

# DISTRIBUTED SIGNAL PROCESSING AND COMMUNICATIONS: ON THE INTERACTION OF SOURCES AND CHANNELS<sup>1</sup>

Thibaut Ajdler\*, Razvan Cristescu\*, Pier Luigi Dragotti<sup>§</sup>, Michael Gastpar<sup>◊</sup>, Irena Maravic\*, Martin Vetterli\*<sup>◊</sup>

\*Laboratory for Audio-Visual Communications, Swiss Federal Institute of Technology Lausanne, 1015 Lausanne, Switzerland

<sup>§</sup>Department of Electrical and Electronic Engineering, Imperial College, London SW7 2BT, UK

<sup>◊</sup>Department of EECS, University of California at Berkeley, Berkeley CA 94720, USA

## ABSTRACT

Distributed ways of communicating, processing, and sensing are replacing more traditional centralized architectures. An early example of this revolution in distributed communications is appearing in the form of sensor networks, which are densely distributed networks of embedded signal sensors, controls and processors. These nodes could be simple signal sensors, but could also be cameras and microphones. In this distributed scenario, there are several interesting topics to investigate that span from traditional signal processing problems (i.e. sampling, compression, detection) to communication and information theory (i.e. transmission protocols, capacity bounds for ad-hoc networks). This paper reviews some recent results on the topic of source representations and distributed source coding and transmission.

## 1. INTRODUCTION

With the advent of sensor networks and other distributed sources and channels, a new paradigm for signal processing and communications is emerging. This has a profound impact on the way we think about signals, the way we process and compress signals, and the way we transport them to the end user.

Before exploring this new paradigm, let us review the classical point-to-point scenario: a source is acquired, processed and communicated, the typical example being speech processing and transmission. From the source point of view, the bandwidth is key for acquisition, and so is the rate-distortion function for compression. From the channel point of view, the bandwidth is critical, as is the capacity-cost function. Underlying both source and channel coding is the separation principle, which allows to exchange bits between the source and the channel.

In the new, distributed scenarios, both the sources and channels are distributed, and separation based coding is usually sub-optimal. Consider the sources: signals are rarely point sources, rather, they are distributed phenomena spatially sampled by acquisition devices (e.g. temperature, sound fields). Thus, the critical value for acquisition is the spatio-temporal bandwidth. Historically, many phenomena were highly undersampled spatially, but the possibility of having dense sensor networks opens up the possibility of sampling at the Nyquist rate.

The structure of the spatio-temporal data is critical to develop efficient compression schemes. Examples of such structures are

already found in stereo images and sounds, and become even more relevant as the number of sensors increases. The compression can either be centralized or decentralized as is more often the case in sensor networks. In the former case, the structure of the data can be fully utilized, while in the latter case, distributed source coding needs to be applied.

The channel for transmission is typically a many-to-many channel, for example with shared bandwidth like in the wireless case. The attempted transmission might be many-to-one (like in many sensor applications), leading to a multi-access channel. In any case, such multiuser channels can lead to difficult questions regarding fundamental limits on capacity, and an effective method to use the channels. But even worse, since the separation principle does not hold, it is unclear that maximizing the capacity of the multi-user channel will ultimately lead to minimizing the distortion for the signals to be transmitted.

Finally, at the receivers, one or many reconstructions of the sensed signals are performed. The quality of such reconstructions will depend on all the steps described above, as well as on the reconstruction algorithms.

To be more specific, let us consider a multicamera system in a room. The structure of the visual data is given by the plenoptic function [1], and sampling results indicate (approximately) the number of cameras required to be able to synthesize any view [2]. Now, if the cameras are not connected by a cheap communication channel, they need to perform "independent" compression but using prior knowledge about correlations or dependencies in the data. The download of the video to a central observer uses a multiaccess channel. Moreover, the channel could be split into successive aggregations, where recoding could be performed. Finally, at the receiver, some arbitrary views (that is, as seen from a position between actual cameras) can be synthesized, possibly using super-resolution techniques.

As can be seen from this example, there is a wealth of interesting and challenging signal processing and communications issues to be addressed. To sum up, the problems central to distributed signal processing and communications concern:

- distributed signal acquisition and sampling
- representation of dependent data (e.g. plenoptic and plenacoustic functions)
- distributed compression of correlated and dependent data
- transmission and joint source-channel coding
- reconstruction of distributed signals.

In the present paper, we will address the above points, and stress the fundamental interaction of sources and channels in distributed scenarios.

<sup>1</sup>This work was supported in part by the Swiss National Competence Center in Research on Mobile Information and Communication Systems, www.nccr-mics.ch.

The outline of the paper is as follows. Section 2 is concerned with distributed signals and their structure, in particular the plenoptic and plenacoustic function, as well as the sampling of these functions. Section 3 reviews correlated source coding, and indicates some recent results on correlated data gathering. Finally, Section 4 discusses the distributed Karhunen-Loève transform.

## 2. THE STRUCTURE OF DISTRIBUTED SIGNALS

Given a certain physical phenomena, the signals that are possible are potentially very structured. Consider the very familiar example of linear time-invariant systems, like for example a room and its associated acoustic field. If the source is a pure sinusoid, then whatever the location of a microphone, it will pick up a sinusoid of the same frequency, the only unknowns being the amplitude and phase. More generally, if many sensors pick up a phenomenon, strong dependencies may exist between that various sensor data. In the sequel, we investigate two cases of particular interest, namely light fields and acoustic fields. In both cases, the first question to solve is the sampling question: is a sufficient representation possible from a countable or finite set of sensors?

### 2.1. The plenoptic function and its sampling

The plenoptic function was first described by Adelson and Bergen [1]. Let us consider the simpler problem of the number of camera along a line that are necessary to synthesize any view along the line. To study this, consider the model of a pinhole camera, which describes an idealization of the projective geometry performed by a camera. Thus, all rays from the real world (more precisely the one visible) pass through the pinhole and are projected on the film plane (which can be in the front or the back of the pinhole). Now consider the following simple set up: the camera can be located anywhere on a line  $t$ , and looks perpendicular to the line. The film plane is parallel to the line  $t$ , and indexed by  $s$  (for now we ignore the second dimension). This set up is shown schematically




Fig. 1. Plenoptic function: setup.

in Figure 1, part (a). In part (b), the mapping of two points is shown, leading to two lines in the plenoptic space. The key point is that the slope of the line depends on the distance of the point from the focal point, and if there is a finite depth of field, there is a finite range of slopes. This is the key insight that leads to a sampling theorem for the plenoptic function, as derived by Shum et al. [2]. In a nutshell, the Fourier transform will be limited angle as well, and if it decays sufficiently fast, a sampling can be done while retaining most of the energy, see Figure 2. From such samples, the plenoptic function can be well interpolated, and thus, arbitrary views can be synthesized. Substantial work along this line has been done see for example [2].

Our work in this area led to both a positive and a negative result. First, if the observed scene is made of simple objects, but not




Fig. 2. Spectrum of the plenoptic function.

necessarily bandlimited, then one can use recent results on sampling of finite rate of innovation signals [3] and sampling of the Radon transform of such objects [4]. In particular, plenoptic sampling is a particular case of the Radon sampling, and it was shown in [4] how a finite number of projections and a finite number of samples allowed to reconstruct perfectly a set of Diracs. Such a result can be seen as a superresolution approach to some parametric plenoptic functions.

As far as bandlimitedness is concerned, it can be shown that in general, the plenoptic function is not bandlimited. For the same argument, consider a smooth bandlimited wall, on which a bandlimited function has been painted. This seems a very well behaved environment. Unfortunately, unless the wall is linear, the resulting plenoptic function will not be bandlimited [5]. The argument relies on time scaling of bandlimited functions, but can be intuitively understood by the following argument: assume the wall is sinusoidal, in which case the plenoptic function relate to frequency modulation. This in turn leads to Bessel function, which are not bandlimited.

### 2.2. Plenacoustic function

Similarly to the idea of the plenoptic function [1, 2], we have introduced the plenacoustic function. It characterizes the sound field in space, e.g. inside a room. We are interested in the room impulse responses to characterize what one would hear at any point in the room. Knowing this information, we can simply calculate the convolution of the sound produced by the source at some point in the room with the room impulse response from the source's position (S) to the listener's position (R). The plenacoustic function is thus parameterized by the following factors:

$$P_a(S \text{ and } R \text{ position, characteristics of room, time}).$$

As it is impossible to store the whole information contained in the plenacoustic function, we need to sample this function. First, we sample the room impulse responses at a certain temporal sampling rate depending on the desired audio bandwidth. Further, by taking an evenly spaced finite number of impulse responses, we uniformly sample the plenacoustic function in space. The study of the plenacoustic sampling answers the following question: "How many microphones do we need to place in space in order to completely reconstruct the sound field at any position in space?"

We consider (without loss of generality) the plenacoustic function along a line in the room. By calculating the 2-dimensional Fourier transform of this data, we obtain Figure 3. The figure shows a bandlimited triangular spectrum. We can see that the spatial frequency ( $\omega_d$ ) support grows for increasing temporal frequencies ( $\omega_t$ ). Spatial sampling of the plenacoustic function at some spatial sampling frequency will lead to aliasing for all the temporal

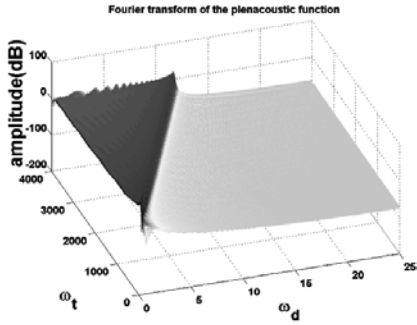


Fig. 3. 2-dimensional Fourier transform of plenacoustic function.

frequencies above some corresponding temporal frequency. Similarly to the Shannon theorem, we can develop a sampling theorem about the necessary spatial sampling frequency in order to reconstruct the plenacoustic function up to some temporal frequency. To reconstruct the plenacoustic function, we use the usual interpolation techniques. We can also show that using a quincunx sampling grid divides the number of samples to process by a factor 2. In [6], we also find a mathematical derivation of the plenacoustic function and experimental results matching the developed theory are presented.

### 3. POWER EFFICIENT DATA GATHERING

In this section, we explore one particular case of source channel interaction, namely the question of how source correlation structure influences the tree building process in a data gathering application.

The scarce resource for autonomous sensor networks is the battery power, which is mostly used when transmitting data. The goal is to find an optimal communication structure, which minimizes the overall power cost of the transmission. The *data gathering problem* considers the case when there is one node to which all data should arrive (the base station), and all other nodes are information sources.

We will make a set of assumptions to make our model tractable to analysis. We assume nodes are fixed, and know their locations. We consider the link-based channel model, which corresponds to the case of e.g. uni-directional antennae. The network measures a property (e.g. temperature) of a random data field, at locations corresponding to the positions of its nodes. We assume measured data is spatially correlated, but i.i.d. in time. Denote by  $X_i$  the random variable that represents the data measured at node  $i$ . The correlation in the data nodes is independent of the distance between the nodes. We will assume moreover that the amount of information at any node can be coded with a fixed number of bits:  $H(X_i) = R, H(X_i|X_j, \dots) = r$  for all  $i, j \neq i$ , with  $0 \leq r \leq R$ . Let  $\rho = 1 - r/R, 0 \leq \rho \leq 1$  be the correlation coefficient. When  $\rho$  is close to 1, the data are strongly correlated; when  $\rho$  is close to 0, the data are independent. Our goal is to obtain a spanning tree for the network graph, that minimizes

$$\rho \sum_{l \in L} d_{ST}(l, S) + (1 - \rho) \sum_{i \in N} d_{ST}(i, S) \quad (1)$$

where  $L$  are the leaf nodes in the tree and  $d_{ST}(i, S)$  the sum of weights of edges connecting node  $i$  to  $S$  on the  $ST$  tree. When  $\rho = 0$  the optimal tree is the shortest path tree (which is known to

be solvable in polynomial time by e.g. a distributed Bellman-Ford algorithm). When  $\rho = 1$ , the optimal solution is a spanning tree for which the sum of paths from the leaves to the base station is minimum (*multiple traveling salesman problem*), and finding it is known to be NP-hard. In [7] we show that the problem is NP-hard in the general case  $0 < \rho \leq 1$ . We only state here the main result, a detailed proof can be found in [7].

#### Definition 1 MINIMUM POWER GATHERING TREE

INSTANCE: A undirected graph  $G = (V, E)$  with lengths  $d_{i,j}$  assigned to the edges  $\{i, j\} \in E$ , a positive integer  $M$ , and a particular node  $S \in V$ .

QUESTION: Does the graph admit a spanning tree  $ST$  such that the total cost of  $ST$  as defined in expression (1) is at most  $M$ ?

**Theorem 1** There is no polynomial time algorithm that solves the MINIMUM POWER GATHERING TREE problem, unless  $P=NP$ .

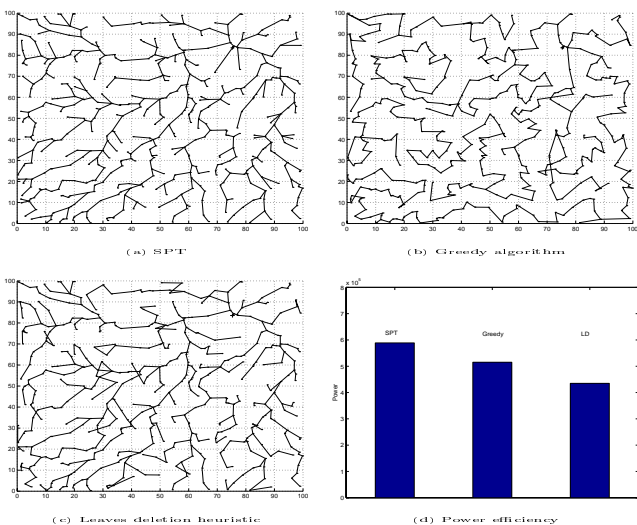
One practical scenario where such a setting is useful is the case of a distributed sensor network, where the sensors measure and/or monitor some environmental data with strong correlation structure (e.g. temperature). In the set-up phase, autonomous sensors are dropped on the area of interest, and they find their respective locations in a distributed fashion. They build a data gathering tree using a distributed algorithm (e.g. the shortest path tree, or the minimum spanning tree). The sensors have to gather data over a long period of time, so power is the main resource that needs to be preserved. The algorithms that take into advantage correlation in the data can outperform standard algorithms for power efficient data gathering. Algorithms to find good approximate tree solutions for expression (1) may provide significant improvements in the power consumption. The heuristic we propose builds in a simple and distributed manner a structure which is more power efficient than standard algorithms. In our simulations with random networks, the power consumption is reduced with up to 40%, so the life time of the network can be increased accordingly. An example is shown in Figure 4.

### 4. DISTRIBUTED COMPRESSION

In this section, we consider the distributed signal processing problem depicted in Figure 5: Multiple sensors (in the figure, three) measure each a part of a set of correlated random variables  $X$ . We assume the covariance matrix of  $X$  to be known and fixed throughout, and for the purpose of this article, we also assume  $X$  to be a vector of jointly *Gaussian* random variables. More general cases are studied in [8]. Sensor  $i$  senses  $M_i$  of the random variables of  $X$ . Each sensor independently furnishes a certain approximation of what it measured to a central decoder whose goal is to produce an estimate  $\hat{X}$  of the entire vector  $X$  in such a way as to minimize the mean-squared error  $E\|X - \hat{X}\|^2$ . We consider two approximation scenarios:

1. Sensor  $i$  furnishes a  $k_i$ -dimensional approximation to its measured data,  $k_i \leq M_i$ . What is the best approximation space? For a given set  $\{k_1, k_2, \dots\}$ , what is the smallest achievable distortion?
2. Sensor  $i$  provides a bit sequence of rate  $R_i$  bits per sample. For a rate-tuple  $\{R_1, R_2, \dots\}$ , what is the smallest achievable distortion?

Special versions of this problem can be solved in closed-form, including the case of the *partial KLT*, by which we denote the case

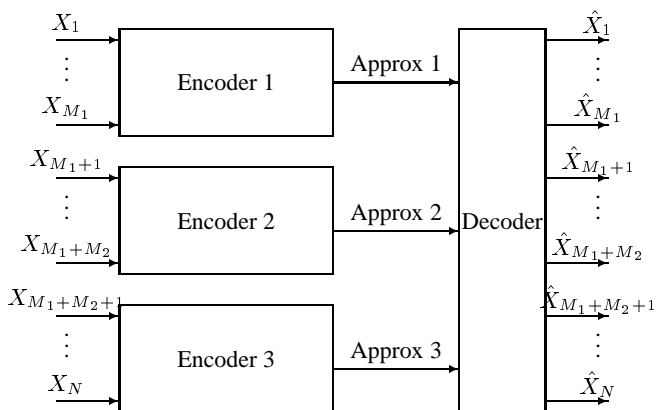


**Fig. 4.** Data gathering trees on a network with  $N = 500$  nodes: (a) shortest path tree, (b) greedy approach, (c) leaves deletion, (d) total power [power unit  $\times$  distance unit<sup>2</sup>]. If the sources transmitted data directly to the base station, the total power would be one order of magnitude larger.

when  $k_j = 0$ ,  $j \geq 2$ , and the *conditional KLT*, by which we denote the case when  $k_j = M_j$ ,  $j \geq 2$ .

For the case of the *partial KLT*, sensor 1 measures the correlated random vector  $X_S = (X_1, X_2, \dots, X_{M_1})$ , with covariance matrix  $\Sigma_S$ . Denote the remaining part of  $X$  by  $X_{S^c} = (X_{M_1+1}, \dots, X_N)$ . Since  $X$  is Gaussian, we can write  $X_{S^c} = AX_S + V$ , where  $X_S$  and  $V$  are independent. The best  $k_1$ -dimensional approximation that sensor 1 can furnish is in a subspace of the KLT of  $X_S$ , i.e., of the eigendecomposition of  $\Sigma_S$ . Out of its eigenvectors, the best  $k_1$  choices are the ones corresponding to the  $k_1$  largest *modified* eigenvalues  $(1 + a_i)\sigma_i^2$ , where  $\sigma_i^2$  are the eigenvalues of  $\Sigma_S$ , and  $a_i = \sum_{j=1}^{M_1} |(AP^{-1})_{ji}|^2$ , where  $P$  is the KLT transform matrix. For the compression problem, a solution can be found along similar lines, see [8].

For the case of the *conditional KLT*, denote the conditional co-



**Fig. 5.** The distributed KLT problem

variance matrix of  $X_S$  given  $X_{S^c}$  by  $\Sigma_{S|S^c}$ , which exists since  $X$  is Gaussian. The best  $k_1$ -dimensional approximation that sensor 1 can furnish is a subspace of the conditional KLT of  $X_S$  given  $X_{S^c}$ , i.e., of the eigendecomposition of  $\Sigma_{S|S^c}$ . Out of its eigenvectors, the best  $k_1$  choices are the ones corresponding to the  $k_1$  largest eigenvalues  $\lambda_i^2$ . For the compression problem, a solution can be found along similar lines, see [8].

These two results can be combined as follows. Suppose that sensors 2, 3,  $\dots$  have fixed their  $k_i$ -dimensional approximation space. What is the best  $k_1$ -dimensional approximation that sensor 1 can furnish? This scenario is equivalent to a scenario where some sensors furnish all of their measured data, i.e.,  $k_j = M_j$ , and others furnish nothing at all, i.e.,  $k_j = 0$ . This is the *combination* of the partial and the conditional KLT. Collect the approximations furnished by sensors 2, 3,  $\dots$  in the vector  $Y'_{S^c}$ , and the discarded (or hidden) parts in the vector  $Y''_{S^c}$ . By the joint Gaussianity of the components of  $X$ , we can write  $Y'_{S^c} = A_1 X_S + A_2 Y''_{S^c} + V$ , where  $V$  is independent of both  $X_S$  and  $Y''_{S^c}$ . The solution for the combined problem can be stated in the following shape: The best  $k_1$ -dimensional approximation that sensor 1 can furnish is in a subset of the eigenvectors of the conditional covariance matrix  $\Sigma_{X|Y'_{S^c}}$ . The best subset is the one corresponding to the  $k_1$  largest modified eigenvalues  $(1 + c_i)\lambda_i^2$ , where  $c_i = \sum_{j=1}^{N-M_1-k_2-k_3-\dots} |(A_1 C^{-1})_{ji}|^2$ , where  $C$  is the conditional KLT transform matrix.

Finally, consider again the problem depicted in Figure 5. The above combination of the partial and the conditional KLT suggests an algorithm to determine the best solution: Initialize each sensor's operation by selecting an arbitrary orthonormal transform and using the first  $k_i$  components as the approximation furnished to the decoder. Then, in turn, each sensor updates its transform and its choice of  $k_i$  components. Numerical studies and the convex shape of the involved functionals suggest a rapid convergence of this procedure. A formal proof of convergence is under way.

## 5. REFERENCES

- [1] E.H. Adelson and J.R. Bergen, "The plenoptic function and the elements of early vision," in *Computational Models of Visual Processing*, pp. 3–20. MIT Press, 1991.
- [2] J.X. Chai, X. Tong, S.C. Chan, and H.Y. Shum, "Plenoptic sampling," in *Proceedings of the conference on Computer graphics*, 2000, pp. 307–318.
- [3] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, June 2002.
- [4] I. Maravic and M. Vetterli, "A sampling theorem for the radon transform of finite complexity objects," in *Proc. of IEEE ICASSP*, 2002.
- [5] D. Marchand-Maillet, M. Do and M. Vetterli, "On the bandlimitedness of the plenoptic function," 2002, in preparation.
- [6] T. Ajdler and M. Vetterli, "The plenacoustic function, sampling and reconstruction," in *IEEE ICASSP*, 2003, submitted.
- [7] R. Cristescu and M. Vetterli, "Power efficient gathering of correlated data: optimization, np-completeness and heuristics," submitted to *MobiHOC* 2003.
- [8] P. L. Dragotti, M. Gastpar and M. Vetterli, "The distributed Karhunen-Loève transform," 2003, to be submitted.