

Chapter 8

Sampling and Discrete Linear Canonical Transforms

John J. Healy and Haldun M. Ozaktas

Abstract A discrete linear canonical transform would facilitate numerical calculations in many applications in signal processing, scalar wave optics, and nuclear physics. The question is how to define a discrete transform so that it not only approximates the continuous transform well, but also constitutes a discrete transform in its own right, being complete, unitary, etc. The key idea is that the LCT of a discrete signal consists of modulated replicas. Based on that result, it is possible to define a discrete transform that has many desirable properties. This discrete transform is compatible with certain algorithms more than others.

8.1 Introduction

Most of the literature on the LCTs, including many chapters of this book, explicitly or implicitly make use of the continuous transform. There are, however, a number of situations in which it is desirable or necessary to use a discrete transform. Most of these are obvious by analogy with situations in which the fast Fourier transform (FFT) is used to numerically approximate the Fourier transform. In addition, the LCT is of increasing relevance in situations where we wish to model optical systems, including those with inherently discrete components such as spatial light modulators or digital cameras. In this chapter, we will explore the relationship of the LCTs with sampling, and the consequences for everything from the definition of a corresponding discrete transform to how to perform sampling rate changes accurately and efficiently.

The impact of this material should be clear to anyone who has had to comprehend the details of calculations involving the FFT. For many users, the FFT is essentially a black box that performs a Fourier transform on their data. This is close enough

J.J. Healy (✉)

School of Electrical and Electronic Engineering, University College Dublin, Belfield, Dublin 4, Ireland

e-mail: john.healy@ucd.ie

H.M. Ozaktas

Department of Electrical Engineering, Bilkent University, 06800 Bilkent, Ankara, Turkey

e-mail: haldun@ee.bilkent.edu.tr

© Springer Science+Business Media New York 2016

241

J.J. Healy et al. (eds.), *Linear Canonical Transforms*, Springer Series in Optical Sciences 198, DOI 10.1007/978-1-4939-3028-9_8

to the truth to make the FFT useful to users with even a primitive understanding of Fourier analysis. Such ‘idiot-proofing’ arises from the definition of the discrete Fourier transform (DFT), carefully chosen to be unitary, complete and have many of the properties of the Fourier transform of continuous signals. Closer examination reveals subtle differences such as circular convolution. Numerous papers from the past two decades have developed parts of a theory of numerical approximation of the LCTs. This chapter will provide a summary of some of these results, and a discussion of how they fit together to build a complete picture of the relationship between the continuous and discrete LCTs.

First, however, we must introduce some notation and certain key ideas which will crop up repeatedly in this discussion.

8.1.1 Linear Canonical Transform and Notation

This material is covered well in numerous other chapters, most particularly Chap. 2, and the reader is referred there for a broad introduction to the transform. However, there may be some notational differences, and so we have included this brief section for clarity.

Given a function of a single variable, $f(x)$, $f : \mathbb{R} \rightarrow \mathbb{C}$, the LCT of that function for parameter matrix $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is given by

$$F_M(y) = \mathcal{L}_M\{f(x)\}(y) = \begin{cases} (1/\sqrt{j b}) \int_{-\infty}^{\infty} f(x) \exp\left(\frac{j\pi}{2b} ax^2 - 2xy + dy^2\right) dx & \text{if } b \neq 0 \\ \sqrt{d} \exp(j\pi c dy^2) f(dy) & \text{if } b = 0. \end{cases}$$

We will briefly discuss the extension to complex y and complex elements of M later, but unless specified otherwise, these are taken to be real. We will also briefly mention transforms of 2 D signals, but omit the transform definition here. Finally, the $b = 0$ case is trivial, and it is generally assumed that $b \neq 0$ without significant consequence. The elements of M are referred to as the ABCD parameters, and M has unit determinant.

8.1.2 Dirac Delta Function

A key idea in understanding the relationship between discrete and continuous signals and systems is the Dirac delta function. Introduced by Paul Dirac, the delta function is commonly used to model a sampling pulse. It is defined as follows:

$$\int f(x) \delta(x - \tau) dx = f(\tau). \quad (8.1)$$

This is called the *sifting* property of the delta function. Note the form of the definition: the delta function is defined by its effect under integration. The following inadequate definition of the delta function can be useful heuristically.

$$\delta(x) = \begin{cases} \infty & \text{if } x = 0 \\ 0 & \text{if } x \neq 0. \end{cases}$$

The delta function may also be viewed as the limit of a series of functions. Consider a rectangle function which has unit area

$$r_L(x) = \begin{cases} 1/L & \text{if } |x| < L/2 \\ 0 & \text{otherwise.} \end{cases}$$

Now consider the limit of this function as $L \rightarrow 0$. The function grows narrower and taller, but always has unit area. Thus we can think of the delta function as being a pulse that is infinitely tall, infinitely thin and has unit area.

8.1.3 Time-Frequency Representations

The Fourier transform of a time-varying signal decomposes the signal into a weighted sum of its frequency components, but it doesn't localize them in time. Time-frequency representations of a signal attempt to overcome this limitation. They have proven useful in a variety of applications, but appear here because of their utility in discussing sampling problems associated with the LCTs.

The reader is referred to the introductory chapters of this book for a review of the Wigner–Ville distribution function (WDF). Briefly, a time-varying function $f(t)$ is mapped to a function of time and frequency, $W(t, k)$. Loosely, this function identifies which frequencies are active and how active they are at any given time. Strictly speaking, it is impossible to speak of an instantaneously occurring frequency component—this would breach Heisenberg's uncertainty principle—and so the WDF is referred to as a pseudo-distribution. In an optical context, the time variable is commonly replaced by a spatial variable without repercussion.

A significant simplification of the WDF which retains a lot of useful information is the *phase space diagram* (PSD), see, e.g., [20, 32]. The idea is to pick an arbitrary rectangle enclosing most of the WDF of a given signal. (The WDF must necessarily be of infinite extent, so the whole signal cannot be encapsulated. Hence, we must use nebulous terms like “most”. We must choose arbitrary limits, such as those that enclose some proportion, η , of the signal power.) The evolution of this simple geometric shape may then be tracked as the signal passed through an optical system. Such tracking is possible because the effect of the LCT (in this case, implemented as the optical system) is to perform a linear co-ordinate transformation in time-frequency [1]. The resulting parallelograms indicate the width and bandwidth of the signal throughout the system. Figure 8.1 illustrates such a PSD.

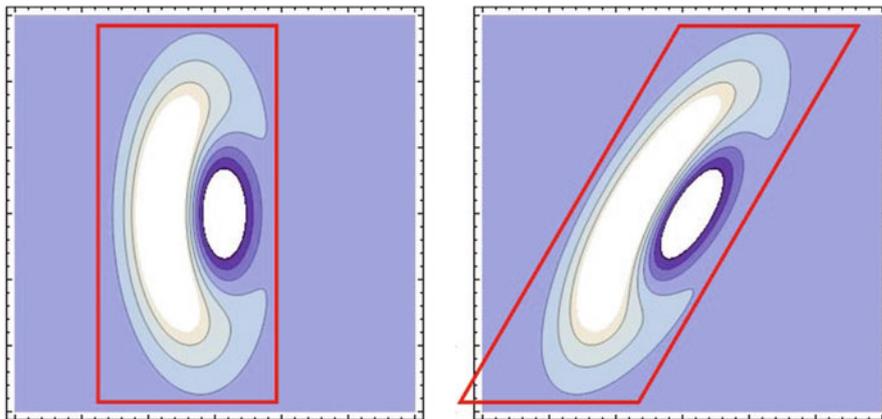


Fig. 8.1 Contour maps of the WDF signal $(3x^2 - 5x + 1)e^{-x^2}$ before (*left*) and after (*right*) free space propagation. The *red boxes* are the PSDs for the same signal

The PSD has been expanded upon by a number of authors [32]. We will encounter some of these at later points in this chapter, but let us mention one of the more relevant examples: Hennelly and Sheridan systematized and applied PSDs to show that by tracking the sampling requirements at each step of an algorithm consisting of a series of LCTs, they could obtain satisfactory results from such algorithms [12]. The flaw in their result is that it does not account for the input to those algorithms being discrete [7]. Shortly, we will see how that may be overcome.

8.2 Sampling

Shannon–Nyquist sampling is so broadly useful and applied that it is natural to attempt to use it in relation to the LCT. However, the underlying assumption—that a signal has finite bandwidth—does not in general lead to the most efficient result, nor is it often the simplest choice. In this section, we will discuss a different assumption, associated with the LCT, which leads to a sampling theorem more general than that of Nyquist. This sampling theorem, which is essential for the theory of a discrete transform, permits more efficient numerical calculations and has physical significance in that it explains higher order diffraction terms in holography [13]. First however, in order to facilitate the reader’s comprehension of this more general sampling theory, let us outline the Shannon–Nyquist sampling theorem, which is associated with the Fourier transform (FT). You will recall that the FT is a special case of the LCT, and so Shannon–Nyquist sampling similarly is a special case of the more general theory.

Suppose we have a signal, $f(x)$, which we want to sample. Shannon–Nyquist sampling starts with the assumption that the FT of $f(x)$, $F(k)$, is zero outside of some finite range of k .

$$F(k) = 0, |k| > \Omega.$$

The width of this finite range of support is then directly related to a sampling rate for the signal, the Nyquist rate. Its reciprocal, T , is called the Nyquist period.

$$r_s = 1/T = 2\Omega.$$

A reconstruction filter is specified which allows recovery of the continuous signal, $f(x)$, from its samples, $f(nT)$ taken at not less than that rate.

$$f(x) = \sum_{n=-\infty}^{\infty} f(nT) \frac{\sin((x - nT))}{\pi(x - nT)}.$$

It is illuminating to think about Shannon–Nyquist sampling in the Fourier domain. Given a signal, $f(x)$, we sample it at a rate T . The resulting spectrum of the sampled signal is as follows:

$$F_s(k) = \sum_{p=-\infty}^{\infty} F(k - pT).$$

Hence, the effect of the sampling operation in the Fourier domain is to create periodic copies of the signal. If the signal had finite support prior to sampling, and if the separation between replicas is sufficiently large, we can obtain the original signal from the copies by filtering. We can interpret the equivalent result for the LCT, which we will examine next, in an identical fashion except that we no longer consider the Fourier domain, but the output domain of the LCT.

8.2.1 Uniform Sampling

Prior to the general LCT case, a number of special cases were obtained. First, Franco Gori derived the special case for the Fresnel transform [4], i.e. a theorem that, given a signal that has compact support in some Fresnel plane, determines a sampling rate such that the signal may be recovered from its samples taken at no less than that rate. In the same paper, Gori also proved that a signal with finite bandwidth could not have finite support in any Fourier domain. The equivalent of this latter result for the LCT is more complicated [5] and depends on the ABCD parameters. Results regarding the compactness and bandlimitedness of different LCTs can be easily understood in terms of the concept of essentially equivalent LCT domains [23]. Later, Xia was the first of several authors to derive the corresponding sampling theorem for the fractional Fourier transform [34].

Ding proved that if a particular LCT of a signal is zero outside some range $|x| < \Omega$, then the signal may be sampled at regular intervals of $T = B/\Omega$ [3]. Ding's

theorem generalizes those of Shannon and Nyquist, Gori, and Xia in the sense that those theorems refer to special cases of the LCT, namely the Fourier, Fresnel and fractional Fourier transforms, respectively. Ding's proof resembles Gori's, though he does not cite Gori. Ding's theorem was later independently derived by several others, including Stern [29], Deng et al. [2] and Li et al. [17]. Stern's proof uses the Poisson formula for the comb function to derive the LCT of a sampled signal. Deng et al.'s approach was to derive the convolution and multiplication theorems for the LCT, and then calculate the LCT of the product of an arbitrary function and a train of delta functions. Li et al.'s proof is similar to Ding's. All of these proofs use techniques from well-known proofs of the Shannon–Nyquist theorem, and ultimately all of them are equivalent.

Similar to the Shannon–Nyquist theorem, Ding's theorem shows that the effect of sampling a signal on the LCT of that signal is to create an infinite number of shifted copies of the signal, with a regular spacing that depends on the sampling rate. The difference is that each replica is modulated, multiplied by a chirp; hence, the LCT of a discrete signal is referred to as 'chirp-periodic'. Again, as with the Shannon–Nyquist theorem, if the LCT of the continuous signal has finite support we can specify a sampling rate that ensures the replicas do not overlap, and it is possible to extract the continuous signal from among the replicas using an appropriate filter.

We now sketch a proof of the theorem. This proof ignores some constants, but the form of the solution is correct. Consider a signal, $f(x)$. Let us assume that $g(x) = f(x) \exp(j\pi ax^2/2b)$ is bandlimited in the Fourier domain, i.e. $G(k) = 0$ if $|k| > \omega$. According to the Shannon–Nyquist theorem, if we sample $g(x)$ (equivalent to sampling $f(x)$ first) with sample period T , its Fourier transform becomes periodic, i.e.

$$G_s(k) = \sum_{p=-\infty}^{\infty} G(k - p/T).$$

If we scale $G_s(k)$ by a factor b , and substitute $y = bk$, we obtain

$$G_s(y) = \sum_{p=-\infty}^{\infty} G(y - p(b/T)).$$

Finally, we multiply this function by a chirp, yielding the LCT of $f(nT)$,

$$F_{M,s}(y) = \exp(j\pi dy^2/2b) \sum_{p=-\infty}^{\infty} G(y - p(b/T)). \quad (8.2)$$

The parameters of this LCT are given by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ d/b & 1 \end{pmatrix} \begin{pmatrix} b & 0 \\ 0 & b^{-1} \end{pmatrix} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ a/b & 1 \end{pmatrix}.$$

In other words, we decomposed the LCT into a series of steps, which are read from right to left as follows: a chirp multiplication, a Fourier transform, magnification or scaling, and a second chirp multiplication. The consequence of Eq. (8.2.1) is that the LCT of a discrete signal is chirp-periodic. The replicas are spaced by b/T , so if the LCT of the original, continuous signal is zero for $|y| > \Omega$, sampling with sample period not less than $T = b/\Omega$, we can filter out the replicas much as in the Fourier case.

The complex parametered case was considered in [19]. By the complex parametered case, we mean not only that the ABCD parameters are complex but so too is the output variable, y . Here, the key result is as follows:

$$\begin{aligned} & \mathcal{L}_M\{\hat{f}(x)\}(y) \\ &= \exp\left(\frac{j\pi dy^2}{2b}\right) \sum_{m=-\infty}^{\infty} F_M\left(y - \frac{mB}{T}\right) \exp\left(-j\pi d\left(y - \frac{mB}{T}\right)/2b\right). \end{aligned}$$

Consequently, if we sample a function with sampling period T , its complex LCT develops regularly spaced, modulated replicas along a line in the complex plane that depends on the b parameter. The spacing of these replicas also depends on the b parameter, and is given by b/T .

The discussion in this section does not account for noise. This issue has been addressed in some recent papers [28, 38]. Neither does it account for quantisation error, which has also been examined recently [25].

8.2.2 Sampling Rate Conversion

A common problem may be stated as follows. Given the samples of a signal which have been taken at a given rate, how do we obtain the samples of the same signal taken at a different rate without damaging the signal? Such sampling rate conversions are often required to downsample signals for computational savings, demodulation and many other applications. Zhao et al. have proposed methods for sampling rate conversions [36] for signals that are bandlimited in some LCT domain. Their work applies to rational number sampling rate changes. Previously published methods for sampling rate conversion in the Fourier domain and the fractional Fourier domain are special cases of their result.

There are two fundamental operations involved: interpolation and downsampling or decimation. Interpolation was defined as follows:

$$g(n) = \begin{cases} f(n/L) & \text{if } \frac{n}{L} = k, k \in \mathbb{Z} \\ 0 & \text{otherwise.} \end{cases}$$

In the LCT domain, this does not alter the signal. Downsampling is defined as follows:

$$g(n) = x(Qn),$$

where Q is the integer downsampling conversion factor. The consequence of this in the LCT domain is as follows:

$$G_T(y) = \frac{1}{Q} \sum_{k=0}^{Q-1} F_T \left(y - \frac{\text{sgn}(b)k}{Q} \right) \exp \left(j2\pi kbd \frac{(\text{sgn}(b)y - \pi k)}{(QT)^2} \right).$$

8.2.3 Nonuniform Sampling

We may have to deal with nonuniform sampling measurements in various settings, e.g. tomography, or due to factors that impinge on our ability to sample regularly, such as jitter, or dropped packets. A number of papers have addressed this issue for a variety of special cases [18, 27, 31, 33, 35]. For example, in [31], Tao et al. discussed two factors that affect the quality of reconstruction. One was truncation, which we will discuss in the following section on the discrete transform. The other is the situation where we have irregularly or nonuniformly sampled data. In this latter case, four special cases are examined in [31].

The first of the four non-uniform cases considered by Tao *et al.* was that of signals for which the locations of the samples are periodic. The second case considered is a special case of the first. The sample locations are again assumed to be periodic, but the locations in each period are assumed to be further divisible into groups with a common inner period. The next case considered is where we nearly have uniform sampling, but it has been degraded by a finite number of samples being taken at incorrect locations.

Finally, a more general case was considered. Given a signal that is bandlimited in some given LCT domain, such that $\mathfrak{L}_T\{f(x)\}(y) = 0$ for $|y| > \Omega$, then the original signal is uniquely determined by its samples taken at locations t_n if

$$\left| t_n - n \frac{b\pi}{\Omega} \right| \leq D < \frac{b\pi}{4\Omega}.$$

The reconstruction formulae for these cases are omitted here, and may be found in [31].

8.3 Exact Relation Between Discrete and Continuous LCTs

We have seen that the LCT of a discrete signal is chirp-periodic. In this section, we will show how to use that fact to define a well-behaved discrete LCT. This section is largely derived from [6, 22].

The obvious way to fashion a discrete transform is simply to sample the input and output domains. The function to be transformed, $f(x)$, becomes a vector of input samples, the n th element of which is given by $f[n] = f(nT_x)$. Similarly, the transformed function becomes the vector F_M , where the m th element is given by $F_M[m] = \mathcal{L}_M\{f(x)\}(mT_y)$. Applying the same substitutions to the LCT kernel allows us to relate the input and output samples by means of a matrix multiplication.

$$F_M = Wf, \quad (8.3)$$

where W is a square matrix with elements,

$$W[n, m] = K \exp\left\{\frac{j\pi}{B}(A(nT_x)^2 + nmT_xT_y + D(mT_y)^2)\right\}.$$

W is obtained by sampling the kernel of the LCT as described in the preceding paragraph. Here, K is some complex constant, which we determine by requiring the discrete transform to be unitary, which concept we now define.

When discussing discrete transforms, one property we are often interested in is whether or not they are unitary. This is often compared with the continuous property of power conservation, and indeed it resembles it, but it fundamentally determines whether or not the transform is invertible. Given a 1-D linear transformation such as that of Eq. (8.3), then unitarity is the following requirement.

$$WW^\dagger = I_N,$$

where I_N is the identity matrix of dimension N , and the \dagger indicates the Hermitian conjugate.

To our knowledge, the first proposed discrete LCT of this kind was by Pei and Ding in [26]. It has since been found that the sampling rates, T_x and T_y , are critical if we want a genuine discrete transform. The idea is elegant: since the continuous LCT of a discrete signal is chirp-periodic, and hence the continuous LCT of a chirp-periodic signal is discrete, then a true discrete LCT transforms a discrete, chirp-periodic signal into another signal of that kind with the underlying periodically replicated functions being related to each other through a continuous LCT [22, 30]. Oktem and Ozaktas showed that such a relation exists with a discrete transform if the number of samples is determined from the knowledge of the signal's extent in the input and output domains, implying a parallelogram-shaped support in phase space [21]. Figure 8.2 demonstrates this using a series of modified PSDs. Healy and Sheridan proposed that it was necessary to include the location of replicas [10] (which neglects the cross-terms, though this is not critical [11]) and later also incorporated the parallelogram bound [6].

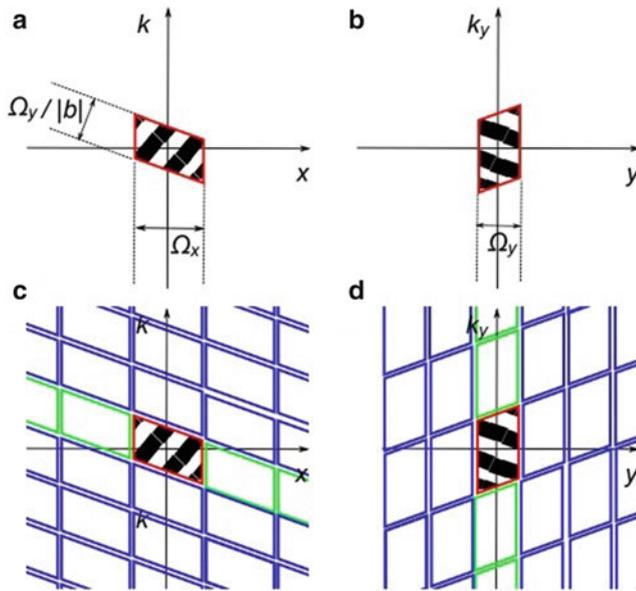


Fig. 8.2 The PSD of a continuous signal that is bandlimited in some LCT domain (a) before and (b) after that LCT. The PSD of a discrete, chirp-periodic signal (c) before and (d) after a discrete LCT. The green replicas alone in (c) constitute a chirp-periodic signal (and are transformed into a discrete signal, exactly periodic in frequency—see the green replicas in (d)), while the blue replicas are a consequence of the signal being discrete

The consequence of this idea of using discrete, chirp-periodic signals is that T_x and T_y should both be chosen using Ding's sampling theorem, i.e.

$$T_x = \frac{\Omega_y}{|b|},$$

and

$$T_y = \frac{\Omega_x}{|b|},$$

where Ω_x and Ω_y are the widths of the signal in the input and output domain, respectively. In practice, the sampling periods are also not independent of one another. Since $\Omega_y = NT_y$,

$$T_x = \frac{N}{|b|} T_y.$$

That is, the three parameters T_x , T_y , and N should be chosen using the preceding three equations together. This choice of sampling rates was first implicitly proposed

in [30] (though without any observation of the consequences). Oktem and Ozaktas later independently proposed this rate, and indicated that this choice provides the minimum possible number of samples to approximate the continuous LCT when finite extents in the input and output domains are assumed [22]. Healy and Sheridan interpreted the results in phase space [6], where chirp-periodic signals form a tiling of parallelograms which packs the replicas optimally, indicating that the sampling rate is optimal under these assumptions. Zhao et al. proved that the transform is unitary, and additionally proved that integer multiples of the Ding sampling rate also result in a unitary transform providing the integer is coprime with the number of samples [37].

A key property of the continuous LCT is the fact that it forms a group, in the sense that

$$\mathcal{L}_{M_2}\{\mathcal{L}_{M_1}\{f(x)\}(y)\}(z) = \mathcal{L}_{M_2M_1}\{f(x)\}(z).$$

For the special case of the Fourier transform, this reduces to a familiar property. The DFT applied twice is equivalent to time reversal, applied three times is equivalent to the inverse DFT, and applied four times is equivalent to the identity transformation. For larger multiples, the pattern repeats, yielding $n \bmod 4$ instances of the DFT. Zhao et al. have recently proven that providing we sample the input and output at a coprime integer multiple of Ding's sampling rate (just as for unitarity), the discrete LCT exhibits the group property similar to the continuous LCT [9, 40].

The literature to date on decomposition based algorithms, e.g. [12, 14, 24] and their two-dimensional and complex extensions [15, 16], has not explicitly used a DLCT, but this discussion may nevertheless have relevance for those algorithms.

The LCT of 2 D signals may often be decomposed into 1D transforms in orthogonal spatial directions, and so has received less attention, but there are cases with independent significance. These 2 D non-separable LCTs have also been shown to be unitary for appropriate sampling rates [39].

8.4 Comparison of Algorithms

We have fulfilled two thirds of our goal: we have a discrete transform and a sampling theorem. However, the discrete LCT is not efficient enough for most purposes. We turn now to algorithms for calculating the discrete transform quickly. This topic is dealt with more comprehensively in Chap. 10 of this book. Here, we focus on the consequences of our new understanding of discrete signals and their LCTs for two such algorithms.

We believe Hennelly and Sheridan were the first to look at systematic ways of accounting for sampling requirements at each stage of LCT algorithms based on decomposition of the matrix of parameters, M , which they described in [12]. Around the same time, Ozaktas et al. demonstrated an Iwasawa-type decomposition

with speed and accuracy comparable to the FFT [24]. Healy and Sheridan later reduced many of the consequences of [12] to two parameters, the well-known space-bandwidth product and a new metric they dubbed the space-bandwidth ratio [8]. However, these results are only accurate in certain circumstances where the replicas due to sampling don't substantially alter the calculations.

One interesting result of the analysis in [8] is that fast LCT algorithms of the kind presented in [6] require at worst the same number of samples as two (and perhaps many more) of the more common algorithms. The algorithm of [6] decomposes the discrete transform matrix into smaller matrices iteratively in the style of the FFT, while the others decompose the matrix of parameters, i.e. decompose the discrete transform into a sequence of simpler discrete transforms. However, in practice, the slowest step in algorithms of the second type are Fourier transforms, and commercial FFT software is fast enough to overcome any advantages of algorithms of the first kind have in terms of the number of samples required.

Let us turn to the consequences of discrete, chirp-periodic signals for the direct and spectral methods [7]. First, let us define these algorithms. The direct method is given by the following decomposition:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ d/b & 1 \end{pmatrix} \begin{pmatrix} b & 0 \\ 0 & b^{-1} \end{pmatrix} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ a/b & 1 \end{pmatrix}.$$

This is the decomposition we used to prove Ding's sampling theorem in Sect. 8.2.1. The spectral method is given by a different decomposition [8].

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ c/a & 1 \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -b/a & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

In both cases, the diagonal matrices are scalings (scaling factor given by the top left character), the anti-diagonals are Fourier transforms (top right entry is 1) or inverse Fourier transforms (top right entry is -1), and the remaining triangular matrices are chirp multiplications with parameter given by the lower left entry.

Before we continue, an aside. In the discussion below, we will consider the discrete LCT to transform one discrete, chirp-periodic signal into another. This can be a source of confusion. The data input to such a discrete transform can be thought of as merely a vector of numbers, or as the samples of a continuous signal under examination. Indeed, the means of obtaining our input data may support one of those interpretations. However, for the purposes of PSD-based analyses of discrete signals, it is useful to think of the vector of numbers as being the samples of a single period of a chirp-periodic signal, and hence of the input as a discrete, chirp-periodic signal. It may be helpful for the reader to consider the special case of the DFT, which can be thought of as relating discrete, periodic functions; necessarily so because the input being discrete means the output is periodic, and the output being discrete means the input is periodic.

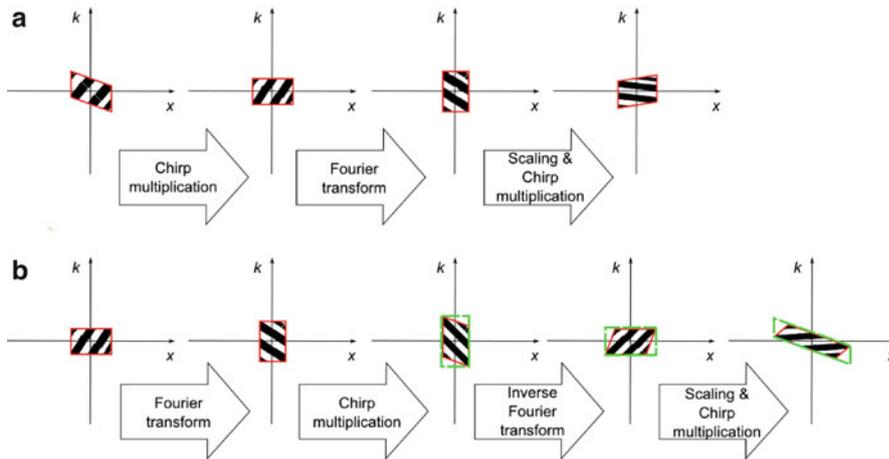


Fig. 8.3 Sequence of PSDs of signal undergoing the stages of (a) the direct method algorithm and (b) the spectral method. Replicas are not shown for convenience. Note the resampling (*green box*) necessary at the third step of the spectral method

Let us consider the direct method first. Given a signal that is discrete and chirp-periodic, the first operation is a chirp multiplication. It is no coincidence that this operation converts the signal into a simple discrete, periodic signal that is suitable for the second operation—a Fourier transform (numerically implemented using an FFT). The signal remains discrete and periodic after the Fourier transform and the third operation—scaling. Finally, the second chirp multiplication converts the signal to chirp-periodic. This sequence is illustrated in Fig. 8.3.

Now let us consider the spectral method. The first operation is a Fourier transform, so we require the signal to be discrete and periodic in order to use the FFT. (How can we require this? We implicitly impose this form on the signal simply by assuming we can sample the input and output at some rates for which aliasing is acceptable.) The resulting signal is also discrete and periodic. The second operation is a chirp multiplication, which converts the signal to being discrete and chirp-periodic. The third operation is a Fourier transform, for which we require a discrete-periodic signal. This mismatch means that we must oversample the signal at this point. Consequently, the spectral method is inherently less efficient than the direct method.

A second mark against the spectral method is that it requires two FFTs to the direct method's one. Why then, does the spectral method remain useful? The two algorithms make different assumptions about the signal. The direct method assumes the signal is bounded in the output LCT domain, while the spectral method assumes it is bounded in the Fourier domain. For certain domains where the ratio a/b is large, such as Fresnel transforms for short distances ($a = 1$, $b = \lambda z$), the signal is much more efficiently bound in the Fourier domain.

There are many other algorithms that decompose the ABCD matrix, and each of these may be analysed in this fashion.

8.5 Conclusion

While by no means an exhaustive treatment of the issues surrounding numerical approximation of LCTs, this chapter has reviewed a number of key ideas. Central to the topic is Ding's sampling theorem, which shows that the LCT of a discrete signal is 'chirp-periodic', and that a signal may be recovered from its samples if it has finite support in some LCT domain such that the periodic replicas do not overlap. The second really critical idea is that a discrete, chirp-periodic signal may be transformed by certain LCTs (those on the manifolds given by certain multiples of a particular ratio of parameters a/b) into discrete, chirp-periodic signals. Completely equivalently, any discrete, chirp-periodic signal may be transformed by LCT into another discrete, chirp-periodically providing the choice of sampling rate is appropriately constrained. However we frame the discussion, the resulting discrete transform exhibits a number of desirable properties including unitarity and additivity. With these properties, the discrete LCT becomes a stable and powerful tool in its own right.

Acknowledgements H.M. Ozaktas acknowledges partial support of the Turkish Academy of Sciences.

References

1. M.J. Bastiaans, Wigner distribution function and its application to first order optics. *J. Opt. Soc. Am.* **69**, 1710–1716 (1979)
2. B. Deng, R. Tao, Y. Wang, Convolution theorems for the linear canonical transform and their applications. *Sci. China (Ser. F Inf. Sci.)* **49**(5), 592–603 (2006)
3. J.J. Ding, Research of fractional Fourier transform and linear canonical transform. Ph.D. thesis, National Taiwan University, Taipei, Taiwan, ROC, 2001
4. F. Gori, Fresnel transform and sampling theorem. *Opt. Commun.* **35**(5), 293–297 (1981)
5. J.J. Healy, J.T. Sheridan, Cases where the linear canonical transform of a signal has compact support or is band-limited. *Opt. Lett.* **33**, 228–473 (2008)
6. J.J. Healy, J.T. Sheridan, Fast linear canonical transforms - FLCTs. *J. Opt. Soc. Am. A* **27**, 21–30 (2010)
7. J.J. Healy, J.T. Sheridan, Reevaluation of the direct method of calculating Fresnel and other linear canonical transforms. *Opt. Lett.* **35**(7), 947–949 (2010)
8. J.J. Healy, J.T. Sheridan, Spacebandwidth ratio as a means of choosing between Fresnel and other linear canonical transform algorithms. *J. Opt. Soc. Am. A* **28**(5), 786–790 (2011)
9. J.J. Healy, K.B. Wolf, Discrete canonical transforms that are Hadamard matrices. *J. Phys. A Math. Theor.* **44**, 265302 (2011)
10. J.J. Healy, B.M. Hennelly, J.T. Sheridan, Additional sampling criterion for the linear canonical transform. *Opt. Lett.* **33**, 2599–2601 (2008)

11. J.J. Healy, W.T. Rhides, J.T. Sheridan, Cross terms of the Wigner distribution function and aliasing in numerical simulations of paraxial optical systems. *Opt. Lett.* **35**(8), 1142–1144 (2010)
12. B.M. Hennelly, J.T. Sheridan, Generalizing, optimizing, and inventing numerical algorithms for the fractional Fourier, Fresnel, and linear canonical transforms. *J. Opt. Soc. Am. A* **22**(5), 917–927 (2005)
13. D.P. Kelly, J.J. Healy, B.M. Hennelly, J.T. Sheridan, Quantifying the 2.5D imaging performance of digital holographic systems. *J. Eur. Opt. Soc. RP* **6**, 11034 (2011)
14. A. Koç, H.M. Ozaktas, C. Candan, M.A. Kutay, Digital computation of linear canonical transforms. *IEEE Trans. Signal Process.* **56**(6), 2383–2394 (2008)
15. A. Koç, H.M. Ozaktas, L. Hesselink, Fast and accurate algorithm for the computation of complex linear canonical transforms. *J. Opt. Soc. Am. A* **27**(9), 1896–1908 (2010)
16. A. Koç, H.M. Ozaktas, L. Hesselink, Fast and accurate computation of two-dimensional non-separable quadratic-phase integrals. *J. Opt. Soc. Am. A* **27**(6), 1288–1302 (2010)
17. B.Z. Li, R. Tao, Y. Wang, New sampling formulae related to linear canonical transform. *Signal Process.* **87**(5), 983–990 (2007)
18. C. Li, B. Li, T. Xu, Approximating bandlimited signals associated with the lct domain from nonuniform samples at unknown locations. *Signal Process.* **92**, 1658–1664 (2012)
19. C. Liu, D. Wang, J.J. Healy, B.M. Hennelly, J.T. Sheridan, M.K. Kim, Digital computation of the complex linear canonical transform. *J. Opt. Soc. Am. A* **28**(7), 1379–1386 (2011)
20. A.W. Lohmann, R.G. Dorsch, D. Mendlovic, Z. Zalevsky, C. Ferreira, Space-bandwidth product of optical signals and systems. *J. Opt. Soc. Am. A* **13**, 470–473 (1996)
21. F. Oktem, Signal representation and recovery under partial information, redundancy, and generalized finite extent constraints. Master's thesis, Bilkent University, 2009
22. F. Oktem, H.M. Ozaktas, Exact relation between continuous and discrete linear canonical transforms. *IEEE Signal Process. Lett.* **16**, 727–730 (2009)
23. F.S. Oktem, H.M. Ozaktas, Equivalence of linear canonical transform domains to fractional fourier domains and the bicanonical width product: a generalization of the space–bandwidth product. *J. Opt. Soc. Am. A* **27**(8), 1885–1895 (2010)
24. H.M. Ozaktas, A. Koç, I. Sari, M.A. Kutay, Efficient computation of quadratic-phase integrals in optics. *Opt. Lett.* **61**(1), 35–37 (2006)
25. A. Ozcelikkale, H.M. Ozaktas, Beyond nyquist sampling: a cost-based approach. *J. Opt. Soc. Am. A* **30**(4), 645–655 (2013)
26. S.C. Pei, J.J. Ding, Closed-form discrete fractional and affine Fourier transforms. *IEEE Trans. Signal Process.* **48**, 1338–1353 (2000)
27. K. Sharma, Approximate signal reconstruction using nonuniform samples in fractional Fourier and linear canonical transform domains. *IEEE Trans. Signal Process.* **57**, 4573–4578 (2009)
28. D. Song, H. Zhao, Stochastic formulation of (a,b,c,d)-bandlimited signal reconstruction. *Circuits Syst. Signal Process.* **34**(6), 2053–2064 (2015). doi:[10.1007/s00034-014-9932-0](https://doi.org/10.1007/s00034-014-9932-0)
29. A. Stern, Sampling of linear canonical transformed signals. *Signal Process.* **86**(7), 1421–1425 (2006)
30. A. Stern, Why is the linear canonical transform so little known?, in *Proceedings of the 5th International Workshop on Information Optics*, ed. by G. Cristóbal, B. Javidi, S. Vallmitjana (Springer, Toledo, 2006), pp. 225–234
31. R. Tao, B.Z. Li, Y. Wang, G.K. Aggrey, On sampling of band-limited signals associated with the linear canonical transform. *IEEE Trans. Signal Process.* **56**, 5454–5464 (2008)
32. M. Testorf, B. Hennelly, J. Ojeda-Castanyeda, *Phase-Space Optics* (McGraw-Hill, New York, 2010)
33. D. Wei, Q. Ran, Y. Li, Reconstruction of band-limited signals from multichannel and periodic nonuniform samples in the linear canonical transform domain. *Opt. Commun.* **284**, 4307–4315 (2011)
34. X.G. Xia, On bandlimited signals with fractional Fourier transform. *IEEE Signal Process. Lett.* **3**(3), 72–74 (1996)

35. Xiao, L., Sun, W.: Sampling theorems for signals periodic in the linear canonical transform domain. *Opt. Commun.* **290**, 14–18 (2013)
36. J. Zhao, R. Tao, Y. Wang, Sampling rate conversion for linear canonical transform. *Signal Process.* **88**(11), 2825–2832 (2008)
37. L. Zhao, J.J. Healy, J.T. Sheridan, Unitary discrete linear canonical transform: analysis and application. *Appl. Opt.* **52**(7), C30–C36 (2013)
38. H. Zhao, R. Wang, D. Song, Recovery of bandlimited signals in linear canonical transform domain from noisy samples. *Circuits Syst. Signal Process.* **33**(6), 1997–2008 (2014)
39. L. Zhao, J.J. Healy, J.T. Sheridan, Two-dimensional nonseparable linear canonical transform: sampling theorem and unitary discretization. *J. Opt. Soc. Am. A* **31**(12), 2631–2641 (2014)
40. L. Zhao, J.J. Healy, C.-L. Guo, J.T. Sheridan, Additive discrete 1D linear canonical transform. *Proc. SPIE* **9599**, 95992E (2015)