

Adaptive Compressive Sensing in Smart Water Networks

Sokratis Kartakis ^{1,*}, George Tzagkarakis ^{2,3} and Julie A. McCann ¹

¹ Imperial College London, UK; E-Mails: s.kartakis13@imperial.ac.uk; j.mccann@imperial.ac.uk

² EONOS Investment Technologies, Paris, France; E-Mail: gtzag@eonos.com

³ Foundation for Research and Technology-Hellas, Institute of Computer Science/SPL, Greece

* Author to whom correspondence should be addressed; E-Mail: s.kartakis13@imperial.ac.uk;
Tel.: +44-207-594-8296.

Published: 10 November 2015

Abstract: Contemporary water distribution networks exploit information communication technologies (ICT) to monitor and control the behavior of water network assets. Limited capability and typically battery powered low-resourced devices, such as smart meters/sensors, have been used to transfer information from the water network to data centers for further analysis. Many water companies deploy devices aiming to last beyond the 10-year mark. This prohibits the use of high-sample rate sensing therefore limiting the knowledge we can obtain from this data. However, data reduction techniques with minimal information loss can overcome this problem. In this paper, we present a self-adaptive scheme that reduces the amount of transmitted data, thus extending the battery life of sensor nodes, while still maximizing the received information to data centers. To achieve these goals, we exploit the power of compressive sensing (CS), which enables significant compaction of the original information content in a few random incoherent projections. Sparsity of the recorded data streams, which is a necessary condition for successful CS reconstruction, is achieved via the transformation of the original data into an appropriate transform domain. Using over 170 days of real high-sample rate water pressure data from 25 sensor nodes of our large scale testbed in the Bristol area, we verify the efficiency of our CS-based algorithm in significantly reducing the data volume, and thus extending the battery life of sensor nodes. In addition, we demonstrate that our system supports self-tuning and automatic reconfiguration as the nature of incoming data changes over time.

Keywords: Wireless Sensor and Actuator Networks; Smart Water Network; Compressive Sensing

1. Introduction

In recent years, there is a trend for water utility companies to create smart water networks [1, 2], in order to improve the quality of service, reduce water waste through balancing the water supply and demand, and minimize maintenance costs by increasing the network resilience. The work presented in this paper is part of a Smart Water project that both monitors water distribution networks (WDN) and controls its valves to optimize water network performance and lifetime over varying demands. A large scale testbed, which consists of 25 sensor nodes, was implemented in Bristol Water utility company water network shown in Figure 1. Each sensing device records, analyzes, and transmits high sample rate water pressure and flow data (up to 128 samples per second) of District Metering Areas (DMAs) to a data processing center for further analysis. These high sample rates result in ever increasing amounts of data, thus making in-node data processing, such as compression, a necessity prior to storage or transmission. Under this context, in the past, we have examined, evaluated, and deployed lossless compression techniques [3] in our sensing devices, which reduce transmitted data and consequently extend battery life proportionally. In this paper, we introduce contemporary lossy compression approaches and examine their efficiency.

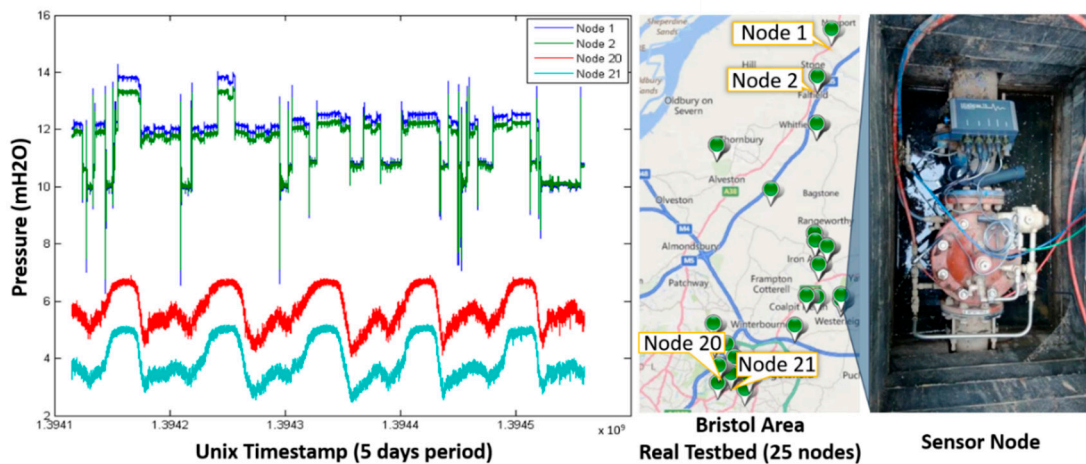


Figure 1. Bristol area testbed and pressure data from 4 sensor nodes.

For decades, the sampling process has been largely dominated by the classical Nyquist-Shannon theories. However, several studies have shown that many natural signals are amenable to highly sparse representations in appropriate transform domains (*e.g.*, wavelets and sinusoids) [4, 5]. Compressive sensing (CS) provides a powerful framework for simultaneous sensing and compression [6], enabling a significant reduction in the sampling and computation costs on a sensor node with limited memory and power resources. According to the theory of CS, a signal having a sparse representation in a suitable transform domain can be reconstructed from a small set of incoherent random projections.

In this study, the advantages of CS are exploited for onboard compression and recovery at a base station of high-resolution pressure data recorded by sensors deployed in a water distribution network. Our experimental evaluation reveals the high performance of our proposed approach, when compared with lossless compression schemes such as [3], in terms of achieving much higher compression ratios, while maintaining highly accurate reconstructions of the original sensor data. Additionally, to our

knowledge, this is the first time that CS is being adapted as a compression technique to Smart Water Network high resolution data.

The rest of the paper is organized as follows. Section 2 describes the methodology of applying CS in water data. Section 3 presents our evaluation, while Section 4 concludes this paper and gives directions for future extensions.

2. Methods

In this section, the design characteristics of the two main components are described, namely, the CS-based module for resolution pressure data (up to 128 samples per sec) compression, which executes on the sensor nodes (see Figure 2 left part), and the decompression module that reconstructs the data, which is implemented on a base station where all the sensor data are gathered (see Figure 2 right part).




Figure 2. Information transformation flow diagram.

2.1. CS-based Data Compression

Let $\Psi \in \mathbb{R}^{N \times L}$ be a matrix whose columns correspond to a transform basis, in general overcomplete (i.e. $N \leq L$). In terms of signal approximation, it has been shown that if a signal $\mathbf{x} \in \mathbb{R}^N$ is K -sparse in basis Ψ , meaning that the signal is exactly or approximately represented by K elements (columns) of this basis, then it can be reconstructed from $M = \mathcal{O}\left(K \log\left(\frac{N}{K}\right)\right)$ non-adaptive linear projections onto a second measurement basis, which is incoherent with the sparsity basis¹ [6, 7]. The general measurement model in the sparsifying transform domain is expressed as

$$\mathbf{y} = \Phi \mathbf{x} = \Phi \Psi^T \mathbf{w} \in \mathbb{R}^M, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^M$ is the vector of compressive measurements, $\Phi \in \mathbb{R}^{M \times N}$ is the measurement matrix whose columns are random vectors with independent and identically distributed (i.i.d.) components (i.e., each random component has the same probability distribution as the others, and the occurrence of one does not affect the probability of the other), Ψ^T denotes the inverse transform, and $\mathbf{w} \in \mathbb{R}^L$ is the sparse coefficient vector. Typical examples of measurement matrices, which are incoherent with any fixed transform basis with high probability (universality property [7]), include random matrices with i.i.d. Gaussian or Bernoulli entries.

2.2. CS-based Data De-Compression

By employing the M compressive measurements and given the K -sparsity property in the transform basis, the original signal \mathbf{x} can be recovered by solving an appropriate constrained optimization problem. Commonly used reconstruction techniques are based on convex relaxation [6] and greedy strategies (e.g.,

¹ the vectors of the measurement basis are not statistically correlated with the vectors of the sparsity basis

Orthogonal Matching Pursuit [8]). In this study, the NESTA algorithm is used, which was shown to achieve a very good trade-off between reconstruction accuracy and computation time [9]. We emphasize that the scope of this paper is to illustrate the efficiency of the CS framework in achieving highly compact, yet very accurate, representations of real high-resolution sensor data recorded in water distribution networks. As such, an exhaustive comparison with the various reconstruction algorithms for finding the optimal solution is left as a separate study.

Focusing again on the optimization problem to be solved for reconstructing the original data, NESTA solves a synthesis-based problem of the following form:

$$\min_{\mathbf{w} \in \mathbb{R}^L} \|\mathbf{w}\|_1 \quad \text{subject to} \quad \|\mathbf{y} - \Phi\Psi^T\mathbf{w}\|_2 \leq \epsilon, \quad (2)$$

where $\|\cdot\|_1$, $\|\cdot\|_2$ denote the l_1 and l_2 norm, respectively. Having estimated the sparse coefficient vector $\hat{\mathbf{w}}$, a reconstruction of the original signal is simply obtained by taking the inverse transform, that is,

$$\hat{\mathbf{x}} = \Psi^T\hat{\mathbf{w}}. \quad (3)$$

In the subsequent experimental evaluation, the discrete wavelet transform (DWT) [4] will be applied on the raw sensor data due to its computational efficiency and sparse representation accuracy. However, we note that the CS-based approach expressed by Eqs. (1) – (3) is generic and can be applied with alternative sparsifying transformations, other than the DWT.

3. Results and Discussion

In this section, the performance of the CS-based approach is evaluated and compared with well-established lossless compression techniques for compressing data streams recorded by a set of sensors deployed in the Bristol Water water distribution network. More specifically, the available dataset consists of high sample-rate pressure data (64 samples per second) from 25 sensor nodes for a 170-day period. For sake of brevity, this section presents the evaluation results for four pressure data streams from an equal number of sensor nodes as shown in Figure 1.

In the first test case, as [3] describes, data compressed using various lossless compression techniques. The applicable compression method for the current hardware infrastructure was MiniLZO [10], which uses sliding window as coding method (LZ77). We repeated the experiments in [3] by dividing each pressure data stream into non-overlapping windows of length $N = 1024$ (or, equivalently, $1024/64 = 16$ sec) and applied MiniLZO algorithm for each data chunk by achieving 55% average compression rate.

Despite the fact that lossless compression allows the original data to be perfectly reconstructed from the compressed data, this is typically achieved at lower compression rates. To overcome these limitations, the acquired data are compressed via highly reduced sets of random measurements (ref. Eq. (1)) and reconstructed by solving the constrained optimization problem expressed by Eqs. (2) – (3). To this end, each pressure data stream is divided again into non-overlapping windows of length $N = 1024$. Then, a multi-scale decomposition in 10 levels is applied on each window using the DWT with the “db4” wavelet function. Subsequently, the resulting wavelet coefficients are projected onto the rows of a measurement matrix Φ with i.i.d. standard Gaussian entries (mean 0, and standard deviation 1). The efficiency of the proposed CS-based scheme is tested for a varying number of measurements, that is rows of Φ , $M = \rho \cdot N$, where $\rho \in \{0.25, 0.35, 0.45, 0.55\}$. This is equivalent to compressing the data in each window at compression rates $\{75\%, 65\%, 55\%, 45\%\}$, respectively.

The reconstruction accuracy is measured in terms of the root mean squared relative error (RMSRE). Specifically, if \mathbf{x} and $\hat{\mathbf{x}}$ are the original and reconstructed data, respectively, their RMSRE is defined as

$$\text{RMSRE}(\mathbf{x}, \hat{\mathbf{x}}) = \sqrt{\frac{1}{N} \sum_{j=1}^N \left(\frac{x_j - \hat{x}_j}{x_j} \right)^2} \quad (4)$$

Figure 3 shows the RMSRE, averaged over all individual windows, as a function of the CS sampling rate for each one of the four nodes. As it can be seen, the reconstruction accuracy is already high even for high compression rates (75%), whilst it improves (that is, the RMSRE decreases) as the compression rate decreases (or, equivalently, the CS sampling rate increases).




Figure 3. Average RMSRE as a function of the CS sampling rate for the four nodes node 1, node 2, node 20, node 21 ($N = 1024$, “db4”).

Finally, to illustrate the approximation accuracy of the original pressure data, Figure 4 shows segments of the original and reconstructed data streams for sensor node 2. As expected, the approximation accuracy improves as the CS sampling rate increases (see Figure 4b); however, a highly accurate reconstruction is achieved even for low sampling rates (see Figure 4a – approximately ± 0.0025 mH₂O error) with a result of significantly higher compression rates than lossless approached which was 55% and consequently proportional battery life extension.

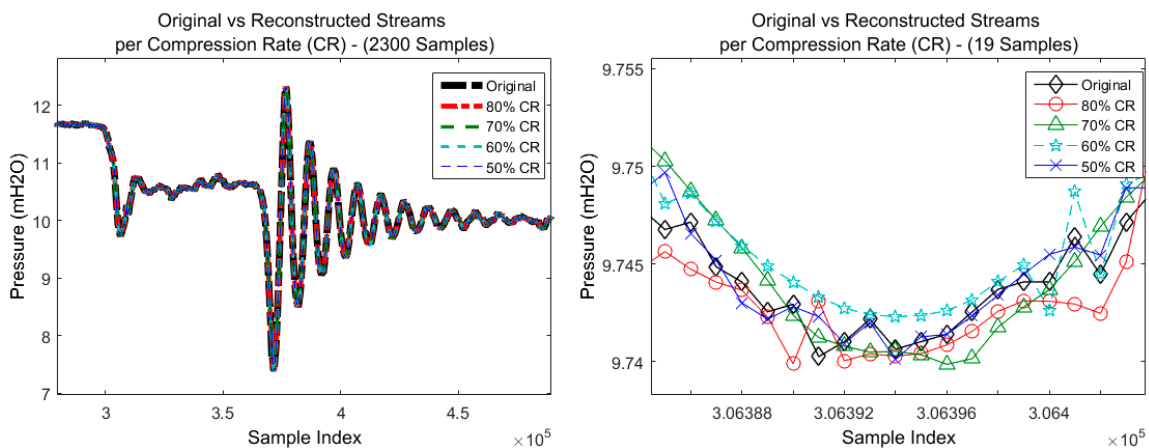


Figure 4. Original and reconstructed segments for sensor node 2 using the NESTA algorithm for (a) 2300 samples (~ 36 sec) and (b) 19 samples (~ 296 msec).

4. Conclusions/Outlook

In this study, the power of compressive sensing was exploited for achieving high compression rates, yet high reconstruction accuracy, of real high sample rate data captured by a set of sensors deployed in Bristol Water water distribution network. The experimental evaluation revealed a superior performance when compared with well-established lossless compression algorithms.

The current approach employs a fixed sparsifying transformation for the sensor data. However, the degree of sparsity can be increased by employing an adaptive sparse representation. To this end, the framework of dictionary learning can be exploited, in conjunction with a joint sparsity assumption to account for correlations either between the windows of each individual data stream or between distinct data streams. Finally, a CS-based approach could be exploited for designing algorithms to detect abnormal sensor behavior by employing the random compressed measurements directly.

Acknowledgments

This work forms part of the Big Data Technology for Smart Water Nets research project funded by NEC Corporation, Japan. The hydraulic pressure datasets used in this paper have been provided by Dr Ivan Stoianov and Mr Asher Hoskins, Dept Civil Engineering, Imperial College London, UK. This work was also partially funded by the European Union's Horizon 2020 research and innovation programme PHySIS under grant agreement No 640174 and the FP7 WISDOM project under grand agreement No 619795.

Conflicts of Interest

The authors declare no conflict of interest.

References

1. Allen, M.; Preis, A.; Iqbal M.; Whittle A. J. Water distribution system monitoring and decision support using a wireless sensor network, in Proc. IEEE SNPD, 2013, pp. 641–646.
2. Stoianov, I.; Nachman, S. M.; Tokmouline, T.; Csail M. Pipenet: A wireless sensor network for pipeline monitoring, in Proc. IPSN 2007, 2007, pp. 264–273.
3. Kartakis, S.; McCann, J. A. Real-time edge analytics for cyber physical systems using compression rates, in ICAC'14, USENIX, pages 153–159, 2014.
4. Mallat, S. *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed. Academic Press, 2009.
5. Daubechies, I. Ten lectures on wavelets, *SIAM*, 1992.
6. Candés, E.; Romberg, J.; Tao, T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, *IEEE Trans. on Information Theory*, 2006, 52, 489 – 509.
7. Donoho, D. Compressed sensing, *IEEE Trans. on Information Theory*, 2006, 52, 1289 – 1306.
8. Tropp, J.; Gilbert, A. Signal recovery from random measurements via orthogonal matching pursuit, *IEEE Trans. on Information Theory*, 2007, 53, 4655 – 466.
9. Becker S.; Bobin. J.; Candés, E. NESTA: A fast and accurate first-order method for sparse recovery, Technical Report, 2009, California Institute of Technology. Available online: <http://statweb.stanford.edu/~candes/nesta/NESTA.pdf>

10. Kraus, J.; Bubla, V. Optimal methods for data storage in performance measuring and monitoring devices. In Proc. of Electronic Power Engineering Conference, 2008.

© 2015 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/4.0/>).