ON DISTRIBUTED SAMPLING OF BANDLIMITED AND NON-BANDLIMITED SENSOR FIELDS

Animesh Kumar, Prakash Ishwar, and Kannan Ramchandran

Department of Electrical Engineering, University of California,Berkeley, CA-94720,
{animesh,ishwar,kannanr}@eecs.berkeley.edu

ABSTRACT

Distributed sampling and reconstruction of a physical field using an array of sensors is a problem of increasing interest in environmental monitoring applications of sensor networks. This work addresses the related sampling framework in the context of both bandlimited and non-bandlimited sensor fields. We show, using a dither-based scheme, that it is possible to reconstruct non-bandlimited fields with a reconstruction accuracy that depends on the available bitrate R and the spectral decay characteristics of the sensor field – we study exponentially decaying spectra as an illustration. For bandlimited fields $f(t)$, the maximum pointwise error D_f decays as $D_f \sim 2^{-a_1 R}$, i.e. exponentially with rate R. For the non-bandlimited case, we show that for fields $u(t)$ with exponentially decaying spectral tails, i.e., $|U(\omega)| \sim e^{-a|\omega|}$, the maximum pointwise error D_u decays as $D_u \sim e^{-a_2\sqrt{R}}(1 + o(R))$ with spatial bit rate R bits/meter. We also show that it is possible to trade off the number of sensors with their precision, while maintaining a similar reconstruction accuracy – a phenomenon that may be dubbed as a *"bit-conservation"* principle underlying the sampling framework.

1. INTRODUCTION

With the inception of sensor networks, high spatial resolution remote sensing of physical fields has become a topic of great interest [1, 2, 3]. In the context of environmental monitoring and other field-reconstruction applications, a key challenge is to instrument "distributed" sampling of the sensor field with a collection of low-precision, low-cost sensing devices having limited computing and communications capabilities. The intimate relation between this problem and that of classical non-uniform sampling has been established recently in [1, 4] for bandlimited sensor fields, for the settings of both deterministic and stochastic fields. In addition to summarizing our results for the bandlimited case, we study the important regime of non-bandlimited fields in this work. Our motivation stems from the observation

that unlike the classical sampling setup, anti-aliasing pre-filtering luxury does not exists in the context of distributed sensor field sampling, where the sensors have to directly sample the field of interest. Given this requirement, what can we say about the reconstruction performance as the sampling density gets larger? What classes of physical fields can be "effectively" sampled in a distributed way? By "effective," we mean that the reconstruction error goes to zero as the sampling density gets large and the sampling bit rate gets large (note that this scaling property is not present in the works of [2, 3]). What is the rate of decay of reconstruction error as a function of the number of bits sent from the sensor field to the collector? These questions will be addressed quantitatively in this paper.

For simplicity, we will confine ourselves to the 1-D model shown in Fig 1. We will use 1-bit A/D precision sensors to highlight the low-precision nature of the sensing devices. We assume that these sensors are densely distributed on a straight line, and there is a central data collector. We consider the field at a particular temporal "snapshots" and analyze the spatial reconstruction behavior for each snapshot.

The main contributions of this paper are:

1). We characterize the effect of oversampling and quantization on amplitude-limited "smooth" fields that are bandlimited or non-bandlimited with suitably decaying spectral profile.

2). We establish how this can be ported to the distributed sampling framework relevant to sensor networks in the context of low-precision sensors, while maintaining similar asymptotic reconstruction performance. We uncover a key "bit-conservation" principle that quantifies the fundamental trade off between sensor precision and oversampling density.

3). We uncover an important "information scaling law" that establishes how information "grows" as a function of number of sensors, and show that the per-node per-snapshot information requirement goes to zero at the rate of $O\left(\frac{\log N}{N}\right)$ for bandlimited fields and $O\left(\frac{\log^2 N}{N}\right)$ for non-bandlimited fields with exponentially decaying spectral tails. The precise utility of this scaling law depends on the underlying routing protocol and transport model used to move this in-
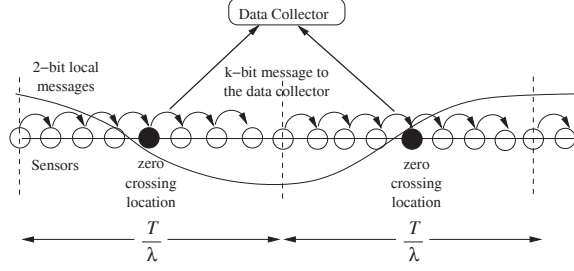
Fig. 1. *We assume that 1-bit sensors are uniformly placed along a straight line. The first zero crossing locations per temporal snapshot are located as follows: each sensor sends a 2-bit message to its right neighbor indicating two things: (i) whether or not a zero-crossing has already been found by some preceding sensor and (ii) if a zero-crossing has not yet been found, the sign of the field plus dither at its location. The sensor which observes a zero crossing for a particular snapshot (black sensors) communicates with the central data collector.*

formation from the sensors to the data-collector. Nonetheless, the fact that the per-node information needs to go zero is comforting and counters the "pessimistic" scaling law results for ad-hoc networks by Gupta and Kumar which (for a specific point-to-point transport model) establish that the per-node transport capabilities go to zero as the network scales[5]. Our results show that even in this regime, it is possible to do "useful" work in sensor networks, i.e. reconstruct a field with arbitrary reconstruction accuracy.

For the rest of the paper lower case letters will denote deterministic fields. Specifically, $f(t)$ will represent a deterministic BL field and $u(t)$ will represent a non-bandlimited (NBL) field. Both classes are assumed to be finite energy and amplitude limited to the level 1. The distortion criterion is the maximum pointwise error between the field and its reconstruction as the distortion criterion, $D_f := ||f - f_{rec}||_\infty$, where $f_{rec}(t)$ is the final reconstruction of the field (using any method).

2. APPLICATIONS TO SENSOR NETWORKS

Figure 1 illustrates the single bit dither-based sampling framework. Let $\lambda > 1$ be a *fixed* oversampling factor and $M = 2^k$ sensors be placed in every $\frac{T}{\lambda}$ length interval. The resulting spatial bitrate is $R = \frac{\lambda \log M}{T} = \frac{\lambda k}{T}$ and the sensor density is $N = \frac{\lambda M}{T}$. T is the Nyquist sampling interval for the case of BL fields. For the case of NBL fields T is chosen by balancing the aliasing and quantization errors. The following results also hold for b-bit sensors and 2^{k-b+1} sensors every $\frac{T}{\lambda}$ length interval.

- Let M, one-bit sensors be placed uniformly at every $\tau = \frac{T}{\lambda M}$ in every interval of the form $[\frac{lT}{\lambda}, \frac{(l+1)T}{\lambda} - \tau], l \in \mathcal{Z}$. The sensors at $\frac{lT}{\lambda}$ are the starting nodes.

- Periodically, the sensors take snapshots of the 1-D spatio-temporal field by comparing the field value to the dither value at their respective locations (the dither values are assumed to be pre-stored in sensors).

- Corresponding to each temporal snapshot of the field, each starting node passes a message to its (right) neighbor, indicating: (i) whether or not a zero-crossing has already been found by some preceding sensor and (ii) if a zero-crossing has not yet been found, the sign of the field plus dither at its location. This requires 2-bits of local messages (see fig 1).

- The first sensor in each $\frac{T}{\lambda}$ interval that detects a sign mismatch between what its left neighbor reports and its own reading records a one. Other sensors record a zero for that snapshot. The local communication for detecting the zero-crossings need not be done in real-time. The sensors can store the signs of the field plus dither over many snapshots and locate the zero-crossings later.

- Finally, each sensor encodes the zero-crossing data using principles of distributed compression [6, 7].

Growth of information: To understand how the sensor data-rate grows with M and T one needs to specify, at some level, a temporal model for the evolution of the random field. However, the temporal sequence of the first zero-crossing locations in every $\frac{T}{\lambda}$-interval contains all the information pertaining to the field. Suppose that the sensors within each $\frac{T}{\lambda}$ interval could collaborate without penalty to jointly encode the zero-crossing information over many temporal snapshots. By design, there is exactly one sensor per $\frac{T}{\lambda}$ interval per snapshot that will record the value one, i.e. only M possible vectors which require no more than $R_T = \log_2(M)$ bits per snapshot per $\frac{T}{\lambda}$ interval. It is possible to achieve a compression efficiency of $\log_2(M)$ bits every $\frac{T}{\lambda}$ interval per snapshot *without any sensor collaboration* using principles of distributed source coding theory [6, 7] which exploits the correlation between sensor outputs. In fact, it is possible for each sensor to encode its zero-crossing data at the rate $R_{sensor} = \frac{1}{M} \log_2(M)$ bits snapshot without talking to other sensors in the same $\frac{T}{\lambda}$ and

III - 926

the central unit will be able to recover the encoded information.

We now show that the information density is finite for a large class of smooth fields, i.e., a finite non-zero reconstruction error $D > 0$ can be achieved by a finite number of sensors in every $\frac{T}{\lambda}$-interval:-

For BL fields $M \sim \frac{1}{D}$, $N \sim \frac{1}{D}$ (since T is fixed by the Nyquist rate) and $R \sim \log \frac{1}{D}$. For fields with exponentially decaying spectra, $M = O\left(\frac{1}{D}\right), T = O\left(\frac{1}{\log 1/D}\right)$, $R = O\left(\log^2 D\right), N = O\left(\frac{\log 1/D}{D}\right)$. For fields with finite second moment in frequency domain, $R = O\left(\sqrt{\frac{1}{D}}\log\frac{1}{D}\right)$, $N = O\left(\frac{1}{D^{\frac{3}{2}}}\right)$ for high rates or small distortions.

Local inter-sensor communication cost: In the above local communication protocol, each 1-bit precision sensor transmits 2-bits to its right neighbor over a distance of τ. Hence, the local sensor communication cost is $\frac{2T}{M\lambda}$ bit-meters for each sensor. In general, for b-bit sensors ($b > 1$) the local sensor communication cost is no more than $(b+1)/(\lambda M_b 2^{b-1})$ where $M_b = \frac{M}{2^{b-1}}$ is the number of b-bit sensors needed to match the asymptotic reconstruction error performance of 1-bit sensors (in the high rate case). The total local communication cost density is no more than $\frac{2}{\lambda}$ bit-meters. Thus, with limited local communication cost, the sampling task can be distributed among the sensors.

Sensor distribution: The *"bit-conservation"* principle holds for BL fields at any rate and for NBL fields at high rates. For these cases, as we increase the A/D precision (b increases), we get a reduction in the number of sensors according to the "conservation of bits" principle while maintaining the same asymptotic error decay profile. Sensors need to be placed only in intervals of the form $\left[\frac{lT}{\lambda}, \frac{lT}{\lambda} + (M_b - 1)\tau\right)$, where T is fixed according to the spectral decay of the field and the target distortion. This leaves intervals over which there is no need to sample the field at all. Hence at high rates for a given reconstruction quality the number of sensing units goes down exponentially with b: 1-bit dithered sampling needs M sensors every $\frac{T}{\lambda}$ interval, b-bit ($b > 1$) dithered sampling needs $M_b = M/2^{b-1}$ sensors, and high precision sampling (see section 3) needs only one sensor every $\frac{T}{\lambda}$ interval. Since the scheme naturally allows "inactive" regions in oversampling, we can have bunched irregular sampling using sensors. This allows for design flexibility in sensor deployment. For example, in rugged terrain or in the presence of occluding obstacles, one would need to use higher precision sensors. Where sensors can be deployed in large numbers it is sufficient to use cheap 1-bit sensors.

Robustness: The dither-based oversampling method also offers robustness to node failures in terms of a graceful degradation of reconstruction error. For example, if every alternate node fails, the effective inter-node separation would increase to 2τ. The effect will vary for different spectral tails.

For BL fields this increase in minimum distance means a loss of 1-bit resolution. However, for NBL fields a decrease in effective τ will affect the balance between the aliasing error and the quantization error and hence the degradation will not be as simple as in the case of BL fields[8].

3. FIELD RECONSTRUCTION AND ACCURACY

A field $f(t)$ bandlimited to $\left[-\frac{\pi}{T}, \frac{\pi}{T}\right]$ can be reconstructed from the uniform samples $f\left(\frac{lT}{\lambda}\right), l \in \mathcal{Z}$ according to

$$f(t) = \sum_{l \in \mathcal{Z}} f\left(\frac{lT}{\lambda}\right) \phi\left(\frac{t}{T} - \frac{l}{\lambda}\right) \tag{1}$$

where $\phi(t)$ is a stable interpolating kernel satisfying $\sum_{l \in \mathcal{Z}} \left|\phi_\lambda\left(\frac{t}{T} - \frac{l}{\lambda}\right)\right| < C < \infty$ and C is independent of the bandwidth $W := \frac{\pi}{T}$[9, 1]. For samples taken at non-uniform sampling locations collected at the locations $\{t_l\}_{l \in \mathcal{Z}}$ satisfying $\sup_{l \in \mathcal{Z}} \left|t_l - \frac{lT}{\lambda}\right| < \kappa < \infty$ and $\inf_{i \neq j} |t_i - t_j| > \delta > 0$, there exist bandlimited interpolating functions $\psi_l(t)$ such that

$$f(t) = \sum_{l \in \mathcal{Z}} f(t_l) \psi_l\left(\frac{t - t_l}{T}\right) \tag{2}$$

where $C' := \sup_{t \in \mathbb{R}} \left[\sum_{l \in \mathcal{Z}} |\psi_l\left(\frac{t-t_l}{T}\right)|\right] < \infty$[9], with C' independent of T as for the uniform case.

3.1. Sampling the field with high precision sensors

Assume that sensors are equipped with k-bit precision uniform scalar quantizer $Q_k(.)$. The field is sampled at the locations $\{\frac{lT}{\lambda}\}_{l \in \mathcal{Z}}$. For BL fields, T is the Nyquist sampling interval and for NBL fields T is chosen by balancing the aliasing and quantization errors. We refer to this sampling framework as *"Nyquist-like"* sampling in what follows. The stable reconstruction of the field obtained from the quantized samples is $\widehat{f}(t) = \sum_{l \in \mathcal{Z}} Q_k\left(f\left(\frac{lT}{\lambda}\right)\right) \phi\left(\frac{t}{T} - \frac{l}{\lambda}\right)$ and $\widehat{u}(t) = \sum_{l \in \mathcal{Z}} Q_k\left(u\left(\frac{lT}{\lambda}\right)\right) \phi\left(\frac{t}{T} - \frac{l}{\lambda}\right)$, where $Q_k(.)$ denotes the uniform scalar quantizer[10]. For BL fields: $D_f = ||f - \widehat{f}||_\infty < C2^{-\frac{TR}{\lambda}}$, i.e. maximum pointwise error decays exponentially in k, the bit budget for every Nyquist interval[9]. For NBL fields, the reconstruction error depends on the spectral decay of the field. For example, for $u(t)$ with $|U(\omega)| \sim e^{a|\omega|}, |\omega| > W_0$, for $T < \frac{\pi}{W_0}$, we get $D_u = ||u - \widehat{u}||_\infty \sim \exp\left[-\sqrt{\pi a \lambda \log 2}\sqrt{R}\right]$[8]. The results will be different for different spectral decay characteristics of the field. However, for fields satisfying a mild condition $\int_{\mathbb{R}} |\omega|^2 |U(\omega)| d\omega < \infty$ it can be shown that for a finite non-zero D_u, only a finite k is needed.

3.2. Sampling the field with 1-bit precision sensors

Let $M = 2^k$ sensors of 1-bit precision be placed uniformly in every interval $\left[\frac{lT}{\lambda}, \frac{(l+1)T}{\lambda}\right)_{l \in \mathcal{Z}}$ with an inter-sensor spacing $\tau = \frac{T}{M\lambda}$. A smooth dither field $d_T(t)$ with $|d_T(lT/\lambda)| =$

$\gamma > 1$ and $\text{sgn}\left\{d_T\left(\frac{lT}{\lambda}\right)\right\} = -\text{sgn}\left\{d_T\left(\frac{(l+1)T}{\lambda}\right)\right\}$ en-sures that the sum field $f(t) + d_T(t)$ (resp. $u(t) + d_T(t)$) crosses the level zero in every interval $\left[\frac{lT}{\lambda}, \frac{(l+1)T}{\lambda}\right)$. A dither field with uniformly bounded slope $||d'_T||_\infty = \frac{\Delta}{T} < \infty$ translates a bounded inaccuracy in space to a bounded inaccuracy in amplitude through the mean value theorem. For example, $d_T(t) = \gamma \cos\left(\frac{\lambda \pi t}{T}\right)$ is a valid dither func-tion. If m_l is the first index such that $[f + d_T]\left(\frac{lT}{\lambda} + m_l \tau\right)$ and $[f + d_T]\left(\frac{lT}{\lambda} + (m_l + 1)\tau\right)$ have different signs, then $|f(t_l) + d_T(t_l)| < ||f + d_T||_\infty \frac{\tau}{2}$, where, $t_l := \frac{lT}{\lambda} + \left(m_l + \frac{1}{2}\right)\tau$. If $\widehat{f}(t) = \sum_{l \in \mathcal{Z}} (-d(t_l)) \psi_l\left(\frac{t - t_l}{T}\right)$ and $\widehat{u}(t) = \sum_{l \in \mathcal{Z}} (-d(t_l))$ $\psi_l\left(\frac{t - t_l}{T}\right)$, then $D_f := ||f - \widehat{f}||_\infty < C'' 2^{-\frac{TR}{\lambda}}$ where C'' is independent of T[9]. For NBL fields the result depends on the spectral decay. For example, fields with exponential decay, $D_u := ||u - \widehat{u}||_\infty \sim \exp\left[-\sqrt{\pi a \lambda \log 2}\sqrt{R}\right](1 + o(R))$[8]. For high rates R, the distortion behaves similar to the "Nyquist-like" sampling.

3.3. The trade-off between the number of sensors and their precision

Let the bit budget per unit length R and k be fixed. As opposed to 1-bit sensors that can detect only a single level crossing, b-bit sensors for $1 < b < k$ can detect a level change among $2^b - 1$ levels. This idea can be used to de-sign a dither based scheme with b-bit precision sensors with sensor density strictly less than 2^k in every $\frac{T}{\lambda}$ interval. The idea is to use a dither function with bounded slope which forces a *level crossing* in an interval of the form $[A_l, B_l] := \left[\frac{lT}{\lambda}, \frac{lT}{\lambda} + (M_b - 1)\tau\right)_{l \in \mathcal{Z}}$, where $M_b := \frac{M}{2^{b-1}}$. Sensors are placed at every $\tau := \frac{T}{\lambda 2^k}$ and thus the position of the first level crossing can be located to an accuracy of $\frac{\tau}{2}$. Let m_l be the first index of level crossing and let q_l be the level crossed in $[A_l, B_l]_{l \in \mathcal{Z}}$. Let $t_l = \frac{lT}{\lambda} + \left(m_l + \frac{1}{2}\right)\tau$, then $\widehat{f}(t) = \sum_{l \in \mathcal{Z}} (q_l - d_b(t_l)) \psi_l\left(\frac{t - t_l}{T}\right)$ is the approximate re-construction. For BL fields $D_f := ||f - \widehat{f}||_\infty < C''' 2^{-\frac{TR}{\lambda}}$, where C''' is a constant independent of T[1]. For the NBL case, the trade-off comes into play at high rates when the slope of the field can be bounded in terms of the decay-ing spectral characteristics. For example, at high-rates for the field $u(t)$ with exponentially decaying spectrum $D_u = ||u - \widehat{u}||_\infty \sim \exp\left[-\sqrt{\pi a \lambda \log 2}\sqrt{R}\right]$. These results can be generalized to wide sense stationary stochastic fields with $\mathcal{L}^p$ distortion measures. We refer the interested reader [4] for the details. We state the bit-conservation principle for BL fields in general and for NBL fields at high rate as

Bit-conservation principle: Let k be the number of bits available per $\frac{T}{\lambda}$-interval. For each $1 \leq b < k$ there exists a sampling scheme with 2^{k-b+1} b-bit A/D converters per $\frac{T}{\lambda}$-interval achieving a worst-case pointwise reconstruction accuracy similar to the "Nyquist-like" sampling.

4. CONCLUSIONS AND FUTURE WORK

We addressed the problem of deterministic oversampling and reconstruction of bandlimited or (smooth) non-bandlimited fields and showed that the maximum pointwise error de-creases to zero as the bitrate goes to infinity. This work summarizes the scope of the distributed sampling frame-work in the case of smooth fields. We showed how the dis-tributed sampling framework can be applied to sensor net-works, with distortion going to zero as the number of sen-sors in the network scales to infinity. We also showed the ex-istence of a "bit-conservation" principle. Ongoing work in-cludes stochastic extensions to non-bandlimited, wide-sense-stationary processes and in general to fields having a finite rate of innovation [11].

5. REFERENCES

[1] P. Ishwar, A. Kumar, and K. Ramchandran, "Distributed Sampling for Dense Sensor Networks: a "bit-conservation" principle," in *IPSN, Proc. of the 2^{nd} Intl. Wkshp., Palo Alto, CA, USA*, LNCS edited by L. J. Guibas and F. Zhao, Springer, NY, 2003, pp. 17–31.

[2] D. Marco, E. J. Duarte-Melo, M. Liu, and D. L. Neuhoff, "On the Many-to-One Transport Capacity of a Dense Wire-less Sensor Network and the Compressibility of its Data," in *IPSN, Proc. of the 2^{nd} Intl. Wkshp., Palo Alto, CA, USA*, LNCS edited by L. J. Guibas and F. Zhao, Springer, NY, 2003, pp. 1–16.

[3] A. Scaglione and S. D. Servetto, "On the interdependence of routing and data compression in multi-hop sensor networks," in *Proc. of the 8^{th} Ann. Intl. Conf. on Mobile computing and networking*, pp. 140–147, ACM Press, 2002.

[4] P. Ishwar, A. Kumar, and K. Ramchandran, "On Dis-tributed Sampling in Dense Sensor Networks: a "bit-conservation" Principle," *Preprint available at www.eecs.berkeley.edu/~animesh/jsac03ishwartetal.pdf*.

[5] P. Gupta and P. Kumar, "Capacity of wireless networks," *IEEE Trans on Info. Thy.*, vol. 46, pp. 388–404, Mar 2000.

[6] D. Slepian, and J. K. Wolf, "Noiseless Coding of Correlated Information Sources," *IEEE Trans on Info. Thy.*, vol. 19, pp. 471–480, july 1973.

[7] S. S. Pradhan and K. Ramchandran, "Distributed Source Coding Using Syndromes (DISCUS): Design and Construc-tion," *IEEE Trans on Info. Thy.*, vol. 49, pp. 626–643, Mar 2003.

[8] A. Kumar, P. Ishwar, and K. Ramchandran, "On Distributed Sampling of Smooth Non-Bandlimited Fields," *submitted to IPSN 2004*.

[9] Z. Cvetković and I. Daubechies, "Single Bit oversampled A/D conversion with exponential accuracy in bit rate," *Pro-ceeding DCC*, pp. 343–352, March 2000.

[10] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Boston: Kluwer Academic, 1992.

[11] M. Vetterli, P. Marzilliano, and T. Blu, "Sampling Signals with Finite Rate of Innovation," *IEEE Trans. Signal Proc.*, pp. 1417–1428, June 2002.