

# Application of Matrix Completion on Water Treatment Data

Sofia Savvaki  
savvaki@csd.uoc.gr

Athanasia  
Panousopoulou  
apanouso@ics.forth.gr

Grigorios Tsagakatakis  
greg@ics.forth.gr

Panagiotis Tsakalides  
tsakalid@ics.forth.gr

Department of Computer Science, University of Crete  
Institute of Computer Science, Foundation for Research and Technology Hellas (FORTH)  
Heraklion, 70013, Greece

## ABSTRACT

Recent advances in Cyber-Physical Systems (CPS) have revolutionized water management in urban areas. Nevertheless, literature reports minor progress in introducing CPS-based systems at industrial water treatment plants, responsible for water purification. Such environments would greatly benefit by adopting CPS technologies in general, and Wireless Sensor Networks (WSNs) in particular. However, WSNs would suffer from a series of industrial monitoring constraints, which would inevitably lead to missing measurements. In this work, we study the problem of efficient estimation of missing water treatment data, collected over a WSN deployed at a water desalination plant, and we examine how redundancies can be used for the recovery of extremely undersampled matrices. We propose the formulation of the problem as an instance of low rank Matrix Completion (MC), and we employ the inexact Augmented Lagrange Multipliers (ALM) algorithm. We consider three key questions related to the performance of our method; namely, recovery from artificially introduced missing entries, single versus collective recovery of measurements matrices, and the real-valued problem of temporal super-resolution. The results highlight that MC is a valid method for estimating missing water-treatment data, even from a very limited number of measurements.

## Keywords

Wireless Sensor Networks, Matrix Completion, Smart Water Networks, Water Treatment

## 1. INTRODUCTION

Smart Water Networks (SWNs) [1] constitute an emerging engineering field which addresses the blend of networked data technologies with water infrastructures, with the objective to deliver sustainable solutions related to water resources utilization. Driven by the application demands of modernizing water quality monitoring, acting upon alerting

events, and improving our awareness of water allocation and consumption, Cyber-Physical Systems (CPS) and Wireless Sensor Networks (WSNs) encapsulate the key-enabling technologies for the next generation of SWN systems [4, 10].

Existing solutions in the arena of SWN are considered extremely useful both for responding in alarming situations [13, 20], as well as for engaging citizens in becoming part of water sustainable policies [9, 16]. In their majority, they are deployed close to urban areas thereby yielding feedback on the actual quality of the water that flows within the pipes, towards the plumbing facilities of the individuals [11, 14, 15, 17]. Despite their significance, such solutions cannot address water purification aspects for human consumption, which is becoming a critical aspect of water management, especially as the natural resources of fresh water become scarcer and the population of urban areas grows. Consequently, the necessity intensifies for expanding the engineering focus towards industrial treatment plants, which are responsible for water purification for human consumption.

Water purification involves the combination of slowly varying physical and chemical processes, for making untreated water suitable for human consumption. Such industrial processes are responsible for monitoring and controlling critical microbiological parameters, based on sparse and sporadic samples, and rely on off-line testing procedures and the involvement of experienced personnel in the control loop.

Employing WSNs in water treatment plants would improve their autonomous character, by introducing novel paradigms of data acquisition and processing. Nevertheless, the respective deployments would be also accompanied by constraints, such as limited computational capabilities, packet losses, and decreased lifetime. From the perspective of the high-level application, these constraints are translated to the lack of sufficient data samples for characterizing different aspects of the water purification process. Apart for improving the robustness and resiliency of the network backbone, modern signal processing techniques can be employed for providing accurate estimations on the missing values. Matrix Completion (MC) is considered as a promising novel approach, capable of recovering measurements in a reliable fashion [8, 19]. The underlying concept of MC is that a complex signal can be recovered from a small number of random measurements, far below the traditional Nyquist-Shannon limit. The key assumption in MC is that the signal is sparse, while sub-sampled matrix-value data are available, therefore making MC appealing for WSN applications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CySWater '15 Seattle, WA USA

Copyright 2015 ACM 978-1-4503-3485-3/15/04

<http://dx.doi.org/10.1145/2738935.2738949> ...\$15.00.

With these considerations on mind, in this work our objective is to assess the efficiency of MC on the framework of water purification, with a strong emphasis on water desalination plants. We depart from the current state of the art by evaluating the efficacy of inexact Augmented Lagrange Multipliers MC theory for WSN-based water treatment data. The evaluation process relies on datasets collected within a water desalination plant. The proposed framework is evaluated against different aspects, which are related to the volume of the available dataset considered for reconstructing missing values, the number of available sensors that are providing measurements, and ultimately the temporal super-resolution aspects that express the relationship of the sampling rate and the operational network characteristics. The presented results provide useful insights on applying the MC framework in water treatment data, and highlight the efficacy of our proposed concept.

The remainder of this paper is organized as follows: In Section 2 we describe the WSN scenarios according to which the collected measurements can be extremely sub-sampled. Section 3 formulates the MC problem. Section 4 extensively describes the adopted methodology and derived results, while the conclusions are drawn in the final Section.

## 2. MISSING WSN MEASUREMENTS

Energy-efficient acquisition and communication of measurements is a critical aspect in WSNs, directly determining the lifetime and usability of the network infrastructure. To achieve strict performance requirements and overcome the limitations of challenging operating environments, we exploit redundancies in data collected over time, which are modeled by low rank measurement matrices. In this work, we focus on three prominent cases where such measurement matrices are extremely under-sampled, *i.e.*, they contain a large number of zero-valued entries.

While frequent sampling offers high quality monitoring of the underlying processes, increasing the frequency of sampling may have a dramatic effect on the lifetime of the network. This effect is attributed to the close relationship between measurement acquisition and energy consumption. Waking-up a node, acquiring a measurement, performing quantization, and storing the data in local memory, are all extremely energy-demanding tasks. If communications with other nodes is also necessary due to, *e.g.*, storage requirements, then the impact on network lifetime is even more pronounced.

A second scenario that entails a large number of missing values is attributed to measurements encoded in packets lost due to communications failures. This scenario occurs typically in industrial environments, where heavy machinery has a detrimental impact on link quality. Furthermore, in multi-hop networks, congestion and duty cycling can also lead to dropped packets and, thus, lost measurements. All of the described scenarios, result in a significant number of missing data that can have a dramatic impact on subsequent tasks, such as detection of unusual events or clustering of the measurements.

A third scenario, is related to the temporal sampling frequency of a WSN. Either by design, or due to clock desynchronization, each sensor may end up sampling the underlying field at a different time instance. We argue that such a limitation could actually be used in our advantage and offer temporal super-resolution of the sampling pro-

cess. To facilitate exposition, consider a network monitoring a field, where the entire collection of measurements can be encoded into a measurements matrix, where rows correspond to sensors and columns to sampling instances. A fully synchronized network of 2 sensors, configured to inquire and record the field every 10 minutes, will produce a  $[2 \text{ (sensors)}] \times [6 \text{ (measurements per sensor)}]$  matrix, with columns corresponding to measurements acquired at 00:00, 00:10, 00:20,... 00:50. However, one sensor could sample at 00:03, 00:13,..., 00:53 while the other could sample at 00:07, 00:17, ..., 00:57. One can exploit the difference in sampling instances to produce a  $[2] \times [12]$  matrix, where columns correspond to timestamps 00:00, 00:05, 00:10, 00:15, ..., 00:55. Such a matrix will naturally miss 50% of its measurements, due to the requested temporal super-resolution.

The above described examples and associated reasoning, serve as a motivation for the introduction of efficient and robust measurements recovery mechanisms. In this work, we consider the recently proposed framework of Matrix Completion to address the aforementioned challenges.

## 3. MATRIX COMPLETION

Given a  $[M] \times [N]$  measurement matrix  $\mathbf{M}$ , recovering the  $M \times N$  entries of the matrix from a smaller number of  $K \ll M \times N$  entries is not possible, in general. However, it was recently shown that the recovery of the complete set of entries in a matrix is possible, provided that both the number of missing entries and the rank of the matrix are appropriately bounded [5], [6]. Formally, let  $\mathcal{A}$  be a linear map from  $\mathbb{R}^{[M] \times [N]} \rightarrow \mathbb{R}^K$ , that selects a subset of the entries in the matrix  $\mathbf{M}$ . The linear map  $\mathcal{A}$ , is defined as a random sampling operator that records a small number of entries from the matrix  $\mathbf{M}$ , that is  $\mathcal{A}_{ij} = \{1 \text{ if } (ij) \in S \mid 0 \text{ otherwise}\}$ , where  $S$  is the sampling set. We can estimate  $\mathbf{X}$  from the original, fully populated, matrix  $\mathbf{M}$ , by solving:

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} \text{rank}(\mathbf{X}) \\ & \text{subject to } \mathcal{A}(\mathbf{X}) = \mathcal{A}(\mathbf{M}) . \end{aligned} \quad (1)$$

Unfortunately, rank minimization is an NP-hard problem and therefore cannot be applied in practice. Recently, a relaxation of the above problem was shown to produce accurate approximations, by replacing the rank constraint with the more computationally tractable nuclear norm, which represents the convex envelope of the rank. The relationship is manifested by the Singular Value Decomposition (SVD) of the  $[M] \times [N]$  measurements matrix, into a product of an orthonormal matrix  $\mathbf{U}$ , a diagonal matrix  $\mathbf{S}$ , and another orthonormal matrix  $\mathbf{V}$ , such that:

$$\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^T. \quad (2)$$

According to the spectral theorem associated with the SVD, the number of singular values, *i.e.*, the diagonal entries of  $\mathbf{S}$ , reveals the rank of the matrix. Low rank matrices, such as the ones produced by spatio-temporally correlated processes, are characterized by a small number of singular values. Furthermore, the rank of a measurement matrix might be artificially increased, due to noise that typically follows an independent distribution. Hence, considering a lower-rank approximation of the matrix results in an implicit de-noising of the sampled data. Based on the SVD analysis of a matrix, the minimization in (1) can be reformulated as:

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} \quad \|\mathbf{X}\|_* \\ & \text{subject to} \quad \mathcal{A}(\mathbf{X}) = \mathcal{A}(\mathbf{M}), \end{aligned} \quad (3)$$

where the nuclear norm is defined as  $\|\mathbf{X}\|_* = \sum |\sigma_i|$ , *i.e.*, the sum of absolute values of the singular values. Recovery of the matrix is possible, provided that the matrix  $\mathbf{M}$  satisfies an incoherence property. The solution of (3) will converge to the solution of (1) with probability  $1 - cn^{-3}$  once  $K \geq Cq^{6/5}r \log(q)$  random matrix entries are obtained, where  $q = \max(M, N)$ .

For the noisy case, an approximate version can be solved [6], by replacing the equality constraint with an inequality constraint given by  $\|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\|_F^2 \leq \epsilon$ , where  $\|\mathbf{X}\|_F^2 = \sum \lambda_i^2$  denotes the Frobenius norm and  $\epsilon$  is the approximation error. The optimization is therefore formulated as:

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} \quad \|\mathbf{X}\|_* \\ & \text{subject to} \quad \|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\| \leq \epsilon. \end{aligned} \quad (4)$$

To solve the nuclear norm minimization problem in (4), various approaches have been proposed. In this work, we employ the Augmented Lagrange Multipliers (ALM) [12], [7] based MC, due to its performance with respect to both computationally complexity and recovery capabilities.

## 4. EVALUATION STUDIES

The objective of this study, is to quantify the effectiveness of the ALM algorithm for MC on WSN-based water desalination data. The data collection took place at the La Tordera pilot desalination plant, owned by Acciona Agua<sup>1</sup>, by deploying a customized WSN solution. The deployment environment and WSN modules are illustrated in Figure 1. Each sensor features an IEEE-802.15.4-based protocol stack, that employs a customized CSMA-based solution for adaptive link scheduling [3] and the IEF Standard for Routing over Low Power Lossy Networks (RPL) [2].



Figure 1: The desalination plant (left) and the WSN module used for the collection of water desalination data (right).

The dataset considered for our evaluation purposes contains water impedance measurements (Ohm), sampled at 10 different channels, by 5 sensors at different stages of the desalination process. WSN measurements were collected for a 3 day period (25th, 26th, 27th October 2014) and the sampling rate of the sensors is 1 measurement per hour per sensor. Hence, our dataset contains in total  $5 \times 10 \times 3 \times 24 = 3600$  measurements. These measurements are grouped by sensors, and we refer to each 'virtual' frequency of each sensor as a *sensing modality*. This allows us to generalize the specific study and its conclusions, to other cases of WSN-based monitoring beyond water treatment.

<sup>1</sup><http://www.acciona-agua.com/>

A key issue regarding the evaluation of performance, is related to the specific metric that we use in order to quantify the results. We employ the commonly used Normalized Mean Square Error (NMSE), defined as the mean squared error between the fully populated and the reconstructed measurements matrix, normalized with respect to the  $l_2$  norm.

In addition, another very important aspect of the experimental process, is related to the amount of missing measurements. We evaluate the MC recovery for different cases of missing measurements, ranging from relatively 'dense' matrices lacking only few measurements to 'sparse' matrices where the majority of measurements are unobserved. Towards this direction we consider the *fill ratio*  $f$ , defined as the ratio of the amount of non-zero elements over the number of all the entries of the measurements matrix of dimensions  $[n1] \times [n2]$ :

$$f = \frac{\#non-zero\ elements}{[n1] \times [n2]}.$$

### 4.1 Effects of measurement matrix dimensions

In this experiment, we investigate the MC's recovery abilities with respect to the value of  $f$  for the measurements matrix. The objective is to assess how the ALM based recovery performance is associated with the size of the data, by evaluating the recovery regarding four different sizes of measurements matrices, generated from the initial fully sampled  $50 \times 72$  matrix. Therefore, we incrementally reduce the dimensionality by a factor of two, by assuming that for each day the number of sampling instances, *i.e.*, columns of the matrix, is halved, while the number of rows, *i.e.*, number of modalities, remains constant. As a result, the measurements matrix dimensions at the second iteration is  $50 \times 36$  (2 hours sampling interval), at third iteration is  $50 \times 18$  (4 hours sampling interval), etc. For each data size,  $f$  was initiated at 0.1 and was iteratively incremented up to 0.9, with a step size equal to 0.02.

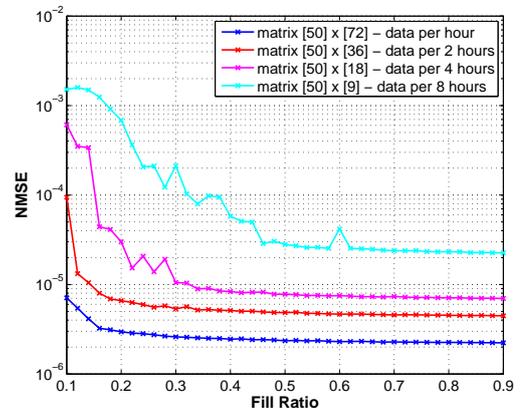


Figure 2: Normalized MSE as a function of  $f$  with respect to 4 different measurements matrix sizes.

Figure 2 illustrates the recovery performance measured by the NMSE as a function of  $f$ . We can observe that the increase in  $f$  has a pronounced effect on the reconstruction quality. Regardless the size of the measurements matrix, the effect of missing values remains crucial to the reconstruction. The results also suggest that the value of NMSE converges at high values of  $f$ , for all 4 sizes of measurements matrix. Moreover, it is clearly shown that, for larger data matrices, the convergence of NMSE is much smoother.

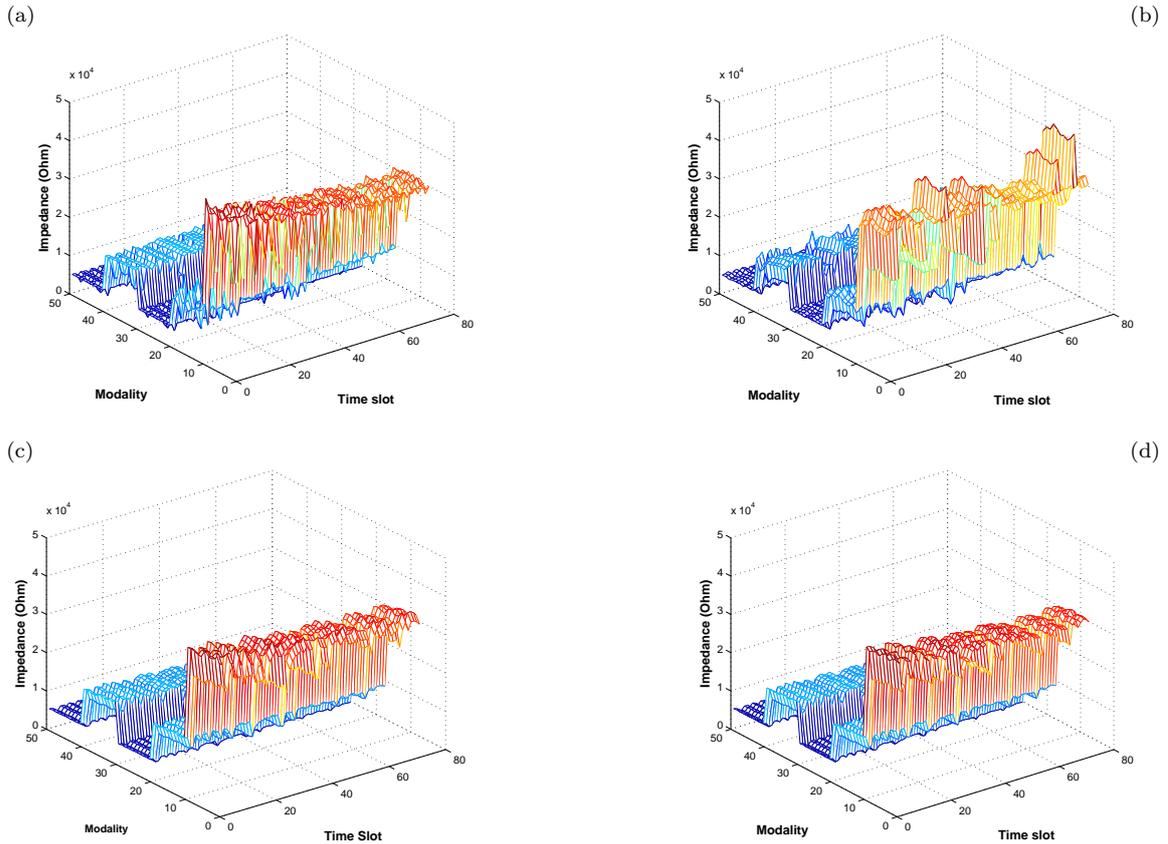


Figure 3: The MC results for: (a) original fully sampled measurements matrix, (b) reconstructed matrix from 10% of the measurements, (c) reconstructed matrix from 50% of the measurements, (d) reconstructed matrix from 90% of the measurements.

Figure 3(a) depicts a fully sampled  $[50] \times [72]$  measurements matrix and Figures 3(b)-3(d) present the MC reconstructed matrices when  $f$  equals to 0.1, 0.5, and 0.9 respectively. It is straightforward to observe, that higher fill ratios lead to more accurate measurements reconstruction, in accordance with both the theoretical models, as well as the quantified results presented at Figure 2. Furthermore, we observe that although the reconstruction at 0.1 exhibits noisy artifact, one can still get a overall sense of the behaviour of the data. Finally, it is shown that even 50% of the measurements, suggesting a dramatic reduction in sampling requirements, can produce very accurate estimations.

## 4.2 Effects of number of sensors

In this experiment, we evaluate the performance of MC based reconstruction when considering data from a single sensor, as opposed to scenarios where collective recovery is performed using measurements from all the sensors. More specifically, we assume that each sensor has access to its locally stored sub-sampled measurements matrix, and cross-examine the MC accuracy on two different options. The first option, defined as the *single sensor recovery* case, considers that MC based recovery takes place on each sensor locally, by employing only its own data. The second option, defined as the *collective recovery* case, incorporates a central processing station and applies MC to the collective measurements matrix corresponding to the concatenation of the individual matrices.

Recovery performance with respect to  $f$ , is shown in Figures 4 and 5, for three sensors with id 1,2, and 5 respectively.

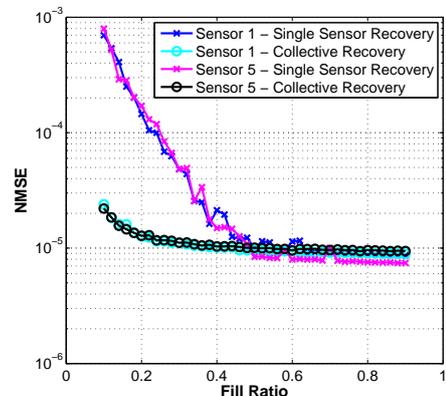


Figure 4: Single sensor case vs. collective reconstruction for  $[50] \times [72]$  measurements matrix from sensors 1 and 5.

As far as sensor 1 is concerned, we observe that high reconstruction quality in the single sensor case can indeed be accomplished using just local measurements. Nevertheless, collective recovery achieves better reconstruction, over all different values of  $f$ . This result suggests that collective MC recovery can fully utilize the correlation that exists among sensors, even if such correlations are not explicitly encoded into the recovery process, thus highlighting the generalization ability of the proposed schemes.

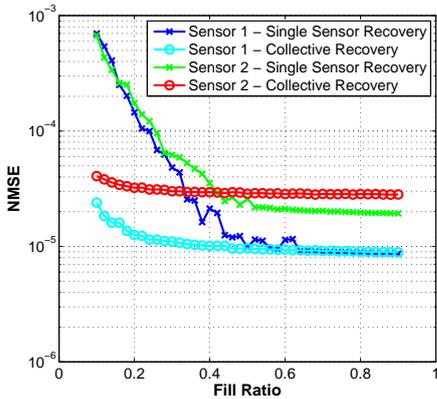


Figure 5: Single sensor case vs. collective reconstruction for  $[50] \times [72]$  measurements matrix from sensors 1 and 2.

Regarding sensors 2 and 5, we observe that, as we move on to higher values of  $f$ , *i.e.*, above 0.4 and 0.5 respectively, the single-sensor ALM matrix completion achieves better reconstruction results. This behavior is different than the case of sensor 1, where the reconstruction quality in the collective recovery scenario is better compared to single sensor case for high values of  $f$ . The experimental results suggest that by exploiting the intra-sensor correlation, collective reconstruction can achieve superior performance compared to the single sensor case for most situations, although at high sampling rates, the performance gain is marginal.

### 4.3 MC for temporal super-resolution

In the final set of experimental results, we attempt to approach the problem of missing measurements from a different, yet more realistic angle. Unlike the previous results, in this case the zero-valued entries are not randomly introduced according to a specified fill ratio. Instead, they naturally arise as we increase the temporal sampling rate beyond the operating characteristics of the network. More specifically, the  $[50] \times [72]$  fully populated measurements matrix corresponds to a sampling rate of 1 measurement every 60 minutes, for all 5 sensors and 10 frequency channels, which is the maximum sampling rate of our testbed. Increasing this sampling rate to, *e.g.*, one measurement every 30 minutes, leads to an increase in the dimensionality of the measurements matrix to  $[50] \times [2 \times 72]$ . This indicates that the number of columns is doubled, while the number of rows remains constant. Since we are operating beyond the temporal sampling capabilities of the network, there are not enough measurements to fill the expanded measurements matrices, according to timestamps on which the measurements were obtained. Subsequently, zero-valued measurements are introduced at cells. Moreover,  $f$  is set to 0.5. Table 1 illustrates the dimensionality and corresponding values of  $f$  for the sampling rates used in this experiment.

In this scenario, it is not possible to calculate the NMSE metric, since there is not a reference measurements matrix. Nevertheless, we conduct this experiment, in order to visually assess the performance of MC recovery. This corresponds to the realistic scenarios, where the reference measurements matrix will not be available. We performed the MC reconstruction experiment for 4 different matrix sizes presented in Table 1. Results are shown in Figures 6(a)-(d).

Since we super-resolve in the temporal domain, the lack

Table 1: Dimensionality and fill ratio for different sampling rates.

Sampling rate	Dimensionality	Fill ratio ( $f$ )
60 minutes	$[50] \times [72]$	1
30 minutes	$[50] \times [144]$	0.5
15 minutes	$[50] \times [288]$	0.25
10 minutes	$[50] \times [432]$	0.16

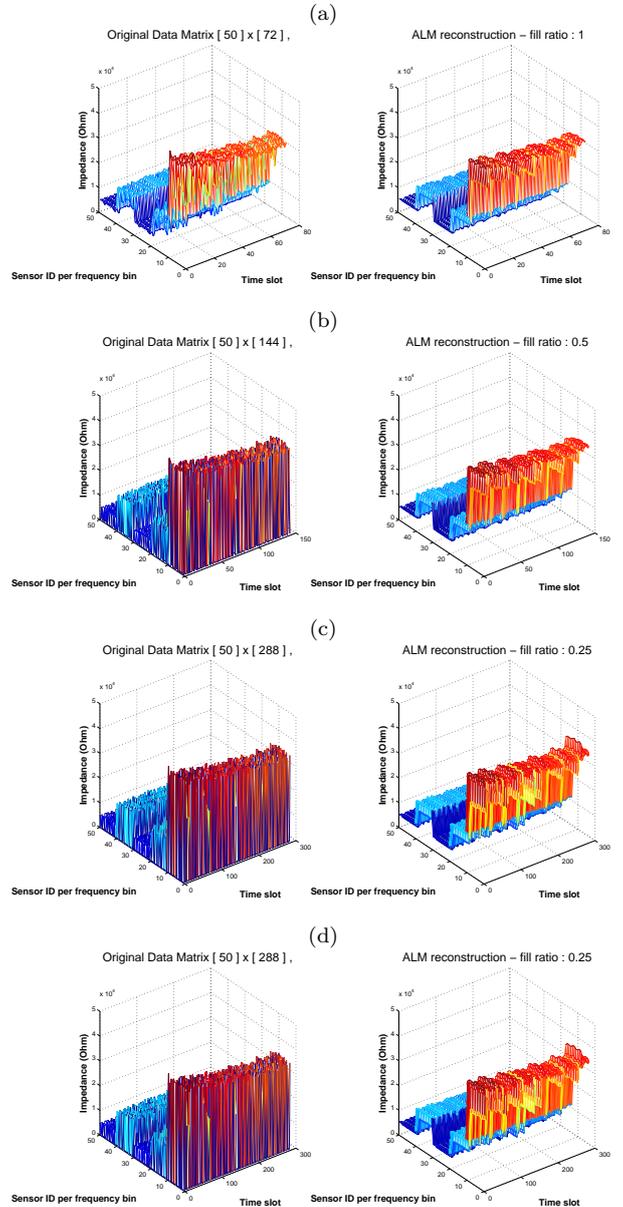


Figure 6: The initial (left) and reconstructed (right) when the dimensions of the measurements matrix are: (a)  $[50] \times [72]$ , (b)  $[50] \times [144]$ , (c)  $[50] \times [288]$ , (d)  $[50] \times [512]$ .

of ground truth data means that we cannot estimate the performance via some error metric. However, the visual observations made suggest that, while the dimensionality of the matrices increases and the fill ratio decreases, the MC reconstructed data maintain their smoothness and distribution to a great extent, compared to the initial full measurements

matrix of dimensions  $[50] \times [72]$ . This gives a fairly good intuition, as far as the efficiency of ALM matrix completion is concerned, on the performance of the proposed scheme in truly lost or unavailable measurements.

## 5. CONCLUSIONS

In this work, we investigated the application of Matrix Completion for the estimation of missing measurements from a water treatment plan. We considered three scenarios of paramount importance for the current operation of a WSN based monitoring paradigm; namely recovery from artificially introduced missing entries, single sensor vs. collective recovery of measurements matrices, and temporal super-resolution. Based on our experimental findings, we can conclude that MC is a viable approach for estimating missing measurements, where such missing measurements are either attributed to lost/unobserved measurements or, in the case of super-resolving, to non-acquired data. Our proposed scheme is very generic by nature and can be easily applied to the recovery of other types of measurements, or be introduced as an energy-efficient sampling architecture.

## Acknowledgments

This work was supported by the HYDROBIONETS project (ICT-2011-7) funded by the European Commission in FP7 (GA-2011-287613) and co-financed by the European Union and Greek national funds through the National Strategic Reference Framework (NSRF), Research Funding Program: “Cooperation-2011”, Project “SeNSE”, grant number: 11SYN-6-1381, and the GSRT O.P. Competitiveness and Entrepreneurship PEFYKA project within the KRIPIS action of the GSRT. We are also grateful to Acciona Agua<sup>2</sup> for providing the premises of La Tordera’s desalination plant, as well as to Ateknea Solutions<sup>3</sup> and CNM (Centre Nacional de Microelectrònica)<sup>4</sup> for assisting in the collection of WSN-based water desalination data.

## 6. REFERENCES

- [1] Smart water networks. <http://www.swan-forum.com/>.
- [2] Routing protocol for low power and lossy networks, 2011.
- [3] D. Alonso-Roman, E. Celada-Funes, C. Asensio-Marco, and B. Beferull-Lozano. Improving reliability and efficiency of communications in wsns under high traffic demand. In *Wireless Communications and Networking Conference (WCNC), 2013 IEEE*, pp. 268–273, 2013.
- [4] C. Beal and J. Flynn. Toward the digital water age: Survey and case studies of australian water utility smart-metering programs. *Utilities Policy*, 32:29–37, 2015.
- [5] E. J. Candès and T. Tao, “The power of convex relaxation: Near-optimal matrix completion”, *Information Theory, IEEE Transactions*, 56(5):2053–2080, 2010.
- [6] Candès, E. and Plan, Y., “Matrix completion with noise”, *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [7] Chen, Caihua, Bingsheng He, and Xiaoming Yuan, “Matrix completion via an alternating direction method”, *IMA Journal of Numerical Analysis*, 32(1):227–245, 2012.
- [8] J. Cheng, Q. Ye, H. Jiang, D. Wang, and C. Wang. Stcdg: An efficient data gathering algorithm based on matrix completion for wireless sensor networks. *Wireless Communications, IEEE Transactions on*, 12(2):850–861, 2013.
- [9] J. Harou, P. Garrone, A. Rizzoli, A. Maziotis, A. Castelletti, P. Fraternali, J. Novak, R. Wissmann-Alves, and P. Ceschi. Smart metering, water pricing and social media to stimulate residential water efficiency: Opportunities for the smarth2o project. *Procedia Engineering*, 89:1037–1043, 2014.
- [10] A. Hauser and F. Roedler. Interoperability—the key for smart water management. 2014.
- [11] T. T.-T. Lai, W.-J. Chen, K.-H. Li, P. Huang, and H.-H. Chu. Triopusnet: Automating wireless sensor network deployment and replacement in pipeline monitoring. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks*, pp. 61–72, 2012.
- [12] Lin, Zhouchen, Minming Chen, and Yi Ma, “The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices”, *arXiv preprint:1009.5055*, 2010.
- [13] S. S. Mane and M. K. Mokashi. Survey on real-time flash-flood monitoring, alerting and forecasting system using data mining and wireless sensor network. *International Journal*, 2(12), 2014.
- [14] N. Metje, D. N. Chapman, D. Cheneler, M. Ward, and A. M. Thomas. Smart pipe: instrumented water pipes, can this be made a reality? *Sensors*, 11(8):7455–7475, 2011.
- [15] B. O’Flynn, F. Regan, A. Lawlor, J. Wallace, J. Torres, and C. O’Mathuna. Experiences and recommendations in deploying a real-time, water quality monitoring system. *Measurement Science and Technology*, 21(12):124004, 2010.
- [16] T. Robles, R. Alcarria, D. Martin, A. Morales, M. Navarro, R. Calero, S. Iglesias, and M. Lopez. An internet of things-based model for smart water management. In *Proceedings of 28th International Conference on Advanced Information Networking and Applications Workshops*, pp. 821–826, 2014.
- [17] I. Stoianov, L. Nachman, A. Whittle, S. Madden, and R. Kling. Sensor networks for monitoring water supply and sewer systems: Lessons from boston. In *Water Distribution Systems Analysis Symposium 2006*, pp. 1–17, 2008.
- [18] N. Tsiftes, J. Eriksson, and A. Dunkels. Low-power wireless ipv6 routing with contikirpl. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN ’10*, pp. 406–407, 2010.
- [19] L. Xiang, J. Luo, C. Deng, A. V. Vasilakos, and W. Lin. Deca: Recovering fields of physical quantities from incomplete sensory data. In *Processings of the 9th IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, pp. 182–190, 2012.
- [20] G. Xu, G. Q. Huang, and J. Fang. Cloud asset for urban flood control. *Advanced Engineering Informatics*, 2015.

<sup>2</sup><http://www.acciona-agua.com/>

<sup>3</sup><http://ateknea.com/>

<sup>4</sup><http://www.imb-cnm.csic.es/>