Sampling Based on Local Bandwidth

Dennis Wei and Alan V. Oppenheim
Massachusetts Institute of Technology
Digital Signal Processing Group
77 Massachusetts Avenue, Cambridge, MA 02139
dwei@mit.edu, avo@mit.edu

***Abstract*—This paper investigates the sampling of continuous-time signals according to time-varying or local bandwidth. A general framework is established for sampling based on local bandwidth. Two specific representations of local bandwidth are subsequently presented. In the first approach, the time-invariant anti-aliasing filter that typically precedes uniform sampling is generalized to a time-varying lowpass filter. A corresponding sampling and re-synthesis method is developed that achieves perfect reconstruction for a class of self-similar input signals; in most cases, however, perfect reconstruction does not result. In the second approach, local bandwidth is represented in terms of time-warping applied to globally bandlimited signals. The use of time-warping is proposed as a pre-conditioning technique to transform the input signal into one that is as close as possible to being globally bandlimited.**

I. Introduction

Discrete-time representations of continuous-time signals are traditionally obtained by means of uniform sampling based on the Nyquist sampling theorem. More generally, extensions of the Nyquist sampling theorem, such as bandpass sampling and the non-uniform sampling theorem, give rise to alternative sampling methods (see [1] for a general survey). These theorems make reference to the bandwidth of the Fourier transform of the signal, which we refer to as the *global* bandwidth. Many signals of interest, however, have frequency content that is often interpreted, at least informally, as being time-varying. For example, a high frequency carrier modulated in frequency by a low frequency signal is often described in terms of instantaneous frequency. Speech and music are other examples that are perceived and referred to as having time-varying frequency content, while the same is true in the spatial domain for natural images.

Signals with time-varying frequency content are commonly characterized using the short-time Fourier transform or other time-frequency distributions. However, time-frequency analysis techniques are typically not directed at the representation of signals by their samples. In this paper, we seek representations for time-varying bandwidth (which we refer to as *local* bandwidth) in the specific context of sampling and reconstructing signals.

Given the association between globally bandlimited signals and time-invariant lowpass filters, it is natural to consider modelling local bandwidth in terms of time-varying lowpass filters. Accordingly, we develop a representation for local bandwidth and a corresponding structure for sampling and reconstruction based on a class of time-varying lowpass filters, drawing upon work by Horiuchi [2] on signals with bandlimited time-varying spectra.

An alternative approach proposed in [3], [4] represents local bandwidth as the result of local time-scaling or time-warping applied to globally bandlimited signals. In this spirit, [4] developed pre-conditioning filters tailored to pre-specified non-uniform sampling grids. For the more general problem of adapting the sampling grid to the signal at hand, a pre-conditioning technique is suggested in which a time-warping is used to transform the signal into an approximately globally bandlimited signal. To determine an appropriate time-warping, a heuristic iterative algorithm based on time-frequency distributions was developed in [5]. In this paper, we introduce an alternative method based on minimizing the energy of the signal above a specified maximum frequency.

The remainder of the paper is organized as follows: Section II establishes a general framework for sampling and reconstruction based on local bandwidth together with desirable properties of such systems. Two specific realizations of the general structure are then discussed. In Section III, we present a model for local bandwidth based on an appropriate class of time-varying lowpass filters and develop a corresponding method for sampling and reconstructing signals. In Section IV, we turn our attention to the representation of local bandwidth based on time-warping. We present a structure for sampling and reconstruction that incorporates time-warping in the pre-conditioning stage with the goal of minimizing the signal energy above a given maximum frequency.

II. General framework

In this section, we establish a basic structure for sampling and reconstruction based on local bandwidth. Figure 1 illustrates the familiar structure typically used in uniform Nyquist sampling. The input signal $f(t)$ is pre-conditioned using an anti-aliasing lowpass filter, limiting its maximum frequency to the cut-off frequency ω_c, before being sampled at or above the Nyquist rate. The signal $s(t)$ is reconstructed from its samples using bandlimited interpolation.

We generalize the structure in Figure 1 to the case of time-varying bandwidth as shown in Figure 2. The sampling grid is now permitted to be non-uniform, while the pre-conditioning and reconstruction filters become potentially time-varying. The

This work was supported in part by the Siebel Scholars Foundation, the Texas Instruments Leadership University Program, and BAE Systems PO 112991.

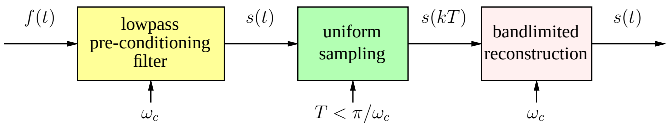

Fig. 1. Structure for uniform Nyquist sampling.

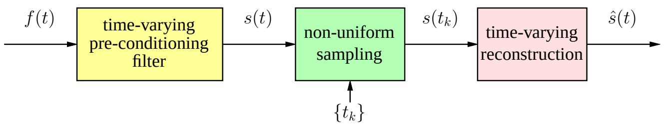

Fig. 2. General structure for sampling based on local bandwidth.

sampling grid and corresponding pre-conditioning are chosen according to the time-varying frequency content of the input signal $f(t)$. Intuitively, we would expect the sampling rate to be higher where $f(t)$ varies more rapidly, and vice versa. As in uniform sampling, the reconstruction filter is ideally designed to re-synthesize the conditioned signal $s(t)$.

Systems designed for sampling according to local bandwidth should ideally possess certain desirable properties. Most importantly, given any input signal, the pre-conditioning filter should yield an output that can be perfectly reconstructed from its samples, i.e., $\hat{s}(t) = s(t)$. We refer to pre-conditioned signals satisfying the perfect reconstruction condition as locally bandlimited signals, as they play a role analogous to globally bandlimited signals in the Nyquist sampling theorem.

Even if perfect reconstruction cannot always be achieved, a requirement often proposed in sampling theory (e.g. [6]) is that of consistent resampling, i.e., $\hat{s}(t_k) = s(t_k)\ \forall\ k \in \mathbf{Z}$. Furthermore, in the limit of constant local bandwidth, the system should reduce to the uniform sampling structure of Figure 1, which does satisfy the perfect reconstruction property. Accordingly, as the local bandwidth approaches a constant, the sampling grid should approach a uniform grid, and the reconstruction error should gradually decrease.

In the next section, we examine a specific realization of the structure in Figure 2 in which the pre-conditioning filter is chosen to be a time-varying lowpass filter. In Section IV, we discuss a system incorporating time-warping in the pre-conditioning stage.

III. Representation Based on Time-Varying Lowpass Filters

In this section, we consider a natural extension of the uniform sampling structure in Figure 1 to the case of time-varying bandwidth. As shown in Figure 3, the time-invariant lowpass filter in the pre-conditioning stage is replaced by a time-varying lowpass filter parameterized by a cut-off frequency $\omega_c(t)$. The function $\omega_c(t)$ is restricted to be continuous and positive with finite minimum and maximum values ω_{min} and ω_{max}, i.e., $0 < \omega_{min} \leq \omega_c(t) \leq \omega_{max} < \infty\ \forall\ t$. The sampling grid and reconstruction filter are then specified in terms of $\omega_c(t)$.

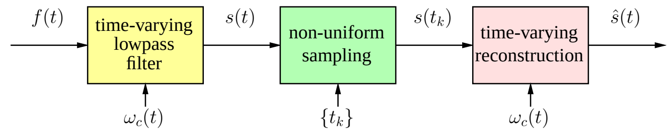

Fig. 3. Structure employing a time-varying lowpass filter in the pre-conditioning stage.

We first consider the definition of the time-varying lowpass filter. There are a variety of ways in which a time-invariant lowpass filter can be generalized to the time-varying case. Our definition is based on the concept of the time-varying frequency response $H(t,\omega)$ for linear time-varying systems, defined as follows: When the input to the system is of the form $f(t) = e^{j\omega t}$, the output is given by

$$s(t) = H(t,\omega)e^{j\omega t}. \tag{1}$$

The time-varying frequency response completely characterizes a linear time-varying system in terms of its responses to complex exponential inputs (a more detailed discussion can be found in [7]). For a time-varying lowpass filter, an intuitively satisfying choice of frequency response is given by

$$H(t,\omega) = \begin{cases} 1, & |\omega| < \omega_c(t), \\ 0, & |\omega| > \omega_c(t). \end{cases} \tag{2}$$

With $f(t) = e^{j\omega_0 t}$ and applying (1) and (2),

$$s(t) = \begin{cases} e^{j\omega_0 t}, & \omega_c(t) > |\omega_0|, \\ 0, & \omega_c(t) < |\omega_0|, \end{cases}$$

i.e., the filter exactly reproduces a complex exponential input when the input frequency is below the cut-off frequency, and produces an output of zero otherwise. The corresponding input-output relation in the time domain is as follows:

$$s(t) = \int_{-\infty}^{\infty} f(\tau)\frac{\sin[\omega_c(t)(t-\tau)]}{\pi(t-\tau)}d\tau. \tag{3}$$

In general, the output of a lowpass filter defined according to (1) and (2) is not globally bandlimited. However, if $\omega_c(t)$ is constant, equations (2) and (3) reduce to the input-output relations of an ideal time-invariant lowpass filter, and hence in this case the output is globally bandlimited. It can be shown that as the range of $\omega_c(t)$ decreases to zero, the output gradually approaches a globally bandlimited signal [8].

The time-varying spectrum $\widetilde{S}(t,\omega)$ of the output $s(t)$ is defined as $\widetilde{S}(t,\omega) = H(t,\omega)F(\omega)$, where $F(\omega)$ denotes the Fourier transform of the input $f(t)$. Since (2) implies that $\widetilde{S}(t,\omega)$ is non-zero only on the finite interval $-\omega_c(t) < \omega < \omega_c(t)$ for each value of t, we refer to $s(t)$ as having a bandlimited time-varying spectrum. Consequently, $s(t)$ may be expressed in terms of the following series expansion derived in [2]:

$$s(t) = \sum_{k=-\infty}^{\infty} \widetilde{s}\left(t, \frac{k\pi}{\omega_c(t)}\right)\varphi_k(t), \tag{4}$$

where

$$\varphi_k(t) = \frac{\sin[\omega_c(t)t - k\pi]}{\omega_c(t)t - k\pi}. \tag{5}$$

The 2-D signal $\widetilde{s}(t,\tau)$ is the inverse Fourier transform of $\widetilde{S}(t,\omega)$ with respect to ω, i.e.,

$$\widetilde{s}(t,\tau) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \widetilde{S}(t,\omega)e^{j\omega\tau}d\omega, \tag{6}$$

and is related to the 1-D signal $s(t)$ by the equation $s(t) = \widetilde{s}(t,t)$.

Based on (4), the reconstruction filter in Figure 3 is chosen so as to yield

$$\hat{s}(t) = \sum_{k=-\infty}^{\infty} s(t_k)\varphi_k(t), \tag{7}$$

where the sampling times t_k are the solutions to

$$\omega_c(t_k)t_k - k\pi = 0, \quad k \in \mathbf{Z}. \tag{8}$$

With this choice of sampling times, the value $s(t_k)$ of the kth sample and the kth coefficient in (4) are equal at $t = t_k$.

Equations (5), (7) and (8) together imply the consistent resampling property since

$$\varphi_k(t_m) = \begin{cases} 1, & k = m, \\ 0, & k \neq m. \end{cases} \tag{9}$$

In addtion, the sampling rate implied by (8) depends on the cut-off frequency $\omega_c(t)$ as intuition would suggest, being higher where $\omega_c(t)$ is large, and vice versa. In the limit as $\omega_c(t)$ approaches a constant, both the sampling times in (8) and the reconstruction in (7) reduce to the uniform sampling and reconstruction of Figure 1.

We note that (8) always has at least one solution for every value of k, since $\omega_c(t)$ is continuous and has a strictly positive lower bound. For those values of k for which there exist multiple solutions to (8), only one of the solutions is chosen to be the kth sampling time.

While the sampling and reconstruction system specified by (7) and (8) ensures consistent resampling, it does not necessarily guarantee perfect reconstruction. The condition for perfect reconstruction can be expressed as follows:

$$\sum_{k=-\infty}^{\infty} \left[\widetilde{s}\left(t, \frac{k\pi}{\omega_c(t)}\right) - s(t_k)\right]\varphi_k(t) = 0. \tag{10}$$

A notable case in which perfect reconstruction does result is for a class of self-similar input signals (defined e.g. in [9]). Specifically, if $F(\omega)$ is of the form

$$F(\omega) = \frac{C}{|\omega|^\gamma}, \tag{11}$$

with $\gamma = 1$, then the quantity in square brackets in (10) vanishes and we have $\hat{s}(t) = s(t)$. More generally, perfect reconstruction can be achieved for any value of γ in (11) by modifying the reconstruction formula in (7) as follows:

$$s(t) = (\omega_c(t))^{1-\gamma} \sum_{k=-\infty}^{\infty} (\omega_c(t_k))^{\gamma-1} s(t_k)\varphi_k(t). \tag{12}$$

In general, however, (10) is not satisfied by many potential inputs to the system of Figure 3, and perfect reconstruction does not result.

The sampling and reconstruction system defined by (7) and (8) does have the property that the reconstruction error decreases as the variation in the cut-off frequency $\omega_c(t)$ becomes more gradual. We define the coefficients $e_k(t)$ of the error signal $e(t) = \hat{s}(t) - s(t)$ by means of an expansion similar to (4):

$$e(t) = \sum_{k=-\infty}^{\infty} e_k(t)\varphi_k(t). \tag{13}$$

The consistent resampling property ensures that $e_k(t_k) = 0$ for all k. For $t \neq t_k$, the following bound applies [8]:

$$|e_k(t)| \leq \frac{1}{\pi}\left\{\int_{\omega_c(t)}^{\omega_c(t_k)} |F(\omega)|d\omega + \lambda_k(t)\sqrt{2\omega_c(t)}\left(\int_0^{\omega_c(t)} |F(\omega)|^2 d\omega\right)^{1/2}\right\}, \tag{14}$$

where

$$\lambda_k(t) = \left(1 - \frac{\sin[\omega_c(t)t_k - k\pi]}{\omega_c(t)t_k - k\pi}\right)^{1/2}.$$

As the range over which $\omega_c(t)$ varies converges to zero, so too does the upper bound in (14).

IV. Representation based on time-warping

In this section, we focus on an alternative representation of local bandwidth based on the time-warping of globally bandlimited signals. We first review the definition of a locally bandlimited signal due to [3], [4] before presenting a structure for sampling and reconstruction.

Consider a signal $f(t)$ for which there exists a continuous, invertible function $\alpha(t)$ such that the signal

$$g(t) = f(\alpha(t)) \tag{15}$$

is bandlimited to ω_0, i.e., ω_0 is the lowest frequency such that $G(\omega) = 0$ for all $|\omega| > \omega_0$. We refer to the transformation $t \to \alpha(t)$ as a time-warping, and to $\alpha(t)$ as the warping function. Equivalently, $f(t)$ may be viewed as the result of applying the inverse time-warping to a globally bandlimited signal $g(t)$, i.e,

$$f(t) = g(\gamma(t)), \tag{16}$$

where $\gamma(t)$ denotes the inverse of $\alpha(t)$. We refer to signals of the form in (16) as locally bandlimited signals since they can be sampled according to local bandwidth (implied in this case by the warping function) and perfectly reconstructed from those samples. The sampling times and reconstruction are given by

$$t_k = \alpha(kT), \tag{17}$$

and

$$f(t) = \sum_{k=-\infty}^{\infty} f(t_k)\frac{\sin(\omega_0\gamma(t) - k\pi)}{\omega_0\gamma(t) - k\pi}. \tag{18}$$

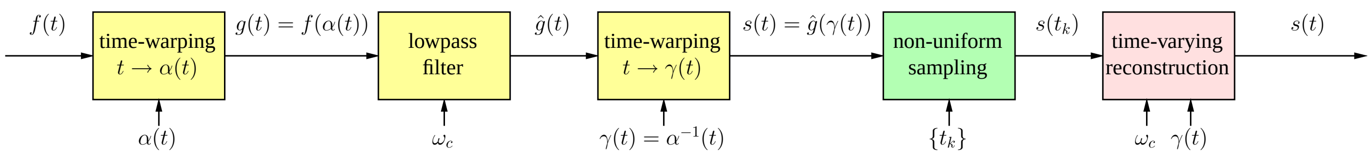

Fig. 4. Structure employing time-warping in the pre-conditioning stage.

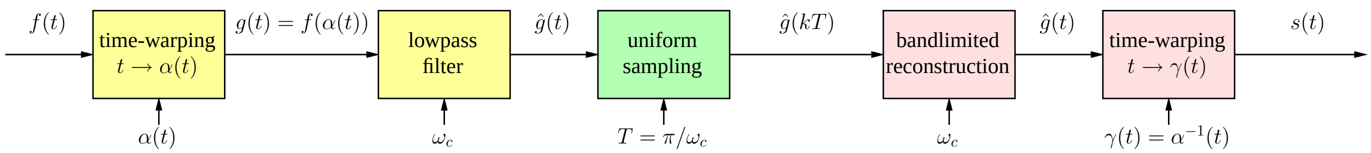

Fig. 5. A modified version of the structure in Figure 4.

In order for the maximum frequency ω_0 to be uniquely specified, the warping functions $\alpha(t)$ and $\gamma(t)$ are normalized to exclude any global time scaling by the following constraint:

$$|\alpha(t) - t| \leq B \quad \forall\, t, \tag{19}$$

for some positive constant B. Equation (19) and the requirement of invertibility constrain $\alpha(t)$ and $\gamma(t)$ to be monotonically increasing functions of t. Thus, we may interpret the time-warping as a local time-scaling that gives rise to time-varying bandwidth.

Based on the preceding definition of a locally bandlimited signal, we now consider the system for sampling and reconstruction depicted in Figure 4. The first step in the pre-conditioning consists of a time-warping, chosen such that the result $g(t)$ is as close as possible in some sense to being globally bandlimited. As a consequence, the loss through the anti-aliasing lowpass filter is minimized. Intuitively, the time-warping serves to counteract the time-varying bandwidth of $f(t)$, yielding an output with nearly constant bandwidth.

Conceptually, the output $\hat{g}(t)$ of the lowpass filter should be subjected to the inverse of the initial time-warping, i.e., that specified by the inverse warping function $\gamma(t)$, to yield a locally bandlimited signal $s(t)$ as the final output of the pre-conditioning. The signal $s(t)$ could then be sampled and perfectly reconstructed according to (17) and (18). The structure in Figure 4 can be simplified practically, however, by combining the operations of time-warping and sampling, and by decomposing the reconstruction filter into a time-invariant lowpass filter followed by a time-warping. Figure 5 shows the resulting structure, the middle part of which now corresponds to the uniform sampling structure of Figure 1.

To determine an appropriate time-warping $\alpha(t)$, Brueller et al. [5] have proposed a heuristic algorithm based on the thresholding of time-frequency distributions. We focus instead on an alternative strategy in which the objective is to minimize the ratio of the energy of $g(t)$ above ω_c to the total energy of $g(t)$. Specifically,

$$\alpha_*(t) = \arg\min_{\alpha(t)} r = \arg\min_{\alpha(t)} \frac{\displaystyle\int_{|\omega|=\omega_c}^{\infty} |G(\omega)|^2 d\omega}{\displaystyle\int_{-\infty}^{\infty} |G(\omega)|^2 d\omega}, \tag{20}$$

where $\alpha(t)$ ranges over all continuous invertible functions satisfying (19). The maximum frequency ω_c is determined by the sampling period T permitted by the system. We note that, under the assumption that $f(t)$ contains an infinite number of zeroes, minimizing the energy above ω_c without normalizing by the total energy leads to a degenerate solution in which the total energy is driven to zero.

As a preliminary approach to the optimization problem in (20), we consider a simplification in which $f(t)$ is assumed to be non-zero only on the finite interval $[0, L]$, and is represented by uniform samples over this interval at a sufficiently high rate for the error due to aliasing to be negligible. Furthermore, we restrict attention to approximations of the optimal warping function in the following class of piecewise linear functions:

$$\alpha(t) = \sum_{p=0}^{P-1} \Delta_p B\left(\frac{t}{\delta} - p\right), \quad 0 \leq t \leq L, \tag{21}$$

where

$$B(t) = \begin{cases} 0, & t < 0, \\ t, & 0 \leq t \leq 1, \\ 1, & t > 1, \end{cases} \tag{22}$$

δ is the horizontal length of every linear segment, and $P\delta = L$. The parameters Δ_p represent the increases in $\alpha(t)$ over each segment, and satisfy $\Delta_p > \epsilon > 0$ to ensure invertibility, and $\sum_p \Delta_p = L$, the finite-length analogue to (19). The operation of time-warping is approximated by interpolating between the samples of $f(t)$ to yield uniform samples of $g(t)$. Under these assumptions, the ratio r in (20) can be approximated in terms of the discrete Fourier transform of samples of $g(t)$, and the minimization of r with respect to the parameters Δ_p can be performed using a gradient-descent algorithm.

The gradient-descent algorithm was evaluated using time-warped sinusoidal signals, which fall into the class of locally bandlimited signals under this representation. Thus, the algorithm is expected to yield the inverse of the time-warping applied to the underlying sinusoid. However, simulation results show that the gradient-descent algorithm suffers from sensitivity to the initial values for the parameters Δ_p. The algorithm converges to a local minimum of r near the initial conditions, which typically does not correspond to the desired inverse time-warping. Future work will consider more sophisticated initialization and/or optimization procedures for the problem in (20).

V. Conclusion

We have presented in this paper a framework for sampling and reconstructing signals on the basis of local bandwidth. Two specific representations for local bandwidth were examined in detail. The first approach based on time-varying lowpass filters led to a sampling and reconstruction system that does not satisfy the property of perfect reconstruction for most input signals. In contrast, the second approach based on time-warping addresses the desire for perfect reconstruction, while leaving the determination of the optimal time-warping as an area for further investigation.

References

[1] A. J. Jerri, "The Shannon sampling theorem – its various extensions and applications – a tutorial review," *Proc. IEEE*, vol. 65, no. 11, pp. 1565–1596, November 1977.

[2] K. Horiuchi, "Sampling principle for continuous signals with time-varying bands," *Information and Control*, vol. 13, no. 1, pp. 53–61, 1968.

[3] J. J. Clark, M. R. Palmer, and P. D. Lawrence, "A transformation method for the reconstruction of functions from nonuniformly spaced samples," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 33, no. 4, pp. 1151–1165, October 1985.

[4] Y. Y. Zeevi and E. Shlomot, "Nonuniform sampling and antialiasing in image representation," *IEEE Trans. Signal Processing*, vol. 41, no. 3, pp. 1223–1236, March 1993.

[5] N. N. Brueller, N. Peterfreund, and M. Porat, "Optimal sampling and instantaneous bandwidth estimation," in *Proceedings of the IEEE International Symposium on Time-Frequency and Time-Scale Analysis*, vol. 1, Pittsburgh, USA, October 1998, pp. 113–115.

[6] M. Unser and A. Aldroubi, "A general sampling theory for non-ideal acquisition devices," *IEEE Trans. Signal Processing*, vol. 42, pp. 2915–2925, November 1994.

[7] T. A. C. M. Claasen and W. F. G. Mecklenbrauker, "On stationary linear time-varying systems," *IEEE Trans. Circuits Syst.*, vol. 29, no. 3, pp. 169–184, March 1982.

[8] D. Wei, "Sampling based on local bandwidth," Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, December 2006.

[9] G. W. Wornell, *Signal Processing with Fractals: A Wavelet-Based Approach*. Upper Saddle River, NJ: Prentice-Hall, Inc., 1996.