

SPATIAL SUPER-RESOLUTION OF A DIFFUSION FIELD BY TEMPORAL OVERSAMPLING IN SENSOR NETWORKS

Yue M. Lu and Martin Vetterli

School of Computer and Communication Sciences
Swiss Federal Institute of Technology Lausanne (EPFL), CH-1015, Switzerland
Email: {yue.lu, martin.vetterli}@epfl.ch; Web: lcav.epfl.ch/~{lu, vetterli}

ABSTRACT

We study the spatial-temporal sampling of a linear diffusion field, and show that it is possible to compensate for insufficient spatial sampling densities by oversampling in time. Our work is motivated by the following issue often encountered in sensor network sampling, namely increasing the temporal sampling density is often easier and less expensive than increasing the spatial sampling density of the network. For the case of sampling a diffusion field, we show that, to achieve trade-off between spatial and temporal sampling, the spatial arrangement of the sensors must satisfy certain conditions. We provide in this paper the precise relationships between the achievable reduction of spatial sampling density, the required temporal oversampling rate, the spatial arrangement of the sensors, and the bound for the condition numbers of the resulting sampling and reconstruction procedures.

Index Terms— Sampling, sensor network, diffusion process, spatial-temporal sampling, super-resolution.

1. INTRODUCTION

A sensor network can be seen as a sampling device, taking spatial-temporal samples of some underlying physical field of interest. If our goal is to reconstruct the original field — or at least a good approximation of it — then the important signal processing questions to answer are where to take the samples (i.e. spatial sensor placement), and when to take the samples (i.e. temporal sampling instants).

While regular multidimensional sampling theory [1] is a well developed field, it usually assumes homogeneity over the dimensions, namely the dimensions are interchangeable (as in images or volumetric data). However, in the case of physical field sampling by sensor networks, the dimensions — space and time — are specific and cannot be interchanged. For example, increasing the spatial sampling rate is often much more expensive than increasing the temporal sampling rate, since the former requires the physical presence of more sensors in the network, whereas the latter is, in theory, only constrained by the communication capacity and energy budget of the network.

M. Vetterli is also with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA. This work was supported by the Swiss National Science Foundation under Grants 200020-103729 and 5005-67322 (NCCR-MICS).

In this work, we study the trade-off between the spatial and temporal sampling densities for the specific problem of reconstructing a field governed by the diffusion equation. Such models are widely used to characterize the spatial-temporal evolution of various physical phenomena (e.g. temperature variation, and the distribution of pollution plumes in the atmosphere).

By exploiting the spatial-temporal correlation offered by the diffusion equation, we show that it is possible to reduce the spatial sampling requirement by taking more samples along time. In other words, we can use a sequence of spatially undersampled and aliased measurements to reconstruct a field with higher spatial bandwidth. Owing to its apparent similarity to the classical super-resolution problem in image processing, we refer to the proposed scheme as the spatial super-resolution reconstruction of a diffusion field.

The rest of the paper is organized as follows. In Section 2, we first briefly overview some relevant concepts on the diffusion equation model, and then precisely state the spatial super-resolution problem we want to address in this paper. We study in Section 3 a general class of periodic nonuniform spatial arrangements of the sensors, and provide precise relationships between the achievable spatial super-resolution factor, the required temporal oversampling rate, spatial sensor arrangement, and the bound for the condition number of the resulting sampling problem. We conclude the paper in Section 4. Due to space limitations, we only present the main results in this paper, and leave the proofs to [2].

2. PROBLEM STATEMENT

2.1. The Linear Diffusion Equation

Consider a field $f(x, t)$ with one¹ spatial variable x and one temporal variable t , satisfying the following linear diffusion (heat) equation

$$\frac{\partial f(x, t)}{\partial t} = D \frac{\partial^2 f(x, t)}{\partial x^2}, \quad \text{for } x \in \mathbb{R}, t > 0, \quad (1)$$

where D is the diffusion coefficient. In the following discussions, we set $D = 1$ by rescaling the time axis. We assume that the sensor network starts observing the field at $t = 0$, by

¹The results presented in this paper can be extended to cases with multiple spatial variables.

which time all sources that originally induced the field have disappeared. Consequently, there is no external source term in the homogeneous diffusion equation (1).

For any $t > 0$, the field $f(x, t)$ satisfying (1) is completely determined by the initial condition $f(x) \stackrel{\text{def}}{=} f(x, t = 0)$. Denoting by $\hat{f}(\omega, t)$ the spatial Fourier transform of $f(x, t)$, we have [3]

$$\hat{f}(\omega, t) = \hat{f}(\omega) e^{-\omega^2 t}. \quad (2)$$

The above formula precisely characterizes the spatial-temporal correlation of the field: for any $t > 0$, $f(x, t)$ is just a filtered version of the initial state $f(x)$, with the filter being a time-varying Gaussian kernel.

2.2. Problem Statement

Suppose that the initial state $f(x)$ can be well-approximated by a bandlimited function. Without loss of generality, we can assume that the (essential) frequency domain support of $f(x)$ is $[-\pi, \pi]$. Denote by $\text{BL}(\pi)$ the space of functions whose Fourier transforms vanish outside of $(-\pi, \pi)$.

We consider the following sampling setup, where we place a sequence of sensors at spatial locations

$$\mathcal{X} \stackrel{\text{def}}{=} \{x_n : n \in \mathbb{Z}\}.$$

Each sensor takes $K \geq 1$ measurements at instants

$$\left\{ \frac{k-1}{K} \tau : 1 \leq k \leq K \right\},$$

where τ/K is the uniform sampling interval along time. Note that $K = 1$ corresponds to the case of pure spatial sampling.

We define the corresponding sampling operator $A_{\mathcal{X}, K}$ to be the linear mapping from any function $f(x) \in \text{BL}(\pi)$ to a sequence of spatial-temporal measurements, i.e.,

$$(A_{\mathcal{X}, K} f)_{n, k} \stackrel{\text{def}}{=} f(x_n, \frac{k-1}{K} \tau), \quad n \in \mathbb{Z}, 1 \leq k \leq K.$$

Our goal is to study the conditions on \mathcal{X} and K such that any function $f(x) \in \text{BL}(\pi)$ can be perfectly reconstructed from its samples $A_{\mathcal{X}, K} f$. Clearly, in that case, a necessary condition is that the sampling operator $A_{\mathcal{X}, K}$ must be an invertible or one-to-one mapping. However, in practice, stronger requirements are needed: we want to be able to reconstruct $f(x)$ in a *numerically stable* way from $A_{\mathcal{X}, K} f$. To guarantee such an algorithm exists, we need to ensure that if $A_{\mathcal{X}, K} f_1$ is “close” to $A_{\mathcal{X}, K} f_2$ then f_1 is “close” to f_2 as well. Furthermore, we want that a small change in the signal f only produces a small change in its sampling data $A_{\mathcal{X}, K} f$. These requirements motivate the following condition on the sampling operator [4].

Definition 1 (Stable sampling) *We call $A_{\mathcal{X}, K}$ a stable sampling operator if there exist constants $0 < \alpha \leq \beta < \infty$ such that for every $f \in \text{BL}(\pi)$,*

$$\alpha \|f\|_{L^2}^2 \leq \|A_{\mathcal{X}, K} f\|_{\ell^2}^2 \leq \beta \|f\|_{L^2}^2. \quad (3)$$

We call α and β stability bounds and the tightest ratio $\kappa = \beta/\alpha$ provides a measure of the stability of the sampling operator.

We can see that stable sampling implies that $A_{\mathcal{X}, K}$ is invertible, whereas the reverse is not necessarily true.

Let $s(\mathcal{X})$ denote the ratio between the Nyquist density of the field $f(x)$ and the “average” spatial density of a particular sensor arrangement \mathcal{X} . For example, in the case of uniform sensor arrangement, i.e., $\mathcal{X} = \{nd : n \in \mathbb{Z}\}$ for some $d > 0$, we have $s(\mathcal{X}) = (\pi/\pi)/(1/d) = d$. Note that for an arbitrary \mathcal{X} , the proper definition of its “average” density can be rather technical [4]. We defer the precise definition of $s(\mathcal{X})$ to Section 3, where we focus on a particular class of nonuniform arrangements, whose densities can be easily calculated.

In this work, we are only interested in cases when $s(\mathcal{X}) > 1$, since otherwise the sensor network would have enough spatial density (as compared to the Nyquist density) and the field reconstruction problem can be directly solved by a spatial-only sampling scheme (i.e. $K = 1$). Intuitively, when $s(\mathcal{X}) > 1$, this quantity indicates how spatially “undersampled” our sensor network is, and — in the case of still being able to perfectly reconstruct $f(x)$ — the spatial “super-resolution factor” we can achieve.

The focus of this paper is to study and answer the following questions.

1. Is it possible to find some \mathcal{X} and K such that $s(\mathcal{X}) > 1$ and yet $A_{\mathcal{X}, K}$ is still a stable sampling operator?
2. If so, what is the largest spatial super-resolution factor $s(\mathcal{X})$ we can achieve?
3. What kind of sensor arrangement \mathcal{X} will allow us to achieve spatial super-resolution?

2.3. The Limitation of Uniform Sensor Arrangement

To see why we need to consider the last question listed above, we study here the case of uniform sensor arrangement, and show a somewhat surprising result.

Proposition 1 *Suppose that we put the sensors uniformly along space, i.e., $\mathcal{X} = \{nd : n \in \mathbb{Z}\}$ for some $d > 0$. If*

$$s(\mathcal{X}) = d > 1,$$

then the sampling operator $A_{\mathcal{X}, K}$ is unstable for arbitrary $K \geq 1$.

In other words, for uniform sensor arrangement, it is *impossible* to achieve spatial super-resolution ($s(\mathcal{X}) > 1$), no matter how many temporal samples we take.

The result of Proposition 1 can be *intuitively* understood by considering the initial state of the field to be $f(x) = \sin(\pi x/d)$. When $d > 1$, its Fourier transform

$$\hat{f}(\omega) = \pi j \left(\delta(\omega + \frac{\pi}{d}) - \delta(\omega - \frac{\pi}{d}) \right)$$

is supported within $(-\pi, \pi)$. It follows from (2) that

$$\begin{aligned} \hat{f}(\omega, t) &= \pi j \left(\delta(\omega + \frac{\pi}{d}) - \delta(\omega - \frac{\pi}{d}) \right) e^{-\omega^2 t} \\ &= \pi j \left(\delta(\omega + \frac{\pi}{d}) - \delta(\omega - \frac{\pi}{d}) \right) e^{-(\pi/d)^2 t} \\ &= \hat{f}(\omega) e^{-(\pi/d)^2 t}, \end{aligned}$$

and thus $f(x, t) = e^{-(\pi/d)^2 t} \sin(\pi x/d)$. Consequently, one can verify that $f(nd, t) = 0$ for all $n \in \mathbb{Z}$ and $t \geq 0$. In other words, there exists a nonzero function $f(x)$ from $\text{BL}(\pi)$ such that $A_{\mathcal{X}, K} f = \mathbf{0}$, and hence the sampling operator is not invertible. A catch in the above discussion though is that $f(x) = \sin(\pi x/d)$ does not have finite energy as required in (3). We leave the formal proof of Proposition 1 to [2].

3. PERIODIC NONUNIFORM SENSOR ARRANGEMENT

3.1. General Conditions on Stable Sampling

The previous example shows that a uniform sensor arrangement cannot “see” the signal $f(x) = \sin(\pi x/d)$, since the sensor locations match exactly with the zero crossings of $f(x)$. Intuitively, one might be able to avoid such situations by introducing certain “nonuniformity” to the set \mathcal{X} .

In this paper, we focus on a particular class of nonuniform sensor arrangement, in which \mathcal{X} consists of a union of $L \geq 2$ different uniform patterns, each shifted by a distance ξ_ℓ . More specifically,

$$\mathcal{X} = \{\xi_\ell + nd : n \in \mathbb{Z}, 1 \leq \ell \leq L\} \quad (4)$$

for some $d > 0$ and $0 \leq \xi_\ell < d$. In this case, the average density of \mathcal{X} is L/d , and hence $s(\mathcal{X})$ — the ratio between the Nyquist density of the field and the average sensor density — can be calculated as

$$s(\mathcal{X}) = \frac{d}{L}.$$

Let $\{y_{k,\ell}[n] = f(\xi_\ell + nd, (k-1)\tau/K)\}_{n \in \mathbb{Z}}$ denote the set of samples taken by the ℓ th shifted lattice, at time instant $(k-1)\tau/K$. Applying the classical sampling formula in the Fourier domain (as obtained from the Poisson summation formula), we can write

$$\begin{aligned} \hat{y}_{k,\ell}(\omega) &\stackrel{\text{def}}{=} \sum_{n \in \mathbb{Z}} y_{k,\ell}[n] e^{-j\omega(nd+\xi_\ell)} \\ &= \frac{1}{d} \sum_{m \in \mathbb{Z}} \hat{f}(\omega + mc) e^{-(\omega+mc)^2 (k-1)\tau/K} e^{j\xi_\ell mc}, \end{aligned} \quad (5)$$

where $c \stackrel{\text{def}}{=} 2\pi/d$.

We make two observations about the equality in (5). First, $\hat{y}_{k,\ell}(\omega)$ is a periodic function of ω , and thus we only need to consider its values within one period, e.g., $[0, c)$. Second, since $\hat{f}(\omega)$ is bandlimited to $(-\pi, \pi)$, then for any given ω , the summation on the right hand side of (5) only involves a finite number of nonzero terms. In fact, denoting by $M(\omega)$ the set of all indices m such that $|\omega + m| < \pi$, we can show that

$$\max_{\omega} |M(\omega)| = \lceil d \rceil,$$

where $\lceil d \rceil$ is the smallest integer greater than or equal to d .

Let $m_i(\omega)$ ($i = 1 \dots \lceil M(\omega) \rceil$) be the i th element in $M(\omega)$. We first introduce the following two matrices: $\mathbf{W}(\omega)$ is an $L \times \lceil M(\omega) \rceil$ matrix with entries

$$\{\mathbf{W}(\omega)\}_{\ell,i} \stackrel{\text{def}}{=} e^{j\xi_\ell m_i(\omega)c},$$

and $\mathbf{D}(\omega)$ is a diagonal matrix whose i th diagonal element is equal to

$$\{\mathbf{D}(\omega)\}_{i,i} \stackrel{\text{def}}{=} e^{-(\omega+m_i(\omega))^2 \tau/K}, \quad \text{for } i = 1 \dots \lceil M(\omega) \rceil.$$

The relation in (5) can now be written as a compact matrix-vector multiplication in the Fourier domain

$$\hat{\mathbf{y}}(\omega) = \mathbf{G}_{\mathcal{X}, K}(\omega) \hat{\mathbf{f}}(\omega),$$

where $\hat{\mathbf{y}}(\omega)$ is an $LK \times 1$ vector formed by $\hat{y}_{k,\ell}(\omega)$ in lexicographic order (i.e. $\hat{\mathbf{y}}_{(k-1)L+\ell}(\omega) = \hat{y}_{k,\ell}(\omega)$), $\hat{\mathbf{f}}(\omega)$ is an $\lceil M(\omega) \rceil \times 1$ vector with $\hat{\mathbf{f}}_i(\omega) = \hat{f}(\omega + m_i(\omega))$, and $\mathbf{G}_{\mathcal{X}, K}(\omega)$ is a block matrix formed by $\mathbf{W}(\omega)$ and $\mathbf{D}(\omega)$ as follows

$$\begin{pmatrix} \mathbf{W}(\omega)\mathbf{D}^0(\omega) \\ \mathbf{W}(\omega)\mathbf{D}^1(\omega) \\ \vdots \\ \mathbf{W}(\omega)\mathbf{D}^{K-1}(\omega) \end{pmatrix}. \quad (6)$$

The matrix $\mathbf{G}_{\mathcal{X}, K}(\omega)$ defined above contains all the information about the sampling operator $A_{\mathcal{X}, K}$. As shown in the following proposition, we can check the stability of the sampling operator by studying the properties of $\mathbf{G}_{\mathcal{X}, K}(\omega)$.

Proposition 2 *The sampling operator $A_{\mathcal{X}, K}$ is stable if and only if*

$$\begin{aligned} 0 < \alpha^* &\stackrel{\text{def}}{=} \text{ess inf}_{\omega} \sigma_{\min}(\mathbf{G}_{\mathcal{X}, K}^*(\omega) \mathbf{G}_{\mathcal{X}, K}(\omega)) / d \\ &\leq \beta^* \stackrel{\text{def}}{=} \text{ess sup}_{\omega} \sigma_{\max}(\mathbf{G}_{\mathcal{X}, K}^*(\omega) \mathbf{G}_{\mathcal{X}, K}(\omega)) / d < \infty, \end{aligned} \quad (7)$$

where $\sigma_{\min}(\cdot)$ and $\sigma_{\max}(\cdot)$ represent the smallest and largest eigenvalue of a matrix, respectively.

An immediate consequence of Proposition 2 is that, for $A_{\mathcal{X}, K}$ to be stable, the matrix $\mathbf{G}_{\mathcal{X}, K}(\omega)$ must have full column rank for almost all ω . This observation leads to the following bound between the number of temporal samples K and the achievable spatial super-resolution factor $s(\mathcal{X})$.

Corollary 1 *If $A_{\mathcal{X}, K}$ is a stable sampling operator, then*

$$K \geq \left\lceil \frac{\max_{\omega} |M(\omega)|}{L} \right\rceil = \lceil s(\mathcal{X}) \rceil.$$

The above result is intuitive: if our sensor network is roughly P -times undersampled in space, then for stable reconstruction, we must perform at least P -times oversampling in time as compensation.

3.2. Theoretical Achievability of Spatial Super-Resolution

Next, we show that spatial super-resolution can indeed be achieved by a wide class of spatial sensor arrangements.

Theorem 1 *For arbitrary choices of d and L with $s(\mathcal{X}) = \frac{d}{L} > 1$, the periodic nonuniform pattern \mathcal{X} as in (4) can lead to a stable sampling operator, if the following two conditions hold.*

1. $K \geq \lceil d \rceil$;
2. For any given integer N with $1 \leq N < d$, there exist ξ_ℓ and ξ_n such that $|\xi_\ell - \xi_n| \neq md/N$ for all $m \in \mathbb{Z}$.

Note that the first condition is easy to satisfy — we just need to take enough temporal measurements. The second condition in Theorem 1 (a special case of this condition was first proposed in [5] for bounded spatial domains) poses certain constraints on the geometry of the spatial pattern \mathcal{X} . Our next result shows that these constraints are in fact almost always satisfied.

Proposition 3 Consider arbitrary choices of the shift vector $\xi = (\xi_1, \xi_2, \dots, \xi_L)^T \in [0, d)^L$. The set of those vectors that do not satisfy the second condition in Theorem 1 has Lebesgue measure zero.

Combining the results of Theorem 1 and Proposition 3, we can reach the following conclusion: it is possible to achieve an arbitrary spatial super-resolution factor with stable reconstruction, as long as we take enough temporal samples such that $(s(\mathcal{X}), K)$ falls inside the dark-gray region of Figure 1. The suitable spatial arrangement of the sensors can be obtained, with probability one, by randomly drawing $(\xi_1, \xi_2, \dots, \xi_L)^T$ according to any continuous probability distribution defined on $[0, d)^L$.

3.3. Practical Achievability of Spatial Super-Resolution

While theoretically very promising, the above message has to be taken with a grain of salt. Note that for stable sampling, Theorem 1 only requires that the condition number $\kappa = \beta/\alpha$ to be finite. In practice, however, the actual value of the condition number makes a big difference in terms of noise amplification and numerical stability, which consequently puts a limit on the achievable spatial super-resolution factors.

Motivated by the above argument, we study the following more practically relevant question. Given a maximum allowable condition number, what are the corresponding achievable spatial super-resolution factors? While it is generally difficult to explicitly compute the exact stability bounds as defined in (7), we derive the following useful estimate.

Theorem 2 The condition number of the sampling operator $A_{\mathcal{X}, K}$ is lower bounded by the following quantity:

$$\kappa = \frac{\beta}{\alpha} \geq \max_{N=1, 2, \dots, (\lceil d \rceil - 1)} \frac{1 + |\sum_{\ell=1}^L e^{j2\pi\xi_\ell N/d}|/L}{1 - |\sum_{\ell=1}^L e^{j2\pi\xi_\ell N/d}|/L}. \quad (8)$$

We can show that, for fixed L , the above lower bound is an increasing function with respect to the spatial super-resolution factor $s(\mathcal{X})$. Consequently, given a certain “budget” on the condition number, we can use (8) to find an upper bound for the achievable super-resolution factors that stay within the “budget” [2].




Fig. 1. Any choice of the spatial super-resolution factor $s(\mathcal{X})$ and temporal sample number K in the dark-gray region is achievable with stable sampling. By contrast, any combination in the light gray region is unachievable, as a result of Corollary 1.

3.4. Reconstruction Algorithms

Up to now, we have only discussed the stability of the sampling operator, without mentioning the actual reconstruction algorithms. Using the techniques of multichannel sampling and reconstruction developed in [6], we can show that, in the case of stable sampling, the original continuous field can be perfectly reconstructed from its samples via a sequence of filtering operations. We leave the details to [2].

4. CONCLUSION

We studied the trade-off between the spatial and temporal sampling densities of a sensor network for the reconstruction of a linear homogeneous diffusion field. We show that it is possible to achieve an arbitrary spatial super-resolution factor with stable reconstruction, as long as we take enough temporal samples and that the spatial arrangement of the sensors satisfies certain conditions. We also provide a useful bound for determining the practically achievable super-resolution factors for a given maximum allowable condition number. As an important area of further research, we are investigating ways to extend the current work to the more general inhomogeneous case, i.e., the sampling and reconstruction of a diffusion field driven by an unknown spatially and temporally varying source.

5. REFERENCES

- [1] D. P. Peterson and D. Middleton, “Sampling and reconstruction of wavenumber-limited functions in N -dimensional Euclidean spaces,” *Inform. Contr.*, vol. 5, pp. 279–323, 1962.
- [2] Y. M. Lu and M. Vetterli, “Super-Resolution reconstruction of a diffusion field by spatial-temporal sampling in sensor networks,” in preparation, to be submitted to *IEEE Trans. Signal Process.*, 2009.
- [3] L. C. Evans, *Partial Differential Equations*, American Mathematical Society, Providence, RI, 1998.
- [4] H. J. Landau, “Sampling, data transmission, and the Nyquist rate,” *Proc. IEEE*, vol. 55, no. 10, pp. 1701–1706, Oct. 1967.
- [5] Y. Sakawa, “Observability and related problems for partial differential equations of parabolic type,” *SIAM J. Control*, vol. 13, no. 1, pp. 14–27, Jan. 1975.
- [6] A. Papoulis, “Generalized sampling expansion,” *IEEE Trans. Circuits Syst.*, vol. CAS-24, no. 11, pp. 652–654, Nov. 1977.