

DISTRIBUTED TARGET LOCALIZATION VIA SPATIAL SPARSITY

Volkan Cevher*, Marco F. Duarte, and Richard G. Baraniuk

Department of Electrical and Computer Engineering
Rice University, Houston, TX 77005

ABSTRACT

We propose an approximation framework for distributed target localization in sensor networks. We represent the unknown target positions on a location grid as a sparse vector, whose support encodes the multiple target locations. The location vector is linearly related to multiple sensor measurements through a sensing matrix, which can be locally estimated at each sensor. We show that we can successfully determine multiple target locations by using linear dimensionality-reducing projections of sensor measurements. The overall communication bandwidth requirement per sensor is logarithmic in the number of grid points and linear in the number of targets, ameliorating the communication requirements. Simulations results demonstrate the performance of the proposed framework.

1. INTRODUCTION

Target localization using a set of sensors presents a quintessential parameter estimation problem in signal processing. Many design challenges arise when the sensors are networked wirelessly due to the limited resources inherent to the sensor network. For example, any inter-sensor communication exerts a large burden on the sensor batteries. Since sufficient statistics are often non-existent for the localization problem, accurate localization requires the full collection of the network sensing data. Thus, the choice of algorithms is usually steered away from those achieving optimal estimation. To increase the lifetime of the sensor network and to provide scalability, low dimensional data statistics are often used as inter-sensor messages, such as local range or bearing estimates at the sensors. Hence, the sensor network localization performance is sacrificed so that the sensors can live to observe another day.

To improve the estimation performance and robustness of the sensor network in the presence of noise over classical maximum likelihood and subspace methods, sparsity based localization have been slowly gaining popularity [1–5]. The main idea in these papers is that under specific conditions [6], the localization estimates can be obtained by searching for the sparsest solution of under-determined linear system-of-equations that frequently arise in localization. In this context,

a vector is called *sparse* if it contains only a small number of non-zero components in some transform domain, e.g, Fourier or wavelets. The ℓ_0 -norm is the appropriate measure of the sparsity, which simply counts the number of non-zero elements of a vector. Unfortunately, minimizing the ℓ_0 -norm is *NP*-hard and becomes prohibitive even at moderate dimensions. At the cost of slightly more observations, it has been proven that ℓ_1 -norm minimization results in the same solution and has computational complexity on the order of the vector dimensions cubed [7, 8].

We formulate the localization problem as the sparse approximation of the measured signals in a specific dictionary of atoms. The atoms of this dictionary are produced by discretizing the space with a localization grid and then synthesizing the signals received at the sensors from a source located at each grid point. We show how this *localization dictionary* can be locally constructed at each sensor. Within this context, the search of the sparsest approximation to the received signals that minimizes the data error implies that the received signals were generated by a small number of sources located within the localization grid. Hence, our algorithm performs successful source localization by exploiting the direct relationship between the small number of sources present and the corresponding sparse representation for the received signals. We assume that the individual sensor locations are known *a priori*; however, the number of sources need not be known. The resulting sparse approximation problem can be solved using greedy methods such as orthogonal matching pursuit [9] or other solvers such as fixed point continuation methods [10].

Since we are interested in distributed estimation over wireless channels where minimizing communications is crucial, we discuss how to solve the localization problem when lower dimensional projections of the sensor signals are passed as inter-sensor messages. To preserve the information content of these messages, a projection matrix must be chosen so that it is incoherent with the sparsifying basis, i.e., the localization dictionary. Fortunately, a matrix with independent and identically distributed (i.i.d.) Gaussian entries satisfies the incoherence property with any fixed basis with high probability [11]. Based on results from compressive sensing (CS) [7, 8], we show that the total number of samples that are needed for recovering the locations of K targets is $\mathcal{O}(K \log(N/K))$, where N is the number of grid points. We also show that the total number of bits encoding the sensor measurements that must be communicated can be made quite small with graceful degradation in performance.

* Corresponding author. This work was supported by the grants DARPA/ONR N66001-06-1-2011 and N00014-06-1-0610, NSF CCF-0431150 and DMS-0603606, ONR N00014-07-1-0936, AFOSR FA9550-07-1-0301, ARO W911NF-07-1-0502, ARO MURI W911NF-07-1-0185, and the Texas Instruments Leadership University Program. E-mail: {volkan,duarte,richb}@rice.edu. Web: dsp.rice.edu/cs.

Given that (i) each sensor has a localization dictionary, (ii) we would like to localize the target locations within a certain resolution as defined by N , and (iii) the total number of targets is much less than the number of grid points (the sparsity assumption), an estimate of multiple target locations within our framework can be done when each sensor receives at least $\mathcal{O}(K \log(N/K))$ samples. This implies that the resolution of the grid and the expected number of targets rather than the number of sensors define the communication bandwidth, hence the proposed localization framework is scalable and is suitable for distributed estimation. Compared to the other distributed estimation methods such as belief propagation [12–14], our approach does not require a refinement process where local message passing is continued until the sensor network reaches convergence. However, such scheme can still be used to improve the localization accuracy within our framework. Compared to the decentralized data fusion [15], our approach does not have data association problems, which is combinatorial in the number of targets K . Moreover, due to the democratic nature of the measurements within our framework, our approach has built in robustness against packet drops commonly encountered in practice in sensor networks. In contrast, when low dimensional data statistics, such as local range and bearing estimates, are used in distributed estimation algorithms, package drops result in significant performance degradation.

Similar to our paper, other sparse approximation approaches to source localization have been proposed before [1–5]. In [1], spatial sparsity is assumed to improve localization performance; however, the computational complexity of the algorithm is high, since it uses the high-dimensional received signals. Dimensionality reduction through principal components analysis was proposed in [2]; however, this technique is contingent on knowledge of the number of sources present for acceptable performance and also requires the transmission of all the sensor data to a central location to perform singular value decomposition. Similar to [2], we do not have incoherency assumptions on the source signals. In [3], along with the spatial sparsity assumption, the authors assume that the received signals are also sparse in some known basis and perform localization in near and far fields; however, similar to [1], the authors use the high-dimensional received signals and the proposed method has high complexity and demanding communication requirements. CS was employed for compression in [4, 5], but the method was restricted to far-field bearing estimation. In contrast, this paper extends the CS-based localization setting to near-field estimation, and examines the constraints necessary for accurate estimation in the number of measurements and sensors taken, the allowable amount of quantization, the spatial resolution of the localization grid, and the conditions on the source signals.

The paper is organized as follows. Section 2 lays down the theoretical background for CS, which is referred in the ensuing sections. Construction of the sensor localization dictionaries is described within the localization framework in Sect. 3. Section 4 describes the spatial estimation lim-

its of the proposed approach such as minimum grid spacing or maximum localization grid aperture. Section 5 discusses communication aspects of the problem including the message passing details and the bandwidth requirements. Finally, simulation results demonstrating the performance of the localization framework are given in Sect. 6.

2. COMPRESSIVE SENSING BACKGROUND

CS provides a framework for integrated sensing and compression of discrete-time signals that are sparse or compressible in a known basis or frame. Let \mathbf{z} denote a signal of interest, and Ψ denote a sparsifying basis or frame, such that $\mathbf{z} = \Psi\theta$, with $\theta \in \mathbb{R}^N$ being a K -sparse vector, i.e. $\|\theta\|_0 = K$. Transform coding compression techniques acquire first \mathbf{z} in its entirety, and then calculate its sparse representation θ in order to encode its nonzero values and their locations. CS aims to preclude the full signal acquisition by measuring a set \mathbf{y} of linear projections of \mathbf{z} into vectors ϕ_i , $1 \leq i \leq M$. By stacking these vectors as rows of a matrix Φ , we can represent the measurements as $\mathbf{y} = \Phi\mathbf{z} = \Phi\Psi\theta$. The main result in CS states that when the matrix $\Phi\Psi$ holds the restricted isometry property (RIP) [8], then the original sparse representation θ is the unique solution to the linear program

$$\hat{\theta} = \arg \min_{\theta \in \mathbb{R}^N} \|\theta\|_1 \text{ s.t. } \mathbf{y} = \Phi\Psi\theta, \quad (1)$$

known as Basis Pursuit [6]. Thus, the original signal \mathbf{z} can be recovered from the measurement vector \mathbf{y} in polynomial time. Furthermore, choosing Φ to be a matrix with independent gaussian-distributed entries satisfies the RIP for $\Phi\Psi$ when Ψ is a basis or tight frame and $M = \mathcal{O}(K \log(N/K))$. Recovery from noisy measurements can be performed using Basis Pursuit Denoising (BPDN), a modified algorithm with relaxed constraints. We employ a fixed point continuation method [10] to solve the BPDN optimization efficiently.

3. LOCALIZATION VIA SPATIAL SPARSITY

In a general localization problem, we have $L + 2$ parameters for each of the targets at each estimation period: the 2D coordinates of the source location and the source signal itself, which has length L . In general, the estimation of these parameters are entangled: the source signal estimate depends on the source location, and viceversa. Our formulation can localize targets without explicitly estimating the source signal, therefore reducing computation and communication bandwidth.

Assume that we have K sources in an isotropic medium with P sensors with known positions $\zeta_i = [\zeta_{xi}, \zeta_{yi}]'$ ($i = 1, \dots, P$) on the ground plane. We do not assume that the number of sources K is known. Our objective is to determine the multiple target locations $\chi_i = [\chi_{xi}, \chi_{yi}]'$ using the sensor measurements. To discretize the problem, we only allow the unknown target locations to be on a discrete grid of points $\varphi = \{\varphi_n | n = 1, \dots, N; \varphi_n = [\varphi_{xn}, \varphi_{yn}]'\}$. By performing this discretization and limiting the number of sources to be localized, the localization problem can be cast as a sparse

approximation problem of the received signal, where we obtain a sparse vector $\theta \in \mathbb{R}^N$ that contains the amplitudes of the sources present at the N target locations. Thus, this vector only has K nonzero entries. We refer to this framework as *localization via spatial sparsity* (LVSS).

Define a linear convolution operator for signal propagation, denoted as $\mathcal{L}_{\chi \rightarrow \zeta}$, which takes the continuous signal for a source at a location χ and outputs the L samples recorded by the sensor at location ζ , by taking into account the physics of the signal propagation and multipath effects. Similarly, define the pseudoinverse operator $\mathcal{L}_{\zeta \rightarrow \chi}^\dagger$ that takes an observed signal at a location ζ and deconvolves to give the source signal, assuming that the source is located at χ . A simple example operator that accounts for propagation attenuation and time delay can be written as

$$\mathcal{L}_{\chi \rightarrow \zeta}(\mathbf{x}) = \left[\frac{1}{d_{\chi, \zeta}^\alpha} \mathbf{x} \left(\frac{l}{F_s} - \frac{d_{ik}}{c} \right) \right]_{l=1}^L,$$

where $d_{\chi, \zeta}$ is the distance from source χ to sensor ζ , c is the propagation speed, α is the propagation attenuation constant, and F_s is the sampling frequency for the L samples taken.

Additionally, denote the signal from the k^{th} source as \mathbf{x}_k . We can then express the signal received at sensor i as $\mathbf{z}_i = \mathbf{X}_i \theta$, where

$$\mathbf{X}_i = [\mathcal{L}_{\chi_1 \rightarrow \zeta_i}(\mathbf{x}_1) \quad \mathcal{L}_{\chi_2 \rightarrow \zeta_i}(\mathbf{x}_2) \quad \dots \quad \mathcal{L}_{\chi_N \rightarrow \zeta_i}(\mathbf{x}_N)]$$

is called the i^{th} sensor's source matrix. Similarly, we can express the signal ensemble as a single vector $\mathbf{Z} = [\mathbf{z}_1^T \dots \mathbf{z}_P^T]^T$; by concatenating the source matrixes into a single dictionary

$$\Psi = [\mathbf{X}_1^T \quad \mathbf{X}_2^T \quad \dots \quad \mathbf{X}_1^T]^T, \quad (2)$$

the same sparse vector θ used for each signal generates the signal ensemble as $\mathbf{Z} = \Psi \theta$.

An estimate of the j^{th} sensor's source matrix \mathbf{X}_j can be determined using the received signal at a given sensor i . If we assume that the signal \mathbf{z}_i observed at sensor i is originated from a single source location, we can then write

$$\hat{\mathbf{X}}_{j|i} = \left[\mathcal{L}_{\chi_1 \rightarrow \zeta_j}(\mathcal{L}_{\zeta_i \rightarrow \chi_1}^\dagger(\mathbf{z}_i)) \quad \dots \quad \mathcal{L}_{\chi_N \rightarrow \zeta_j}(\mathcal{L}_{\zeta_i \rightarrow \chi_N}^\dagger(\mathbf{z}_i)) \right].$$

Furthermore, we can obtain an estimate $\hat{\Psi}_i$ of the signal ensemble sparsity dictionary by plugging in the source matrices estimates into (2).

Thus, by having each sensor transmit its own received signal \mathbf{z}_i to all other sensors in the network (or to a central processing unit), we can then apply a sparse approximation algorithm to \mathbf{Z} and $\hat{\Psi}_i$ to obtain an estimate of the sparse location indicator vector $\hat{\theta}_i$ at sensor i . By using CS theory, we can reduce the amount of communication by having each sensor transmit $M = \mathcal{O}(K \log(N/K))$ random projections of \mathbf{z}_i instead of the L -length signal.

4. RESOLUTION OF THE GRID

The dictionary obtained in this fashion must meet the conditions for successful reconstruction using sparse approximation algorithms. A necessary condition was posed in [16]:

Theorem 1 [16] *Let $\Psi \in \mathbb{R}^{L \times N}$ be a dictionary and ψ_j denote its j th column. Define its coherence $\mu(\Psi)$ as*

$$\mu(\Psi) = \max_{1 \leq j, k \leq T, j \neq k} \frac{|\langle \psi_j, \psi_k \rangle|}{\|\psi_j\| \|\psi_k\|},$$

Let $K \leq 1 + 1/16\mu$ and let $\Phi \in \mathbb{R}^{M \times L}$ be a matrix with i.i.d. Gaussian-distributed entries, where $M \geq \mathcal{O}(K \log(N/K))$. Then with high probability, any K -sparse signal θ can be reconstructed from the measurements $y = \Phi \Psi \theta$ through the ℓ_1 minimization (1).

Thus, the coherence of the dictionary used by the sensor controls the maximum number of localizable sources. Define the normalized cyclic autocorrelation of a signal \mathbf{z} as

$$R_{\mathbf{z}}[m] = \frac{\sum_{n=1}^L z(t_n) z(t_{\text{mod}[(n+m), L]})}{\|\mathbf{z}\|^2}.$$

Then $\mu(\Psi_i)$ depends on $R_{\mathbf{z}}[m]$, since

$$\frac{|\langle \Psi_{i,j}, \Psi_{i,k} \rangle|}{\|\Psi_{i,j}\| \|\Psi_{i,k}\|} = \frac{\sum_{p=1}^P \frac{R_{\mathbf{z}_i} \left[\frac{F_s}{c} (d_{\chi_j, \zeta_p} - d_{\chi_k, \zeta_p} - d_{\chi_j, \zeta_i} + d_{\chi_k, \zeta_i}) \right]}{(d_{\chi_j, \zeta_p} d_{\chi_k, \zeta_p})^\alpha}}{\sqrt{\left(\sum_{p=1}^P d_{\chi_j, \zeta_p}^{-\alpha} \right) \left(\sum_{p=1}^P d_{\chi_k, \zeta_p}^{-\alpha} \right)}}.$$

The coherence μ will thus depend on the maximum value attained by $R_{\mathbf{z}_i} \left[\frac{F_s}{c} (d_{\chi_j, \zeta_p} - d_{\chi_k, \zeta_p} - d_{\chi_j, \zeta_i} + d_{\chi_k, \zeta_i}) \right]$; we assume that the cyclic autocorrelation function is inversely proportional to its argument's absolute value. The coherence then depends on the minimum value of the function's argument. In the location grid setting, this minimum is approximately $\Delta/2D$, with Δ denoting the grid spacing, and D denoting the maximum distance between a grid point and a sensor. Such maximum distance D is dependent on both the extension of the grid and the diameter of the sensor deployment.

In summary, to control the maximum coherence, it will be necessary to establish lower bounds for the localization resolution – determined by the grid spacing – and upper bounds for the extension of the grid and the diameter of the sensor deployment.

5. INTER-SENSOR COMMUNICATIONS

Compared to distributed estimation algorithms that use a single low dimensional data statistic from each sensor, the sparsity based localization algorithms [1–3] require the collection of the observed signal samples to a central location. Hence, for a sensor network with single sensors, a total of $P \times L$ numbers must be communicated as opposed to, for example, P received signal strength (RSS) estimates. Since L is typically a large number, the lifetime of a wireless sensor network would be severely decreased if such a scheme is used.

Considering the lifetime extension, the performance degradation in target localization is considered a fair tradeoff.

Starting with the knowledge of the localization dictionary Ψ at any given sensor, CS results state that to perfectly recover a K sparse vector in N dimensions, $\mathcal{O}(K \log(N/K))$ random projections of \mathbf{Z} are needed. This can easily be achieved at each sensor by multiplying the sensed signal by a pre-determined random projection matrix before communication, effectively resulting in a block diagonal measurement matrix structure [17]. Thus, the dominant factor of the communication bandwidth becomes the number of grid points, as opposed to the number of sensors. As an example, consider $L = 1000$, $N = 100^2$, $K = 5$, $P = 100$: $P \times L = 10^5$ vs. $K \log(N/K) \approx 38$. When compared to distributing the full sensor network data, there is a significant reduction; however, LVSS is still not competitive with sending an RSS estimate per sensor.¹ However, when RSS estimates are sent in the presence of multiple targets, signal interference effects and data association issues decrease the localization performance. In general, the estimated localization dictionary is noisy; hence, a larger number of measurements is needed.

Another way of understanding the minimum required inter-sensor communications is to use information theoretic arguments to determine the minimum number of bits required to localize K -coordinates in an N -dimensional space: we need $K \log_2 N$ bits to encode this information. Since we process the received signals to obtain θ , we can only lose entropy. Thus, the resulting $K \log_2 N$ number of bits of our analysis presents a lower bound the number of bits that each sensor needs to receive for localization over a grid size N to determine K target locations. Even when quantization is considered for the $\mathcal{O}(K \log(N/K))$ measurements needed by LVSS, there is an evident gap between the lower limit and the LVSS requirement, since LVSS recovers both the location of the nonzero coefficients *and* their values.

The aforementioned gap can be explored via quantization of the CS measurements. It is known within the CS framework that compressive measurements are robust against quantization noise as the CS reconstruction is robust against additive noise [18]. Thus, we obtain two degrees of freedom to determine the message size required in LVSS. In practice with simulated and field data, we have found that assigning 1-bit to encode the mean of the absolute values of the compressive measurements is effective in recovery (see also [4]). In this quantization scheme, the sensors pass the sign of the compressive measurements as well as the mean of their absolute values. Hence, the inter-sensor messages would incorporate one additional number, which also needs to be quantized, encoding the quantization level, along with 1-bit messages encoding the sign of the measurements.

6. SIMULATIONS

Our objectives in this section are two fold. We first demonstrate the distributed estimation capabilities of the proposed

¹We assume that there is no communication overhead. If there is some overhead in sending messages, then LVSS becomes competitive.

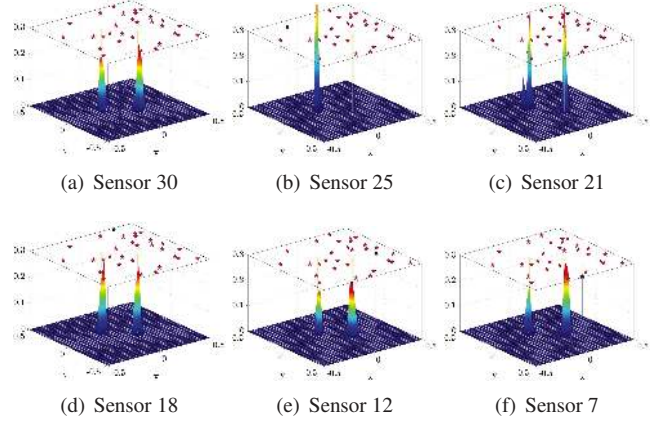


Figure 1: Distributed estimation results: each sensor obtains localization estimates independently from random measurements received from all sensors in the network. The results are similar for most sensors.

framework. We then examine the effects of the inter-sensor communication message sizes and signal-to-noise (SNR) ratio on the performance of the algorithm.

Our simulation setup consists of $P = 30$ sensor nodes sensing two coherent targets that transmit a standard signaling frame in MSK modulation with a random phase shift. The sent signals have length $L = 512$ and a unit grid of $N = 30 \times 30$ points is used for localization, where the speed of propagation $c = 1$. For each simulation, a fixed point continuation solver [10] was used for the sparse approximations. The algorithm employs a parameter μ that weights the goodness of fit of the solution against its sparsity; this parameter is fixed for all simulations at all sensors. We note that when the number of compressive measurements change, adjusting this parameter can improve the localization results.

In the first experiment, we study the dependence of the localization performance on the choice of sensor. We fix the number of measurements per sensor $M = 30$ and set the SNR to 20dB. Figure 1 illustrates the sparse approximation results at a representative subset of the sensors. In the figure, the sensors are represented by filled stars at the ceiling, and the ground truth for the source locations is represented by the yellow asterisks. The surface plots show the output of the sparse approximation, each normalized so that they sum up to 1, defining a PDF of the multitarget posterior. The figure shows consistent localization PDFs at the different sensors. Within the 30 sensor network, a few sensors miss one of the targets (e.g., sensor 25). Note that these PDFs are calculated at each sensor independently after receiving the compressive measurements from the network.

In the second experiment, we study the dependence of the localization performance on the number of measurements per sensor M and the SNR. For each combination of these parameters, we performed a Monte Carlo simulation involving 100 realizations of a uniformly random sensor deployment, as well as 50 realizations of Gaussian noise per deployment. The location estimates in each Monte Carlo run are obtained using K -means clustering on the estimated θ ,




Figure 2: Results from Monte Carlo simulations. Top row: $M = 2$, bottom row: $M = 10$. From left to right, SNR = 0dB, 5dB, 30dB.

with the number of clusters equal to the number of targets. Figure 2 shows scatter plots for the localization estimates for the different setups, together with the root mean square (RMS) error and the likelihood of divergence (Div.) in the sparse approximation algorithm. Intuitively, the figure shows improvement in performance as the SNR or the number of measurements increases. Reducing the number of measurements, however, increases the likelihood of divergence in the sparse reconstruction. For Fig. 2(a-c) each sensor receives $58 > K \log(N/K) \approx 12$ measurements for localization. In general, this increase is due to the noisy localization dictionary estimates in the presence of multiple targets and the block-diagonal nature of the measurement matrix.

7. CONCLUSIONS

Our fusion of existing sparse approximation techniques for localization and the CS framework enables the formulation of a communication-efficient distributed algorithm for target localization. LVSS exhibits tolerance to noise, packet drops and quantization, and provides a natural distributed estimation framework for sensor networks. The performance of the algorithm is dependent on both the number of measurements and the SNR, as well as the observed signal, the sensor deployment and the localization grid. Furthermore, the algorithm performance can be improved by increasing the number of measurements taken at each of the sensors, providing a tradeoff between the communication bandwidth and the accuracy of estimation. Future work will investigate the fundamental limits of localization within the sparsity framework and compare the sparsity based localization algorithms with other state-of-the-art distributed localization algorithms to provide a Pareto frontier of the localization performance as a function of communications. We also plan to study the inclusion of signal sparsity into our framework.

REFERENCES

[1] I. F. Gorodnitsky and B. D. Rao, "Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum

norm algorithm," *IEEE Transactions on Signal Processing*, vol. 45, no. 3, pp. 600–616, 1997.

- [2] D. Malioutov, M. Cetin, and A. S. Willsky, "A sparse signal reconstruction perspective for source localization with sensor arrays," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 3010–3022, 2005.
- [3] D. Model and M. Zibulevsky, "Signal reconstruction in sensor arrays using sparse representations," *Signal Processing*, vol. 86, no. 3, pp. 624–638, 2006.
- [4] V. Cevher, A. C. Gurbuz, J. H. McClellan, and R. Chelappa, "Compressive wireless arrays for bearing estimation," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Las Vegas, NV, Apr. 2008.
- [5] A. C. Gurbuz, V. Cevher, and J. H. McClellan, "A compressive beamformer," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Las Vegas, NV, 2008.
- [6] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM Journal on Scientific Computing*, vol. 20, pp. 33, 1998.
- [7] D. L. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, September 2006.
- [8] E. J. Candès, "Compressive sampling," in *Int. Congress of Mathematicians*, Madrid, Spain, 2006, vol. 3, pp. 1433–1452.
- [9] J. Tropp and A. C. Gilbert, "Signal recovery from partial information via orthogonal matching pursuit," *IEEE Trans. Info. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.
- [10] E. T. Hale, W. Yin, and Y. Zhang, "Fixed-point continuation for ℓ_1 minimization: Methodology and convergence," Tech. Rep. TR07-07, Rice University Department of Computational and Applied Mathematics, Houston, TX, 2007.
- [11] E. J. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, no. 3, pp. 969–985, June 2007.
- [12] K. Murphy, Y. Weiss, and M. I. Jordan, "Loopy belief propagation for approximate inference: An empirical study," *Proceedings of Uncertainty in AI*, pp. 467–475, 1999.
- [13] J. Yedidia, W. T. Freeman, and Y. Weiss, "Generalized belief propagation," *Advances in Neural Information Processing Systems*, vol. 13, pp. 689–695, 2001.
- [14] A. T. Ihler, *Inference in Sensor Networks: Graphical Models and Particle Methods*, Ph.D. thesis, Massachusetts Institute of Technology, 2005.
- [15] J. Manyika and H. Durrant-Whyte, *Data Fusion and Sensor Management: A Decentralized Information-Theoretic Approach*, Prentice Hall, Upper Saddle River, NJ, 1995.
- [16] H. Rauhut, K. Schnass, and P. Vandergheynst, "Compressed sensing and redundant dictionaries," *IEEE Trans. Info. Theory*, vol. 54, no. 5, pp. 2210–2219, May 2008.
- [17] D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk, "Distributed compressed sensing," Available at <http://www.dsp.rice.edu/cs>, 2005.
- [18] E. J. Candès and J. Romberg, "Encoding the ℓ_p ball from limited measurements," in *Proc. IEEE Data Compression Conference (DCC)*, Snowbird, UT, March 2006.