COMPRESSIVE SAMPLING OF LIDAR: FULL-WAVEFORMS AS SIGNALS OF FINITE RATE OF INNOVATION

Juan Castorena, Charles D. Creusere

New Mexico State University
Klipsch School of Electrical and Computer Engineering
Las Cruces, NM 88003

ABSTRACT

The 3D imaging community has begun a transition to full-waveform (FW) LIDAR systems which image a scene by emitting laser pulses in a particular direction and capturing the entire temporal envelope of each echo. By scanning a region, connected 1D profile waveforms of the 3D scenes can be readily obtained. In general, FW systems capture more detailed physical information and characteristic properties of the 3D scenes versus conventional 1st and 2nd generation LIDARs which simply store clouds of range points. Unfortunately, the collected datasets are very large, making tasks like processing, storage, and transmission far more resource-intensive. Current compression approaches addressing these issues rely on collecting large amounts of data and then analyzing it to identify perceptual and statistical redundancies which are subsequently removed. Collecting large amounts of data just to discard most of it is highly inefficiently. Our approach to LIDAR compression models FW return pulses as signals with finite rate of innovation (FRI). We show in this paper that sampling can be performed at the rate of innovation while still achieving good quality reconstruction. Specifically, we show that efficient sampling and compression can be achieved on *actual* LIDAR FW's within the FRI framework.

Index Terms— LIDAR, full-waveform, finite rate of innovation, compressive sampling, sparsity.

1. INTRODUCTION

In general, third generation pulsed LIDAR systems project an energy pulse into a scene and measure the reflected full waveform (FW) signal. Each of these reflected signals provides range measurements of the objects intercepted by the laser pulse along a specified direction. By scanning through a specified region using a series of emitted pulses and observing their reflected FW signals, connected 1D profiles of 3D scenes can be readily obtained. The advantages of 3rd generation LIDAR over conventional 1st and 2nd generation systems that store only 3D range point clouds is that the shapes of the waveforms provide additional insight into the surface structure which can provide improved characterization and classification. Unfortunately, massive amounts of data have to be collected to obtain detailed topographical information about a scene which in turn poses problems for processing, storage and transmission. To resolve these issues, a number of compression approaches have been developed in the literature. They generally require, however, the initial acquisition of large amounts of data only to later discard most of it by exploiting redundancies [1], thus sampling very inefficiently. Our main goal is thus to apply efficient and effective compressive sampling algorithms that can be implemented with low computational complexity on FW LIDAR systems.

During the last few years, the theory of compressive sampling was introduced, reminding us all that the Nyquist rate is a sufficient but not a necessary condition for signal reconstruction. Within this framework, a signal that is inherently sparse in some domain can be sampled very efficiently with rates depending on the sparsity of the signal. Here, we explore this idea by considering FW echoes as signals with finite rates of innovation; thus, our approach compresses a scene by compressively sampling temporal FW signals. In 2010 [2] we applied this concept to simulated FW LIDAR signals to determine the complexity of received waveforms. In the process, we found that very precise FW reconstructions can be achieved using a sampling rate that depended on the FW signal innovations. However, the previously proposed model did not account for non-symmetrical waveform returns typically present in large footprint pulsed LIDAR. Here, we extend the work in [2] by incorporating a model to account for non-symmetrical waveform shapes and use it to compressively sample and reconstruct actual FW LIDAR signals.

2. LIDAR DATA

2.1. Full-waveform signal descriptions

Third generation LIDAR systems typically emit a pulse towards a particular location to extract the ranges by measuring the return echoes. This emitted pulse may cause multiple echoes from intercepting objects laying along the conical 3D volume traversed by the pulse. By sampling the echoes, which we also refer to in here as modes, one extracts a record of the interactions of the emitted pulse with the intercepted

The authors would like to thank AlanVan Nevel and Dean Cook from NAVAIR China Lake, CA for providing the LIDAR waveforms used here.

objects. We refer to the temporal envelope of the complete echo return pulse as the full-waveform (FW) signal.

Each of the modes appearing in the FW signals corresponds to reflections from a single object or from a superposition of several intercepted objects. The shape of each of these echoes may not be the same as that of the emitted pulse. For example, a mode might be skewed and/or its width elongated when the pulse hits a surface whose orientation is not normal to the flight direction of the pulse. In addition, FW's generated by large footprint pulses are more likely to exhibit non-symmetrical and more complex mode shapes in comparison to those corresponding to smaller footprint areas [3].

The problems of compression or processing (e.g., classification) generally require the models for the FW signals. Optimal waveform modeling is not straightforward, however. The number of modes is unknown, and a given mode's shape might be controlled by the sum of echoes from multiple objects. Typical approaches that have been proposed are based on Gaussian curve fitting [4]. Unfortunately, the symmetry of Gaussian curves does not preserve the shape of the skewed and non-symmetrical modes that may be present in the FW echoes. In [3] the authors proposed the use of a marked point process to fit generalized Gaussian, Nakagami and Burr curves since it can model symmetric and skewed modes. Unfortunately, the larger set of curves available increases the number of parameters that needs to be found, thus increasing the complexity of the computation. For this reasons, we instead apply an FRI-based model that approximates FW signals with high quality and yet can be applied efficiently with low computational complexity.

2.2. LIDAR dataset

The LIDAR dataset obtained from NAVAIR China Lake, CA was collected using the VISSTA ELT LADAR system. The shot rate of the system is of 20 Khz (i.e, pulse emission rate). Each time a pulse of 1.5 ns duration at full width half maximum is transmitted. Echoes are measured at a sampling rate of 2Ghz and quantized using an 8-bit A/D converter. The used dataset was collected by imaging a pickup truck through a chain link fence, both positioned perpendicular to the pulse transmission path. To illustrate this more clearly, Figure 1 shows the 3D point cloud obtained by processing the FW signals collected by the system.

Fig. 1. An example of the LIDAR point cloud.

Table 1. Non-uniform linear spline approximation

Number of waveforms	NMSE	NMAE
1000	1e-3	3.87e-4

3. FULL-WAVEFORM MODEL

Based on the aforementioned descriptions, we model FW signals as the sum of a bandlimited and a non-uniform linear spline. The bandlimited component of the signal describes the general signal level while the non-uniform linear spline models the multiple return modes that result when objects are intercepted by the laser pulse. We chose this model to preserve the shape of FW signals, a matter of significant importance to the analyst in determining the physical properties and characteristics of the intercepted objects.

To obtain the nonuniform linear spline approximation of the modes, we first determine the parts of the signal with amplitudes larger than an appropriate predefined threshold value. This value is selected so that FW noise is not considered as false modes. After the peaks of the modes are selected, we then determine the support of each mode by detecting the sample numbers at which the FW signal experiences a zero crossings. The resulting segments of the signal are considered to be the modes and are approximated using non-uniform linear splines. The order of the spline is selected to be one which we have found provides an approximated signal of sufficiently high quality. Splines of higher order such as quadratic and cubic could also be selected, however, but we have not done so here because the approximation gain achieved was not significant. The bandlimited part of the signal is determined by restricting the signal to a predefined bandwidth and has been selected based on observations of the FW signals.

To illustrate the performance of this process, we approximate a FW signal and plot the result in Figure 2a where the entire signal is made positive. Figures 2b and 2c are the same approximations zoomed in around the two extracted modes. Note that the model preserves the shape of the signal with good perceptual quality. The residual noise resulting from the subtraction of the approximation from the FW signal resembles white noise. Thus, FW signal is effectively modeled by a bandlimited plus linear spline signal buried in white noise. To validate our model, we apply it to a total of one thousand FW signals and compute the resulting overall normalized mean squared error (NMSE) and the overall normalized maximum absolute error (NMAE). These metrics are computed only over the portions of the FW signal considered as modes. Normalizations are applied to these metrics such that the highest possible values are one. The results are shown in table 1. Note that the resulting errors are overall very small, thus validating its usage for modeling FW signals.

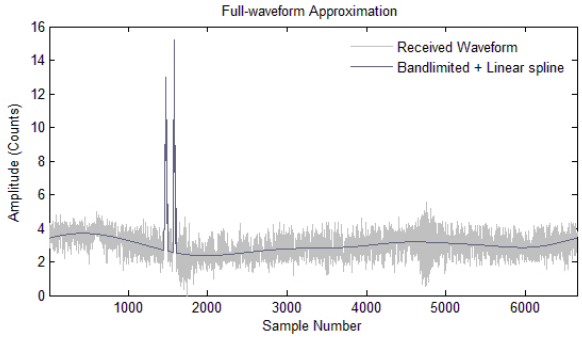

(a) Full-waveform approximation

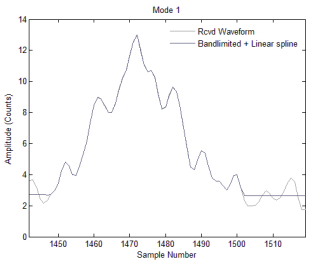

(b) Mode 1

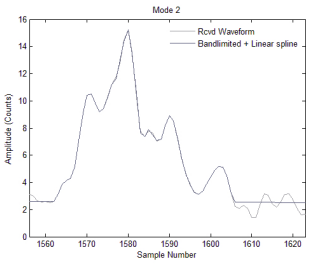

(c) Mode 2

Fig. 2. Full-waveform linear spline approximation.

4. FULL-WAVEFORMS AS A SIGNAL OF FINITE RATE OF INNOVATION

4.1. Review of signals with Finite Rates of Innovation

The concept of FRI developed in [5] essentially states that signals with a finite number of degrees of freedom $2K$ per unit of time τ can be sampled efficiently at a rate greater than or equal to the rate of innovation $\rho = 2K/\tau$. At the heart of this framework are signals of length τ that can be represented by the a_k weighted stream of K Diracs given by:

$$x(t) = \sum_{k=1}^{K} = a_k \delta(t - t_k). \tag{1}$$

The non-bandlimited signal $x(t)$ is fully specified by K locations t_k and weights a_k. As such, if we can somehow observe information relating to these signal parameters, we can then sample very efficiently at a rate much lower than the Nyquist rate and still achieve high quality reconstruction. The method of retrieving the signal parameters starts by observing that the Discrete Fourier transform (DFT) coefficients of (1) are given by

$$X[m] = \frac{1}{\tau} \sum_{k=1}^{K} a_k \exp^{-j2\pi m t_k/\tau}. \tag{2}$$

Observe that with the form of (2), all the $2K$ signal parameters required to completely determine (1) can be determined with just $2K + 1$ contiguous DFT coefficients $X[m]$. This system of equations is nonlinear in the coefficients t_k and can be solved using the annihilating filer method [5]. In general, the annihilating filter with coefficients $H[m]$ for $|m| \leq K$ is one that annihilates K exponentials (i.e., $H[m] * X[m] = 0$). Under this constraint, the zeros of the filter coefficients which occur at $e^{j2\pi m t_k/\tau}$ determine uniquely (under an appropriate transformation) the locations t_k. The fact that only $2K + 1$ contiguous DFT coefficients $X[m]$ are required enables one to sample uniformly using the sinc function of bandwidth B to obtain the measurements

$$y_n = \langle x(t), \mathrm{sinc}(B(nT - t)) \rangle \tag{3}$$

for $n = 1, 2, ..., N$, with DFT coefficients

$$Y[m] = \begin{cases} \tau X[m] & \text{if } |m| \leq \lfloor B\tau/2 \rfloor \\ 0 & \text{for other } m \in [-N/2, N/2]. \end{cases} \tag{4}$$

Note that with a bandwidth of $B \geq \rho$, all of the required coefficients $X[m]$ are still kept by using only $2K+1$ samples y_n and can thus be used to find locations t_k. Using (4) now changes the annihilating filter convolution to $H[m] * Y[m] = 0$. This operation can be recasted by arranging the coefficients $Y[m]$ into a Toeplitz matrix as

$$A = \begin{bmatrix} Y[-L_1] & Y[-L_1-1] & \cdots & Y[-L] \\ Y[L_1+1] & Y[-L_1] & \cdots & Y[-L+1] \\ \vdots & \vdots & \vdots & \vdots \\ Y[L-1] & Y[L+2] & \cdots & Y[-L_1-1] \end{bmatrix} \tag{5}$$

and computing the vector H of DFT coefficients $H[m]$ for which $A \cdot H = 0$. The coefficient index in (5) is chosen to fit the selected FW signal model, where $L = L_1 + L_2$. In the noiseless case, one can always chose values of $L_1 = 0$ and $L_2 = K$ to retrieve the locations t_k. We leave these quantities as variables at this point and define their requirements for the reconstruction of signals modeled by (7) buried in noise in the following subsections. Once the t_k's have been found, the weights a_k can be obtained by using K of the sample values obtained from equation (3) to construct a Vandermonde matrix and applying least squares.

4.2. Non-uniform splines

A signal $x(t)$ is a non-uniform spline of degree R with knots at t_k if its $(R+1)$th derivative is given by the stream of Diracs $x^{(R+1)}(t) = \sum_{k=1}^{K} = a_k \delta(t - t_k)$ with DFT coefficients

$$X^{R+1}[m] = \left(\frac{j2\pi m}{\tau} \right)^{R+1} X[m], m \in \mathbb{Z} \tag{6}$$

In this form, non-uniform splines of degree R can be sampled at or greater than the rate of innovation using the same procedure described in section 4.1 for streams of Diracs to obtain an accurate reconstruction of the signal $x^{(R+1)}(t)$. After recovery of this signal is achieved, a simple $R + 1$ times integration suffices to reconstruct $x(t)$.

4.3. Cadsow's denoising

Ar low SNR, Cadsow's denoising can be applied for improve performance. In general, this algorithm takes advantage of the fact that sampling at a rate higher than the rate of innovation results in a matrix A of dimensions larger than the rank K. As such, we can force the smallest $L_2 - K$ singular values to zero, obtain the corresponding low rank approximation A' and average the coefficients along the diagonals to reconstruct the Toeplitz matrix. This process is then repeated iteratively until the ratio of the $K+1$ and K singular values is less than some threshold value. The number of iterations required to denoise is usually small (less than ten) [6].

4.4. Full-waveforms as FRI signals

We adopt the model described in section 3 to approximate FW signals. Thus, in general the continuous signal $x(t)$ is modeled as

$$x(t) = x_{bl}(t) + x_{modes}(t) \tag{7}$$

where $x_{bl}(t)$ is the bandlimited portion and $x_{modes}(t)$ is the part pertaining to the modes of the FW signal. Sampling and reconstruction of signals that can be modeled by a bandlimited plus a non-uniform splines was proposed in [7] within the FRI framework. Here, we follow their procedure to obtain reconstructions at or above the rate of innovation. We begin by computing the DFT coefficients of (7) given by

$$X[m] = \begin{cases} X_{bl}[m] + X_{modes}[m] & \text{if } |m| \leq L_1 \\ X_{modes}[m] & \text{if } L_1 < |m| \leq L \end{cases} \tag{8}$$

where L_1 is the bandwidth of $x_{bl}(t)$, $L_2 \geq K$ is the number of knots used for the spline approximation and $L = L_1 + L_2$. The algorithm for sampling and reconstruction of FW signals can be then summarized as follows:

- Sample $x(t)$ at a rate of $2L/\tau$ and find $2L$ contiguous DFT coefficients $Y[m]$.
- Use the coefficients in (8) corresponding to $L_1 < |m| \leq L$ to construct the Toeplitz matrix given in (5).
- Apply Cadzow's denoising if low SNR.
- Obtain the locations t_k and weights a_k of $x_{modes}^{(R+1)}(t)$ by applying the annihilating filter method.
- Integrate $x_{modes}^{(R+1)}(t)$ $R+1$ times to recover $x_{modes}(t)$.
- Compute $X_{bl}[m] = X[m] - X_{modes}[m]$ for $|m| \leq L_1$ to obtain $x_{bl}(t)$ and use equation (7) to recover $x(t)$.

5. RESULTS AND DISCUSSION

We measure the performance of our algorithm on several FW signals acquired using the system described in section 2.2. Because the important portion of the FW signal for the analyst is $x_{modes}(t)$, we measure performance only over the support of the modes. One of the metrics used to asses the performance is the NMSE given by

$$\text{NMSE} = \frac{1}{\gamma_1} \|x_{modes}[n] - \widehat{x}_{modes}[n]\|_2^2 \tag{9}$$

where $\gamma_1 = |\operatorname{supp}(x_{modes}[n])| \cdot \|x_{modes}[n]\|_2^2$ and $\widehat{x}_{modes}[n]$ is the FRI estimation. In addition, the NMAE given by

$$\text{NMAE} = \frac{1}{\gamma_2} \max(|x_{modes}[n] - \widehat{x}_{modes}[n]|) \tag{10}$$

where $\gamma_2 = \max(x_{modes}[n])$ is also used as a metric of performance. The normalizations carried out in equations (9) and (10) were selected so that the highest possible values NMSE and NMAE are one. In addition, we assess the compression efficiency using the compression ratio (CR):

$$CR = (1 - \frac{2L+1}{N_f}) \cdot 100\% \tag{11}$$

where L and N_f denote the number of compressive samples and the number of Nyquist samples, respectively.

Table 2. FRI reconstruction of FW with 26dB SNR

	K	NMSE	NMAE	CR (%)
Resolution 1	10	0.35e-3	0.144	99.4
Resolution 2	20	0.027e-3	0.054	99.1
Resolution 3	35	0.027e-3	0.053	98.7
Resolution 4	52	0.0031e-3	0.0178	98.2

An example applying the FRI approach to sampling and reconstruction is shown in Figure 3. On this real FW obtained from the dataset described in subsection 2.2, we performed sampling and applied the method described in section 4.4 with Cadsow's denoising for reconstruction. Since the rate of innovation is generally unknown, we assume values of $K = 10, 20, 35$ and 52. Reconstructions using these values results in FW signals of varying resolution. Performance results for the FW signal with an SNR of 26dB is summarized in table 2. Note that both the NMSE and NMAE are very small even in coarse resolutions. However, the morphology of the mode is not well preserved. To illustrate this, figure 3b shows the zoomed version of the FW signal at the mode. In this figure, the real FW signal and its corresponding reconstructions are shown for resolutions obtained with values of $K = 10$ and 52. Note that the reconstruction corresponding to the coarsest resolution loses a few of the peaks in the mode. As computed in table 2, there is a deviation of as much as 14.4% from the true FW value. Even at higher resolutions, however, a very high compression ratio is achieved. The reconstruction corresponding to $K = 52$ is also shown in Figure 3. This reconstruction preserves the morphology of the true FW signal with two being perceptually indistinguishable.

Summarizing, we found that high quality reconstructions which preserve FW morphology can be achieved by considering FW's as signals of finite rates of innovation. Within this

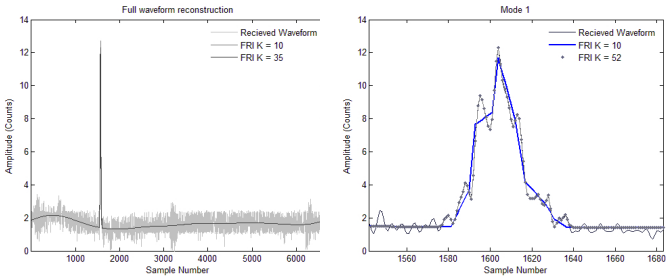

(a) Full-waveform (b) Mode 1

Fig. 3. Full-waveform FRI reconstruction.

framework, we have achieved effective sampling by acquiring samples at or above the rate of innovation ρ. In general, high compression ratios are achieved at the cost of small errors. In the FW signal example with SNR = 26 dB, a CR in the range of $[98 - 99\%]$ achieved with NMSE's as high as 0.35e-3. Thus, these results show that we can efficiently sample FW signals at the rate of innovation and achieve very precise reconstructions.

In addition, we have found also that we can achieve distinct resolutions by changing the sampling rate. A higher innovations number K implies that we increase the rate of innovation ρ accordingly. The higher the sampling rate, the larger the number of samples that we collect thus increasing the number of estimated knots used in the linear spline approximation. An increased number of knots improves the resolution of the reconstruction, thus refining the NMSE and Nmax at the cost of decreased compression ratio. Fortunately, we found in this research that very precise reconstructions can be achieved at low innovation numbers (e.g $K = 61$ overall) for the dataset described in subsection 2.2.

Finally, in [2] the authors used a moment-based FRI method to determine FW complexity characterizations as numeric estimates of the scene complexity. The intended use for these characterizations was to determine, on the fly, a number of FW scans that need to be randomly collected across the entire 3D scene to achieve acceptable reconstruction quality for the entire range map. In fact, such a number was theoretically determined in [8]. The approach, was applied successfully on simulated data by using the number of innovations K as a measure of complexity. Application of the method described here to estimate FW complexity on real data is straightforward. The more complex morphology of the modes and the higher number of modes implies that a larger K is required to achieve a pre-specified reconstruction quality threshold. Thus, as in [2], the number of innovations K can be used to obtain FW characterizations of the scene complexity

6. CONCLUSION

In this paper, we studied the problem of efficient sampling and compression of FW in 3rd generation LIDAR systems. The approach implemented to address this important issue is based on considering FW's as signals with finite rates of innovation. This is made possible by modeling the FW signal as the sum of a bandlimited signal and linear non-uniform spline. The first finding of this paper shows that we can effectively model FW signals in this manner. After appropriate modeling, the FRI procedure was applied to efficiently sample and reconstruct real FW signals under certain resolution constraints imposed by K. The performance metrics show that we can achieve high resolutions by sampling only about 2% of the samples one would typically acquire at the Nyquist rate. In short, we show that we can efficiently and effectively sample FW signals and still achieve precise reconstructions which preserve signal morphology at very low rates.

7. REFERENCES

[1] S. Laky, P. Zaletnyik, and C. Toth, "Land classification of wavelet-compressed full-waveform lidar data," *International Archives of Photogrammetry and Remote Sensing*, vol. XXXVIII, no. 3A, pp. 115–119, September 2010.

[2] J. Castorena, C.D. Creusere, and D. Voelz, "Using finite moment rate of innovation for lidar waveform complexity estimation," in *IEEE Conference Record of the Fourty Fourth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*. IEEE, 2010, pp. 608–612.

[3] C. Mallet, F. Lafarge, M. Roux, U. Soergel, F. Bretar, and C. Heipke, "A marked point process for modeling lidar waveforms," *IEEE Transactions on Image Processing*, vol. 19, no. 12, pp. 3204–3221, December 2010.

[4] M.A. Hofton, J.B. Minster, and J.B. Blair, "Decomposition of laser altimeter waveforms," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 38, no. 4, pp. 1989–1996, July 2000.

[5] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, June 2002.

[6] T. Blu, P.L. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 31–40, March 2008.

[7] Y. Hao, P. Marziliano, M. Vetterly, and T. Blu, "Compression of ecg as signal with finite rate of innovation," in *In Proc. 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2005, pp. 7564–7567.

[8] J. Castorena and C.D. Creusere, "Random impulsive scan for lidar sampling," in *IEEE International Conference on Image Processing*, 2012, Accepted paper.