$ and ends with a frequency of $3/4$. It has the same chirprate: $c = 3/4 - 1/4 = 1/2$, but it will violate the Nyquist limit because part of the chirp exceeds the fractional frequency of $1/2$, and will therefore give rise to aliasing.

Ideally we would like this transform to have nice rectangular boundaries for convenient viewing on a video display, so we overcome the Nyquist problem by tilting the parameter space 45^o. The new chirplets are then given by:

$$\begin{aligned} &\mathcal{C}_{0,(f_{end}+f_{beg})/2,0,(f_{end}-f_{beg})/2,0}\, g(t) \\ &= g(t)\, e^{j2\pi(\frac{f_{end}-f_{beg}}{2}t + \frac{f_{end}+f_{beg}}{2})t} \end{aligned} \tag{17}$$

where g denotes the mother chirplet. The change of coordinates from the CF plane to the FF plane is given by $f_{beg} = f_c - c$ and $f_{end} = f_c + c$. When the analysis interval (window) is of finite duration, f_{beg} may be taken to be the instantaneous frequency of the chirp at the beginning of the analysis interval (time window) and f_{end} the instantaneous frequency at the end of this interval. Since the new parameterization involves two frequency coordinates, we will refer to the resulting parameter space as the "frequency-frequency" (FF) plane. Figure 7 shows the FF plane computed from a harmonic oscillation.

The value of the function defined on the FF plane, evaluated at the origin gives a measure of how strong the chirp component from 0 to 0 (the DC component) is. The value at coordinates (0,1/2), for example, gives the strength of the component of a chirp going from a frequency of 0 to 1/2. Values of the FF plane in the upper left half (above and to the left of the diagonal $f_{beg} = f_{end}$) correspond to upchirps; those to the lower right correspond to downchirps. The values of the FF plane along the diagonal line, $f_{beg} = f_{end}$, define the Fourier transform of the original time-domain signal; the windowed version of the signal may be entirely reconstructed from only the diagonal of the complex-valued FF plane.

3.2 A Simple Example With a Single Chirp Component

In this first example, we allow an object to fall onto a small radar unit[10]. The resulting Time-Frequency distribution is shown as a contour plot in Fig. 8(a). We extract the portion of the recorded data which contains the object when it is in free fall (from the time after it was released, to just before the time it hit the radar horn). From this portion of the time series (the corresponding TF distribution appears in Fig. 8(b)), we compute the FF-plane through the CCT, which is simply a correlation between the signal and a family of chirplets parameterized in terms of beginning and ending frequencies. Its density plot appears as an image in Fig. 9. The response has a very high peak, as evidenced from Fig. 10.

[9] We do not attempt to address issues of discretization in this paper, except to the extent to which they have influenced the development of the continuous chirplet transform.

[10] For this experiment, we positioned the radar horn facing upward and held a volleyball two meters above the horn, and released the ball after the recording began. We recorded only the in-phase component, and ignored the

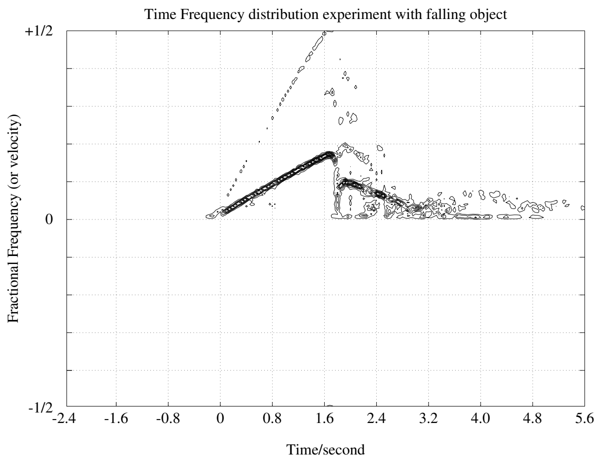

(a)

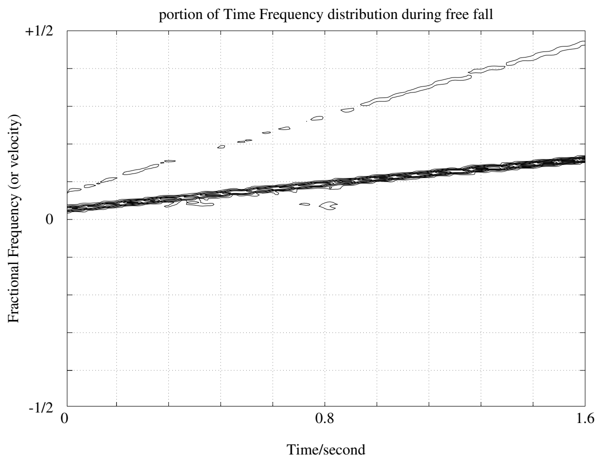

(b)

Figure 8: Illustrative TF example: actual data from a uniformly acclerating object (falling ball). Third harmonics are visible, due to non-linearities and slight clipping in radar. (a) Note the spurious effects as the ball bounces around after it has fallen. (b) Detail of portion of data for which object is in free fall.

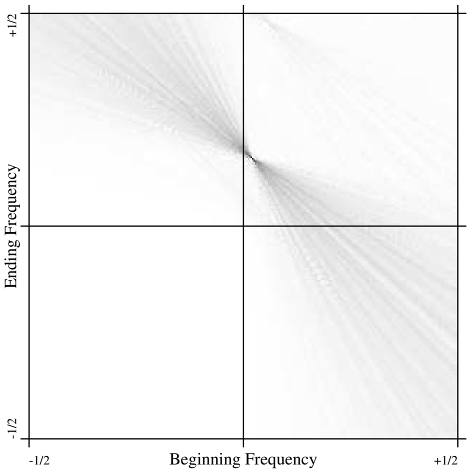

Figure 9: Frequency-Frequency (FF) plane of the chirplet transform taken for radar data from uniformly accelerating object. Note the location of the peak, indicating a near-zero initial velocity, and a much higher final velocity.

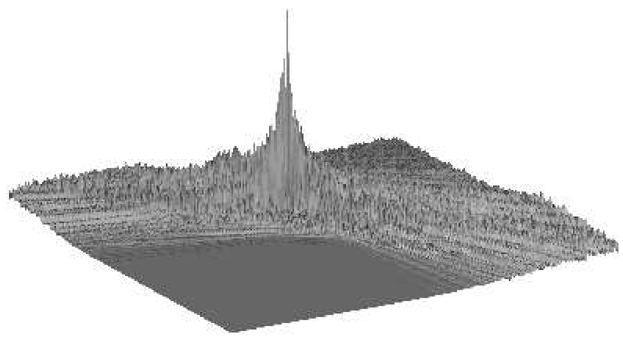

Figure 10: Shaded surface of Frequency-Frequency (FF) chirplet plane for radar return of falling object. The localization in the FF chirplet plane, for uniformly accelerating objects is even more visible here. Also, note the absense of negative frequency components (lower quadrant).

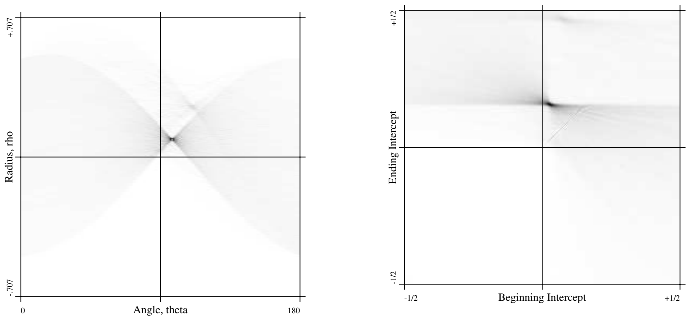

(a) (b)

Figure 11: (a) Radon transform of the TF (Wigner) distribution of the radar return from a uniformly accelerating (falling) object. Since the Doppler return of the continuous wave radar is a linear FM chirp, the TF distribution had a single linear component. A sharp localization in Radon space resulted (except for the smaller peak which is due to radar non-linearities, mainly third harmonic distortion). (b) FF plane of autochirplet transform: a new parameterization of the Radon transform allows its parameters to take on a new physical significance when the input "image" is the TF plane. The abscissa has the meaning of *beginning frequency* and the ordinate represents the *ending frequency*. Notice the diagonal slanted bowtie shape, and the similarity to the bowtie shape in the FF plane of Fig. 10.

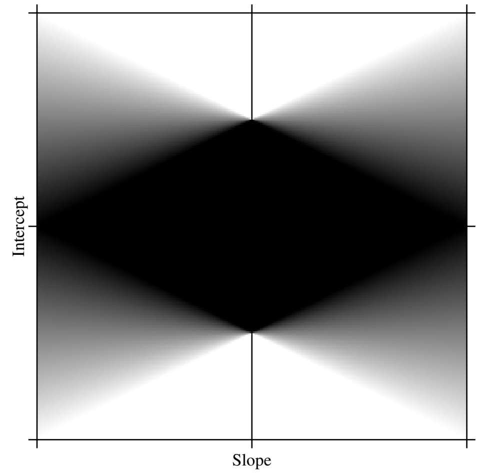

Figure 12: The "Nyquist problem" revisited: Radon transfom computed from an identically non-zero image. The commonly used slope-intercept parameterization of the Radon transform results in the ◇-shaped region similar to our "Nyquist boundary" in the CF plane.

3.3 Relationship Between Auto Chirplet FF Plane and Radon Transform

Conceptually, each point in the FF plane corresponds to a chirp component in the original signal, which also corresponds to a linear portion of the time-frequency (TF) plane. The Radon transform (also known as the Hough transform) is formulated as a family of line integrals through a two-dimensional function. It is known for its ability to extract straight lines from images. For a good survey paper on the Radon transform, see Illingworth and Kittler [58]. This property allows us to use it as an alternate means of computing the FF plane of the chirplet transform, by using the TF plane as our input image.

The Radon transform provides us with a simple means of computing the FF plane of the autochirplet transform, by using the Wigner distribution, and arriving at a transform space that tells us basically the same information as the chirplet FF-plane, except that we benefit from the greater resolution of the Wigner distribution. It is well known that the cross components of the Wigner distribution are of an oscillatory nature, while the autocomponents give a net positive contribution. Therefore, since the Radon transform is integrating along lines, the *cross terms* of the Wigner distribution are cancelled out along each line, so that the points in the Radon transform of the Wigner distribution only "see" the autocomponents of the Wigner distribution (Fig. 11(a)).

The Radon transform is usually computed from the *normal equation* of a line:

$$x\cos(\theta) + y\sin(\theta) = \rho \tag{18}$$

as an integral along each of these lines in the original space. The parameter space is sampled uniformly, in the (θ, ρ) coordinates. It is easier to compare the Radon transform of the Wigner distribution with the chirplet CF plane (Fig. 10) if, by first, without loss of generality, we normalize f_{beg} and f_{end} to be on the interval from $-1/2$ to $1/2$ and the TF distribution to have time and frequency coordinates on the same interval from $-1/2$ to $1/2$. Then we make the substitution:

$$\sin(\theta) = \rho / f_{avg} \tag{19}$$

and

$$\tan(\theta) = 1/f_{diff} \tag{20}$$

where $f_{diff} = f_{end} - f_{beg}$ and $f_{avg} = (f_{end} + f_{beg})/2$.

A simpler (perhaps equally well known) form of the Radon transform parameterizes the lines in terms of their slopes and intercepts. This parameterization has the advantage that it maps lines to points, and points to lines, while it has the disadvantage that there is a singularity when lines of infinite slope (vertical lines) are encountered. Because of the Nyquist limit, however, we do not have this problem when the input to the Radon transform is a time-frequency distribution. Thus we may be tempted to use the slope-intercept form of the Radon transform, except that we would prefer to have a parameterization in that matches the FF plane rather than the CF plane, for reasons previously discussed. The "Nyquist boundaries" we referred to earlier are most evident if we simply consider the discrete Radon transform of a matrix of identically non-zero values (Fig. 12), where we can observe the same diamond shape which initially prompted us to use f_{beg} and f_{end} rather than f_{diff} and f_{avg}.

We may overcome the problems associated with boundaries by defining a new version of the Radon transform, where we use the following pair of parameters:

- Beginning intercept, f_{beg}: the leftmost ordinate on the line (the ordinate for an abscissa of $-1/2$).

quadrature component of the radar. The sampling rate was 8kHz.

- Ending intercept, f_{end}: the rightmost ordinate on the line (the ordinate for an abscissa of $+1/2$).

In Fig. 11(b) we show the autochirplet FF plane calculated from the falling-object data, using the new parameterization of the Radon transform.

3.4 Nondilational Chirplet Transform

We do not address discretization issues in this paper. However, it is worth noting, that in practice, we generally wish to compute the chirplet transform of a discrete-time signal, and it is sometimes the case that the mother chirplet is also discrete-time and has no closed-form mathematical description. Thus, *dilation* would require resampling, and *contraction* would require antialiasing. In this case, the largest subspace we might obtain would be the subspace that omits both dilation and tiling-size, leaving us with the four-dimensional parameter space:

$$\mathbf{S}_{t_c,f_c,0,c,d} = \left\langle \mathcal{C}_{t_c,f_c,0,c,d}\; g \,\middle|\, s \right\rangle \tag{21}$$

3.5 Warbling Chirplet: Analysis of Signals of Oscillating Frequency

Suppose we choose a windowed sinusoidal FM signal for our mother chirplet. Such a signal has a frequency that periodically rises and falls (much like the *vibrato* of musical instruments or the wail of a police siren).

Within time-frequency space, conventional Doppler radar spectrograms treat the motion of objects as though their velocities (Doppler frequencies) were piecewise constant (constant over each of the short time intervals), whereas the chirplet transform attained a certain advantage by generalizing to a piecewise constant acceleration model.

Originally, we had further extended the linear FM chirplet bases to piecewise quadratic, and piecewise cubic FM – piecewise polynomial approximations to the time-frequency evolution of Doppler returns. However, looking more closely at the underlying physics of floating objects, which was our main motivation that led to our discovery of the CCT, we observed a somewhat sinusoidal evolution of the Doppler signals.

If you have ever watched a cork bobbing up and down at the seaside, you would notice that it moves around in a circle with a distinct periodicity. It moves up and down, but it also moves horizontally. Looking out at a target with a radar, for example, we see the horizontal component of motion (which is essentially a scaled version of the Hilbert transform of the vertical movement). This sinusoidal[11] horizontal movement results in a sinusoidally varying frequency in the Doppler return.

We wish to end up with the instantaneous frequency of the basis function being given by:

$$f = \beta \cos(2\pi f_m t + p) + f_c \tag{22}$$

where f_c is the center (carrier) frequency, p (which varies on the interval from 0 to 2π) is the relative position of one of the peak epochs in frequency, with respect to the origin, and f_m is the modulation frequency. If we are analyzing a discrete signal, $s[nT]$, we also note that $|\beta + f_c|$ must be less than $1/2$, otherwise the frequency modulation is not bounded by the Nyquist limit.

[11] Here we are simplifying the description. The dynamics of the sea are more fully described in [59].

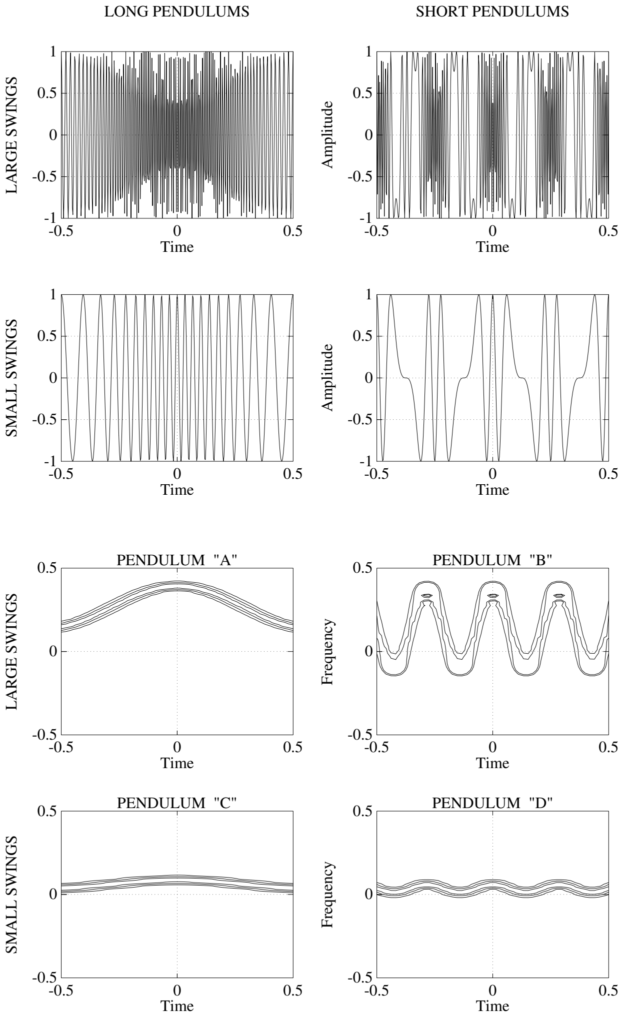

Figure 13: Four examples of warbling chirp functions. Windows have been eliminated for clarity.

Integrating to get the phase, we get:

$$\phi = \frac{\beta sin(2\pi f_m t + p)}{f_m} + 2\pi f_c t \tag{23}$$

which gives us the family of chirplets defined by:

$$g_{f_m,\beta,f_c} = Ae^{j(\frac{\beta sin(2\pi f_m t+p)}{f_m})+j2\pi f_c t} \tag{24}$$

which may be appropriately windowed, such as with a Gaussian, as was done in (3).

In Fig. 13, we show four examples taken from a family of chirplets that were derived from a warbling mother chirplet. We show them in both the time-domain, and the TF domain, annotated in terms of the pendulum model described now.

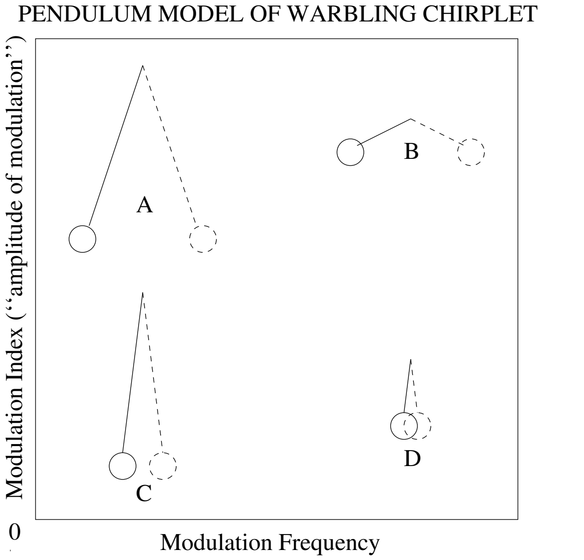

Figure 14: Four pendulums depicted at locations in the dilation-dilation chirplet plane corresponding to where each would produce maximal energy in this distribution.

Pendulums swinging to and fro in front of a radar (assume the amplitude of the swing is small compared to the length of the string) produce a signal which is very similar to that produced by radar returns from floating objects. Suppose the velocity of a pendulum, as a function of time, is given by:

$$v = \beta \cos(2\pi f_m t) \tag{25}$$

(the position is given by $\beta sin(2\pi f_m t + p)/f_m$)

The demodulated radar Doppler signal would then be given by

$$e^{j\frac{\beta sin(2\pi f_m t+p)}{f_m}} \tag{26}$$

which may be analyzed using the family of chirplets given by (24).

A pendulum with a long string, swinging with large amplitude in front of the radar will produce a time series which, will have most of its energy in the upper left hand portion of the space (low f_m and high β). A density plot of the transform, computed from the time series will show a strong peak in the upper left region, with the peak located at the coordinates corresponding to the particular frequency of swinging (f_m) and amplitude β. A pendulum with a small swing, and a short length, will appear as an energy concentration in the lower right corner of the pendulum parameter space.

Figure 14 shows where four pendulums would appear as peaks in this pendulum parameter space. Each of these four points in the space corresponds to the four examples of Fig. 13.

In Fig. 15 we show the STFT computed from an actual radar return from a pendulum.

Using the warbling mother chirplet, we also computed the "dilation-dilation" ($f_m\beta$) plane of the chirplet transform (Fig. 16) for the pendulum data.

The members of the chirplet family given by (24) may

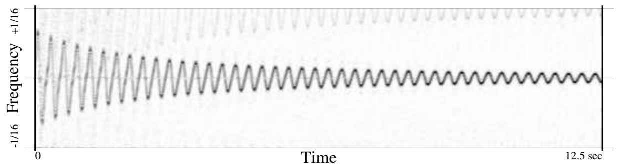

Figure 15: Time-frequency distribution of radar return from a pendulum (computed using proposed rectangular-tiling method). Note the nearly pure (except for a small decay) sinusoidal pattern in TF plane.

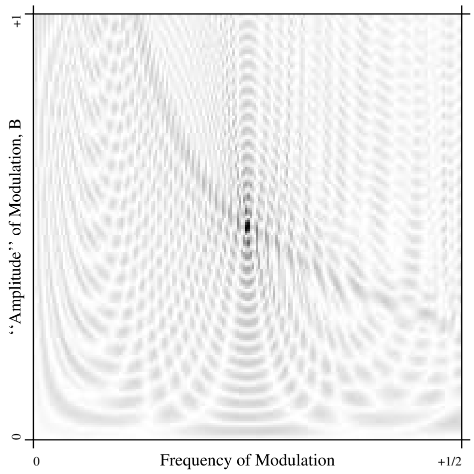

(a)

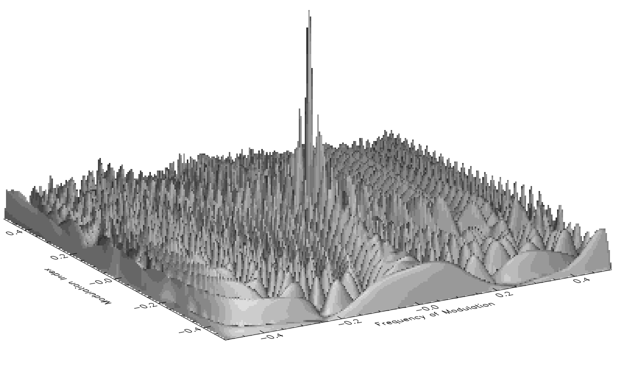

(b)

Figure 16: The dilation-dilation ($\Delta_t \Delta_f$) plane of the chirplet transform computed with a warbling mother chirplet. The signal being analyzed is a Doppler radar return from a swinging pendulum. (a) Density plot (b) Shaded surface plot

be regarded as being related to each other by affine coordinate transformations in the time-frequency plane if we use the rather abstract notion of instantaneous frequency. Consider four functions from (24), denoted A, B, C, and D, corresponding to the four signals depicted in Fig. 13 within the plane formed by time and instantaneous frequency. While there is no way to actually calculate this plane, we can consider these functions as being defined by sinusoids in this plane. Then A is a freqeuncy-dilated version of C, and B is a frequency-dilated version of D. While there is no practical means of dilating frequency without contracting time[12], or vice-versa, we may denote the abstract operator:

$$\mathcal{C}_{0,0,0,0,0,\Delta_t\Delta_f}\ g(t) \tag{27}$$

as an operator that would magnify the time-frequency distribution of $g(t)$.

When we write, for example,

$$\mathcal{C}_{0,0,0,0,0,2}\ g(t), \tag{28}$$

we mean to *replace* g(t) with another function that occupies twice the area in the TF plane. In general, such a function probably does not exist. We noted, in the case of the prolate chirplet family, that we could, however, vary the time-bandwidth product of the tiling by replacing the family of mother chirplets with a new family that had a different value of NW. By similar reasoning, within the context of the warbling chirplet we interpret (28) to mean "replace g(t) with a new sinusoidal-FM function that has $\sqrt{2}$ times the modulation index and $1/\sqrt{2}$ times the modulation frequency", so that, we obtain an equal dilation by $\sqrt{2}$ along each of the time and instantaneous frequency axes. The result is a dilation of both the time and instantaneous frequency axes by a factor of $\sqrt{2}$. The law of composition, identity, and inverse, within this six-parameter "group" is given by the usual two-dimensional affine group law [46].

Therefore, we may write the warbling chirplet transform in terms of the six affine coordinate transformations in the TF plane:

$$\mathbf{S}_{f_m,\beta,f_c} = \langle \mathcal{C}_{0,f_c,\frac{1}{f_m\beta},0,0,\frac{\beta}{f_m}}\ g(t)|s(t)\rangle \tag{29}$$

and refer to the subspace (Fig. 16) defined along the f_m and β axes as the "dilation-dilation" plane, or the $\Delta_t\Delta_f$ plane.

4 CONCLUSION

We have presented the chirplet transform, which may be viewed as a generalization of both the short-time Fourier transform (STFT) and the wavelet transform (WT). These generalizations are based on the fact that both the STFT and WT can be written as inner products of the signal under analysis with versions of a single analysis primitive (window/wavelet) acted on by various operators. In the case of the wavelet, these operators result in 1-D affine coordinate tranformations of the time-axis. In the case of the chirplet, these operators result in 2-D affine coordinate transformations of the TF plane (of the time-domain function that they operate on, if one prefers to regard the operators as acting in the time-domain). The family of chirplets is the result of a family of TF-affine coordinate transformation operators acting on a single window/wavelet (the "mother chirplet"). The chirplet transform is the resulting signal representation on this family of chirplets.

1. As is well-known, taking the Fourier transform of a one dimensional function results in a complex-valued function of a single variable.
2. As is also well-known, the STFT results in a function of two variables: time and freqency. The wavelet transform results in a complex function of two variables: time and scale.
3. The combined TFS transform results in a complex function of three variables: time, frequency and scale.
4. The Gaussian chirplet transform (GCT) results in a complex function of four variables: time, frequency, scale, and "chirprate".
5. Another complex-valued four-dimensional parameter space is given by: time, frequency, "chirprate" and "dispersionrate". This space has the interesting property that it does not require dilation of the mother chirplet, and may therefore be applied to discrete mother chirplets that do not have a mathematical description (e.g. no need for interpolation or antialiasing).
6. The full continuous chirplet transform (CCT) that can be obtained using only a single mother chirplet results in a complex function of five variables: time, frequency, scale, chirprate, and dispersionrate.
7. The multiple-mother-chirplet transform (e.g. using the prolate family) results in a real function of six variables: time, frequency, scale, chirprate, dispersionrate, and TF tile size. The coordinate axes of this six-dimensional parameter space correspond to the six affine coordinate transformations in the TF plane: translation along each of the time and frequency axes, change in aspect ratio, shear along each of the time and frequency axes, and change in area occupied in the TF plane. The last of these six dimensions is discretized, while the other five are continuous.

The chirplet transform allows for a unified framework for comparison of various time-frequency methods, because it embodies many other such methods as lower-dimensional subspaces in the chirplet analysis space. For example, the wavelet transform, the short-time Fourier transform (STFT), the "frequency-frequency" transform, and the scale-frequency transform are planar slices through the proposed multi-dimensional chirplet parameter space, while many adaptive methods [60], [61], [62] are either collections of arbitrary points or two-parameter curved surfaces (manifolds) taken from the multi-dimensional chirplet parameter space. In addition to unifying some of the existing methods, the chirplet transform provides us with a framework for both formulating and evaluating entirely new subspace transforms.

[12] Although there are devices, known as pitch transposers, that attempt to perform such an operation in a highly nonlinear way.

As pointed out in 1.1, many others have contributed directly or indirectly to the development of the chirplet transform. In many ways, however, we have taken its development further toward becoming a useful signal processing tool for practical engineering problems, as evidenced by the material presented in this paper.

ACKNOWLEDGMENT

The authors wish to express their gratitude to the following individuals for their valuable assistance: Rosalind Picard, Irving Segal, Shawn Becker, and Kris Popat of the Massachusetts Institute of Technology; Douglas Jones of the University of Illinois at Urbana-Champaign; Richard Baraniuk of Rice University; and the anonymous reviewers, whose careful efforts resulted in a substantially improved presentation.

References

[1] I. Daubechies. *Ten Lectures on Wavelets.* Number 61 in CBMS-NSF Series in Applied Mathematics. SIAM Publications, Philadelphia, 1992.

[2] C. E. Heil and D. F. Walnut. Continuous and discrete wavelet transforms. *SIAM Review*, 31(4):628–666, 1989.

[3] S. G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 11(7):674–693, 1989.

[4] I. Daubechies. The wavelet transform, time-frequency localization and signal analysis. *IEEE Trans. on Inf. Theory*, 36(5):961–1005, 1990.

[5] G. Strang. Wavelets and dilation equations: A brief introduction. *SIAM Review*, 31(4):614–627, 1989.

[6] Olivier Rioul and Patrick Flandrin. Time-scale energy distributions: A general class extending wavelet transforms. *IEEE Trans. on Signal Processing*, 40(7):1746–1757, July 1992.

[7] D. Gabor. Theory of communication. *J. Inst. Elec. Eng.*, Vol.93(Part III):429–457, 1946.

[8] William Siebert. A radar detection philosophy. IT-2(3), September 1956.

[9] William Siebert. Statistical theories of radar synthesis. October 1956.

[10] William Siebert. Woodward's uncertainty function. April 15, 1958.

[11] Merrill I. Skolnik. *Introduction to Radar Systems, Second Edition.* McGraw-Hill, second edition, 1980.

[12] G.B. Folland. *Harmonic Analysis in Phase Space.* Princeton University Press, Princeton, NJ, 1989.

[13] A. Grossmann and T. Paul. Wave functions on subgroups of the group of affine cannonical tranformations. Lecture notes in physics, No. 211: Resonances — Models and Phenomena, pages 128–138. Springer-Verlag, 1984.

[14] A. Papoulis. *Signal Analysis.* McGraw-Hill, 1977.

[15] D.L. Jones and T.W. Parks. Time-frequency window leakage in the short-time fourier transform. *Circuits, Systems, and Signal Processing*, 6(3), 1987.

[16] H.H. Szu and J.A. Blodgett. On the Locus and Spread of Pseudo-Density Functions in the Time-Frequency Plane. *Philips J. Res.*, 37:79–110, 1982.

[17] A. Janssen. On the Locus and Spread of Pseudo-Density Functions in the Time-Frequency Plane. *Philips J. Res.*, 37:79–110, 1982.

[18] A. Berthon. Operator Groups and Ambiguity Functions in Signal Processing. In J.M. Combes, editor, *Wavelets: Time-Frequency Methods and Phase Space.* Springer Verlag, 1989.

[19] S. Haykin, B. W. Currie, and V. Kezys. Surface-based radar: coherent. In S. Haykin, E. O. Lewis, R. K. Raney, and J. R. Rossiter, editors, *Remote Sensing of Sea Ice and Icebergs*, pages 443–504. John Wiley and Sons, 1994.

[20] D. Slepian and H.O. Pollak. Prolate spheroidal wave functions,Fourier analysis and uncertainty, I. *Bell System Technical Journal*, 40:43–64, January 1961.

[21] H.J. Landau and H.O. Pollak. Prolate spheroidal wave functions,Fourier analysis and uncertainty, II. *Bell System Technical Journal*, 40:65–84, January 1961.

[22] D. Slepian and H.O. Pollak. Prolate spheroidal wave functions,Fourier analysis and uncertainty, III: The dimension of essentially time-and band-limited signals. *Bell System Technical Journal*, 41:1295–1336, July 1962.

[23] D. Slepian. Prolate spheroidal wave functions,Fourier analysis and uncertainty, IV:Extensions to many dimensions;generalized prolate spheroidal functions. *Bell System Technical Journal*, 43:3009–3058, November 1964.

[24] D. Slepian. Prolate spheroidal wave functions,Fourier analysis and uncertainty, V:The discrete case. *Bell System Technical Journal*, 57:1371–1430, may-jun 1978.

[25] D. Slepian. On bandwidth. *Proc. IEEE*, 64:292–300, March 1976.

[26] Steve Mann and Simon Haykin. The chirplet transform — a generalization of Gabor's logon transform. *Vision Interface '91*, June 3-7 1991.

[27] D Mihovilovic and R.N. Bracewell. Adaptive chirplet representation of signals on time–frequency plane. *Electronics Letters*, 27(13):1159–1161, June 20, 1991.

[28] D Mihovilovic and R.N. Bracewell. Whistler analysis in the time-frequency plane using chirplets. *Journal of Geophysical Research*, 97(A11):17199–17204, November, 1992.

[29] B. Torresani. Wavelets associated with representations on the affine Weyl-Heisenberg group. *J. Math. Phys.*, 32(5):1273–1279, May 1991.

[30] J. Segman and W. Schempp. *Two methods of incorporating scale in the Heisenberg group.* 1993. JMIV special issue on wavelets.

[31] R Wilson, A D Calway, and E R S Pearson. A generalised wavelet transform for Fourier analysis: the multiresolution Fourier transform and its application to image and audio signal analysis. *IEEE Trans. on Information Theory*, 38(2):674–690, March 1992. ftp://ftp.dcs.warwick.ac.uk/reports/isp-IT38.

[32] R Wilson, A D Calway, E R S Pearson, and A R Davies. An introduction to the multiresolution Fourier transform. Technical report, Department of Computer Science, University of Warwick, Coventry CV4 7AL UK., 1992. ftp://ftp.dcs.warwick.ac.uk/reports/rr-204/.

[33] R. G. Baraniuk. *Shear madness: Signal-dependent and metaplectic time-frequency representations.* PhD dissertation, University of Illinois at Urbana-Champaign, Department of Electrical and Computer Engineering, August 1992.

[34] Richard Baraniuk and Doug Jones. Shear madness: New orthonormal bases and frames using chirp functions. *Trans. Signal Processing, vol. 41*, December 1993. Special Issue on Wavelets in Signal Processing.

[35] D. Mendlovic H. Ozaktas, B. Barshan and L. Onural. Convolution, filtering, and multiplexing in fractional fourier domains and their relation to chirp and wavelet transforms. *JOSA A.* to appear.

[36] Steve Mann and Simon Haykin. Time-frequency perspectives: The chirplet transform. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, San Francisco, CA, March 23-26, 1992. IEEE.

[37] Steve Mann and Shawn Becker. Computation of some projective chirplet transform (PCT) and metaplectic chirplet transform (MCT) subspaces, with applications in signal processing. *DSP World*, November 1992.

[38] Steve Mann and Simon Haykin. The Adaptive Chirplet: An Adaptive Wavelet Like Transform. *SPIE, 36th Annual International Symposium on Optical and Optoelectronic Applied Science and Engineering*, 21-26 July 1991.

[39] Steve Mann and Simon Haykin. Chirplets and Warblets: Novel Time–Frequency Representations. *Electronics Letters*, 28(2), January 1992.

[40] S. Mann. Wavelets and chirplets: Time–frequency perspectives, with applications. In Petriu Archibald, editor, *Advances in Machine Vision, Strategies and Applications.* World Scientific, 1992.

[41] R.G. Baraniuk and D.L. Jones. New dimensions in wavelet analysis. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, San Francisco, CA, March 23-26, 1992. IEEE.

[42] J. Cunningham and S. Haykin. Neural network detection of small moving radar targets in an ocean environment. *Workshop on Neural Networks for Signal Processing*, September 1992.

[43] Foley vanDam Feiner Hughes. *Computer Graphics, PRINCIPLES AND PRACTICE.* THE SYSTEMS PROGRAMMING SERIES. Addison-Wesley, second edition, 1990.

[44] I.E. Segal. Foundations of the theory of dynamical systems of infinitely many degrees of freedom. *Matematisk-fysiske Meddelelser*, 31(12):1–39, 1959.

[45] Victor Guillemin and Shlomo Sternberg. *Symplectic techniques in physics.* Cambridge University Press, 1984. MIT course text book.

[46] M. Artin. *Algebra.* Prentice Hall, 1991.

[47] I. E. Segal. Foundations of the theory of dynamical systems of infinitely many degrees of freedom. *Matematisk-fysiske Meddelelser*, 31(12):1–39, 1959.

[48] Ram G. Shenoy and Thomas W. Parks. The Weyl correspondence and time–frequency analysis. *IEEE Trans. Sig. Proc.*, 42(2):318–331, February 1994.

[49] D.J. Thomson. Spectrum estimation and harmonic analysis. *Proc. IEEE*, 70(9):1055–1096, September 1982.

[50] Alan V. Oppenheim and Ronald W. Schafer. *Discrete–Time Signal Processing.* Prentice Hall, 1989.

[51] David Slepian. Some comments on Fourier analysis,uncertainty and modelling. *SIAM Review*, 25(3):379–393, July 1983.

[52] Peter J. Burt and Edward Adelson. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31:532–540, April 1983.

[53] Leon Cohen. Time-frequency distributions — a review. *Proceedings of the IEEE*, 77(7):941–981, 1989.

[54] L. Cohen. *Time-Frequency Analysis.* Prentice-Hall, 1995.

[55] Jayakumar Ramanathan and Pankaj Topiwala. Time-frequency localization via the Weyl correspondence. Technical Report MTP-92B0000003, MITRE, Bedford, Massachusetts, September 1992.

[56] P. M. Woodward. *Probability and Information Theory with Applications to Radar.* McGraw-Hill, 1953.

[57] J. Bertrand and P. Bertrand. Affine time-frequency distributions. *To appear in Time-Frequency Analysis - Methods and Applications, B. Boashash (ed.).*

[58] J Illingworth and J Kittler. A survey or the Hough Transform . *Computer Vision, Graphics, and Image Processing*, August 10 1987.

[59] W. J. Pierson and L. A. Moskowitz. A proposed spectral form for fully developed wind seas based on the similarity theory of S. A. Kitaigorodskii. *Journal of Geophysical Research*, 69(24):5181–5203, 1964.

[60] Steve Mann and Simon Haykin. Adaptive "Chirplet" Transform: an adaptive generalization of the wavelet transform. *Optical Engineering*, 31(6):1243–1256, June 1992.

[61] R. G. Baraniuk and D. L. Jones. A signal dependent time-frequency representation: Optimal kernel design. *IEEE transactions on signal processing*, 41(4):1589–1602, April 1993.

[62] R. G. Baraniuk and D. L. Jones. Signal–dependent time–frequency analysis using a radially gaussian kernel. *signal processing*, 32(3):263–284, June 1993.