Title: The Application of Compressed Sensing to Detecting Damage in Structures

Authors: David Mascarenas
Donald Hush
James Theiler
Charles Farrar

INQUIRIES TO: David Mascareñas
Los Alamos National Laboratory Engineering Institute
PO Box 1663 MS T001
Los Alamos, NM, 87544
505-663-5227
dmascare@gmail.com

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

ABSTRACT

Detecting damage in a large scale structure such as a bridge or dam requires the collection of data at multiple time and length scales. The collection of these data generally requires the structure under observation to be instrumented with a variety of sensors, each with a unique sampling rate. One of the principal challenges to the structural health monitoring (SHM) community is to take this large, heterogeneous set of data, and extract information that allows the estimation of the remaining service life of a structure. Another important challenge is to collect relevant data from a structure in a manner that is cost effective, and respects the size, weight, cost, energy consumption, and bandwidth limitations placed on the system. Both of these challenges have proven to be formidable hurdles to the wide-scale implementation of SHM systems. In this work we explore the suitability of compressed sensing to address both challenges.

Recently compressed sensing has presented itself as a candidate solution for directly collecting relevant information from sparse, high-dimensional measurements. The main idea behind compressed sensing is that by directly collecting a relatively small number of coefficients it is possible to reconstruct the original measurement. The coefficients are obtained from linear combinations of (what would have been the original direct) measurements. At first glance it would appear that this should not be possible because it would require solving an underdetermined linear system of equations. However, it has been shown that if the solution is sparse in some basis, it is possible to find the solution using ℓ_1 norm regularization. Conveniently, most signals found in nature are indeed approximately sparse (in some basis) with the notable exception of random noise. Therefore, the findings of the compressed sensing community hold great potential for changing the way SHM data is collected.

In this work a digital version of a compressed sensor is implemented on-board a microcontroller similar to those used in embedded SHM sensor nodes. The sensor node is tested in a surrogate SHM application requiring acceleration measurements. Currently the prototype compressed sensor is capable of collecting compressed coefficients from measurements and sending them to an off-board processor for reconstruction using L1 norm minimization. A compressed version of the matched filter known as the smashed filter, has also been implemented on-board the sensor node, and its suitability for detecting structural damage will be discussed.

INTRODUCTION

Data for structural health monitoring applications is generally collected using a distributed sensor network. Distributed sensor networks made up of nodes with hard-wired data and communication lines are generally have high installation costs. The goal is to transition to low-power, wireless sensor networks featuring minimal installation costs [1]. Two of the major problems with these types of sensor networks are the conservation of energy and bandwidth. Compressed sensing techniques hold promise to help address both of these problems. By collecting compressed coefficients, the signal of interest can be represented using a fraction of the

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

measurements required by traditional Nyquist sampling. The result is reduced energy consumption for data collection, storage and transmission. In addition, the bandwidth required to transmit the data from a signal is also significantly reduced. The focus of this work is to evaluate the applicability of compressed sensing techniques to expand the capabilities of wireless sensor networks for structural health monitoring applications.

BACKGROUND OF COMPRESSED SENSING

Compressed sensing has been a prolific research topic over the last few years. Excellent tutorials covering the basics of compressed sensing can be found in [2] and [3]. To summarize, a signal of interest x can be represented as:

$$x = \sum_{i=1}^{N} s_i \psi_i \quad \text{or in matrix form as } X = \Psi s \tag{1}$$

Where Ψ is an orthonormal basis and "s" is the representation of the signal in the Ψ domain. In the case of compressed sensing we are interested in the case where x is compressible in some domain. That is, the number of significant non-zero elements of s is equal to k and $k << N$. A measurement matrix "Φ" is then introduced to produce compressed sensing coefficients y.

$$y = \Phi x = \Phi \Psi s = \Theta s \tag{2}$$

Where Φ has $M<N$ rows. At this point it is important to note that this equation represents an underdetermined system of linear equations. One of the major breakthroughs of the compressed sensing community was the finding that assuming $k<<M$ it is possible to recover x from y assuming the matrix Φ possesses certain properties. The direct formulation of this problem is finding the vector s with minimal ℓ_0 norm. Unfortunately ℓ_0 norm minimization is numerically unstable and NP-complete. It has been shown though that the associated ℓ_1 norm regularization problem [5] can be solved to recover compressible signals from the compressed coefficients y. In this work, the ℓ_1 norm regularization approach will be explored for recovering the signal x from compressed coefficients y.

EXPERIMENTAL SETUP

In order to evaluate the applicability of compressed sensing for embedded structural health monitoring sensor nodes, a digital prototype of a compressed sensor node was built. The prototype consisted of an ATmega1281 microcontroller, an ICP accelerometer, and the associated amplification and ICP circuitry required to interface the analog-to-digital converter (ADC) of the microcontroller to the ICP accelerometer. The accelerometer was then attached to the second floor of a representative 3 story structure as shown in Figure 1. The accelerometer was oriented to measure the transverse vibration of the 3-story structure. An electro-magnetic shaker was then attached to the base of the three story structure to provide a source of excitation. The

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

excitation to the structure was a sine wave with a frequency of 30.7 Hz which corresponds to the first resonant frequency of this structure. In order to introduce damage into the structure, a rubber bumper was used to induce a nonlinear response when the relative transverse displacement between the second floor and the base would exceed a threshold value. The signal from the accelerometer was sampled by the ATmega1281 with a 10 bit ADC at a sampling rate of about 3000 Hz. In this work, 256 point time series were collected by the Atmega1281 and subsequently converted into compressed measurements by the microcontroller. The elements of the measurement matrix Φ were chosen to be either ± 1. The generation of the measurement matrix Φ was accomplished using a linear feedback shift register (LFSR) similar to that mentioned in [4]. The ± 1 measurement matrix was selected in order to allow the generation of the compressed coefficients on-board the microcontroller using integer arithmetic. The ATmega1281 was placed onboard an STK500 evaluation board. A base station laptop was then connected to the STK500 in order to facilitate the debugging of the compressed sensing algorithms, and to expedite the collection of compressed coefficients from the embedded sensor node. For the purpose of this experiment, the ATmega1281 would transmit both the compressed measurements as well as the original signal when it was queried for a measurement. By collecting both pieces of data the reconstructed signal derived from the compressed coefficients could be compared to the original signal in order to evaluate the performance of compressed sensing techniques.

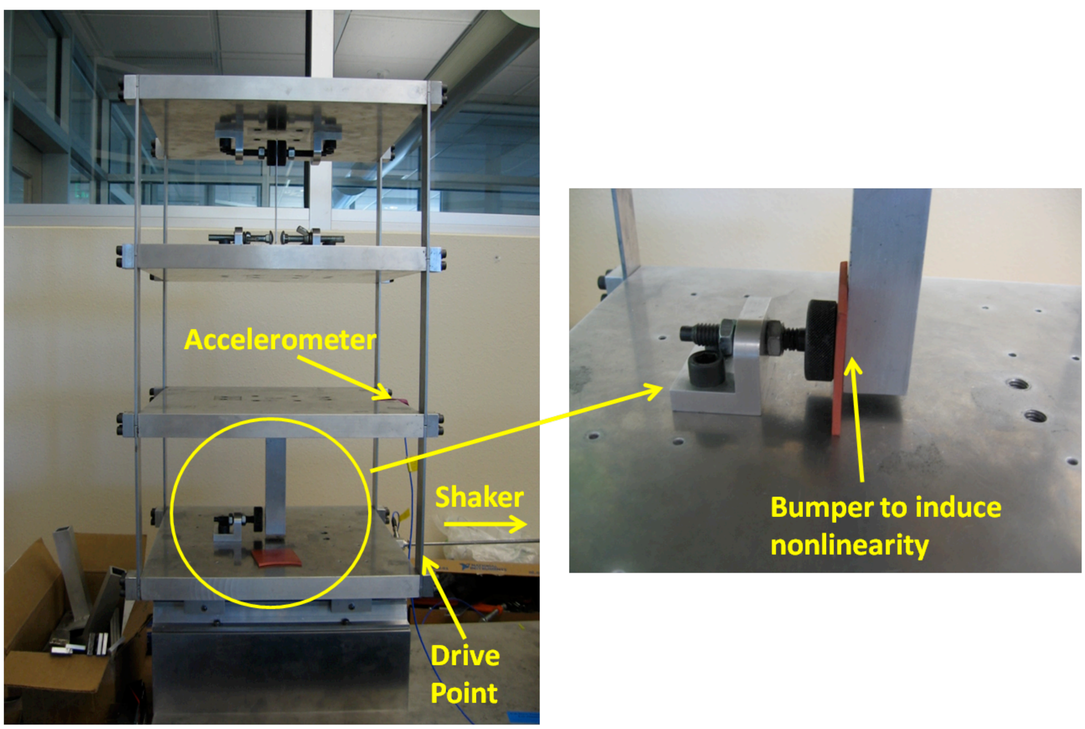

Figure 1- Representative 3 story structure used to evaluate the compressed sensor node.

RESULTS FROM ℓ_1 NORM REGULARIZATION EXPERIMENTS

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

The Structure described in the previous section was placed in a configuration so that it could assume either a damaged state or an undamaged state. In the undamaged state, the output from the accelerometer should assume a sine wave with a frequency of 30.7 Hz that corresponds to the shaker excitation frequency. In the damaged state, the rubber bumper interacts with the structure and the frequency content of the resulting output from the structure is more widely distributed across the spectrum. For this work we assume that the resulting signals from the structure should be sparse in the Fourier basis. Data was collected from the structure in both the damaged and undamaged states, and was subjected to the compressed sensing measurement process. The resulting compressed measurements were then cast into the ℓ_1 norm regularization framework to attempt recovery of the original signal. The ℓ_1 norm regularization problem can be written as:

$$\underset{s}{Minimize}\left(\|y - \Phi\Psi s\|_2 + \gamma |s|_1\right) \tag{3}$$

The ℓ_1 norm regularization problem trades off between the size of the residual and the sparsity of s. The ℓ_1 norm regularization was implemented using the CVXMOD software [6]. 256 pt time series data was collected from the structure using the ATmega1281 and was transmitted to the base station laptop. The compressed coefficients were calculated using the measurement matrix generated by the LFSR. The resulting compressed coefficients were subjected to the ℓ_1 norm regularization with $\gamma = 0.01$. The value of γ was selected heuristically. The ℓ_1 norm regularization problem was solved using 64, 128, 160, and 200 compressed coefficients for both the damaged and undamaged cases. These correspond to compression factors of 25%, 50%, 63%, and 78% respectively. The resulting reconstructions can be found in Figure 2 and Figure 3. The values of the Fourier basis coefficients for the undamaged and damaged cases are found in Figure 4 and Figure 5 respectively. It should be noted that the sine wave measured in the undamaged case has a slight non-linearity caused by clipping in the signal conditioning electronics at low-voltages. This situation was unavoidable in order to excite with an appropriate amplitude to induce nonlinearities in the damage case.

From the reconstructed signals in Figure 2 and Figure 3 we can see that the nature of the output signal does not begin to become apparent until about 160 compressed measurements are taken. The reconstructed signals generated from a lower number of compressed coefficients also tend to exhibit significant high-frequency components that do not show up in the original signal. As the number of

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

compressed coefficients increases the magnitude of the non-existent high-frequency components tends to decrease and the accuracy of the reconstruction improves.

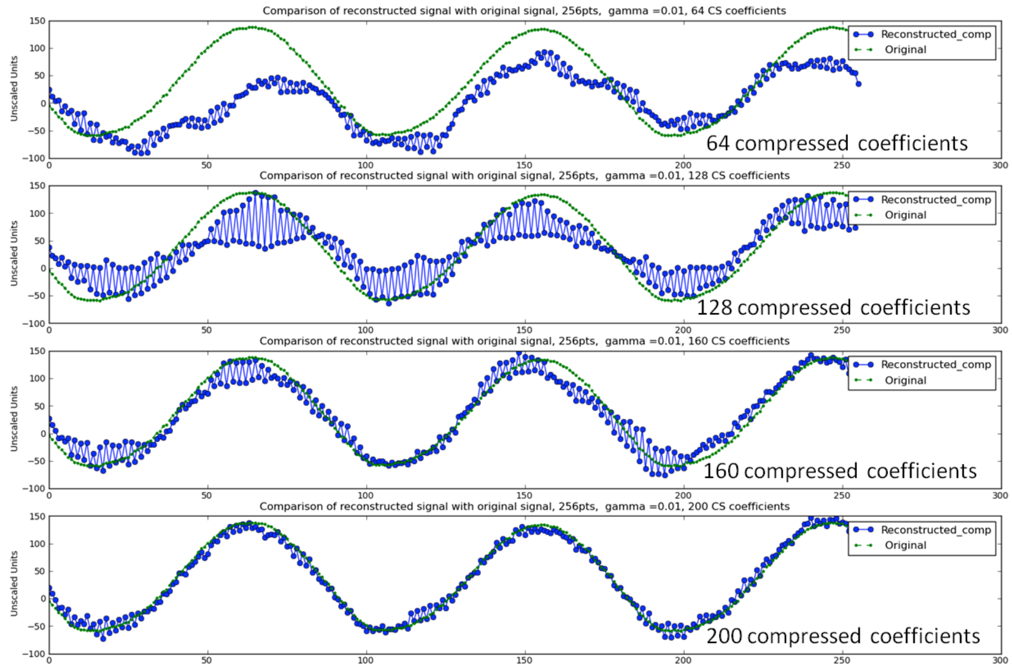

Figure 2 - Nondamaged case compressed sensing reconstruction

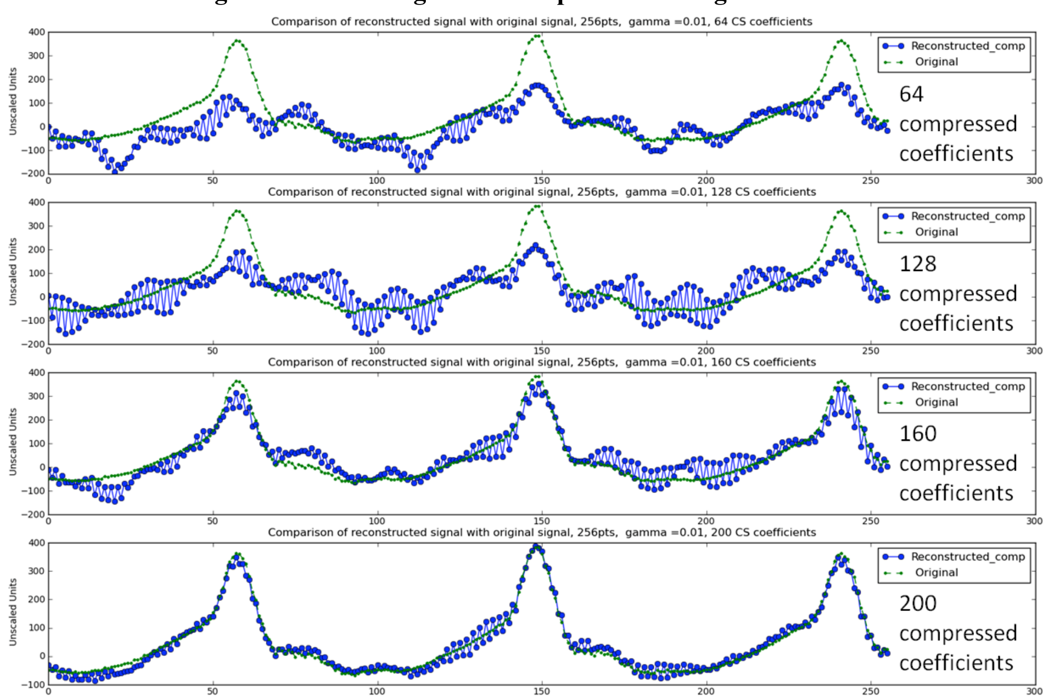

Figure 3 - Damaged case compressed sensing reconstruction

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

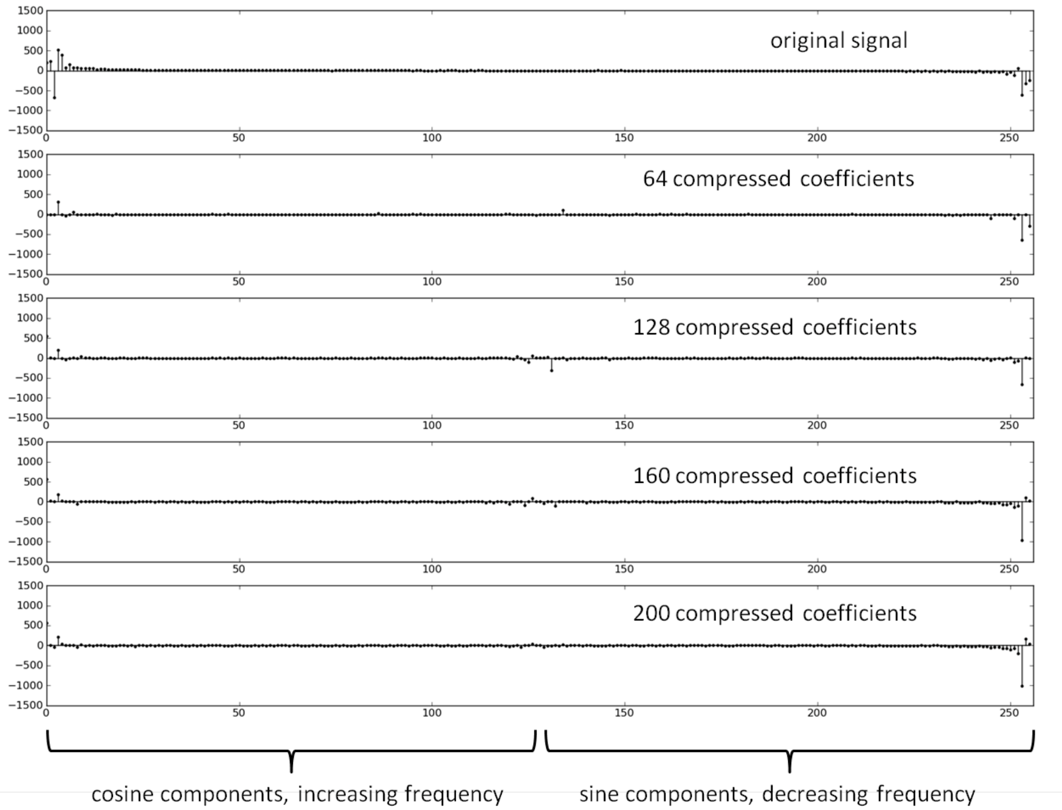

Figure 4 - Fourier coefficients from undamaged structure, original and reconstructed signals.

Next consider the plots of the Fourier basis coefficients displayed in Figure 4 and Figure 5. In both the damaged and undamaged cases the ℓ_1 norm regularization problem finds at least either the sine or cosine component of the original signal. From these plots it is easy to see the inclusion of false, high-frequency components as evidenced by the non-zero components present in the middle of the Fourier basis plots. As the number of compressed sensing coefficients included in the plots is increased, the magnitude of the false high-frequency components decreases, and the accuracy of the reconstructed signal also improves. At this time it is not entirely clear why the reconstructed signals tend to exhibit significant contributions from high frequency components. One possibility is that the reconstructions are amplifying noise present in the signal. Future work is going to consider methods for improving the accuracy of the reconstructed signals.

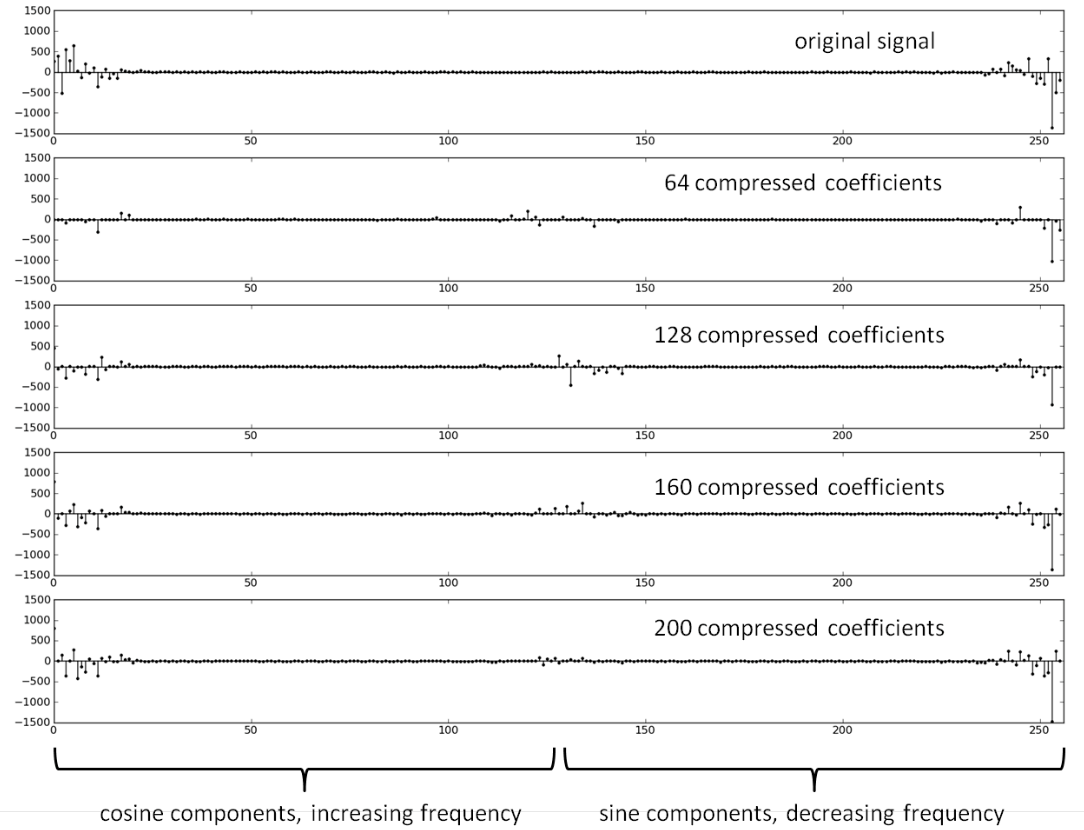

Figure 5 - Fourier coefficients from damaged structure, original and reconstructed signals.

SMASHED FILTER

The initial ℓ_1 norm regularization results showed that a fairly significant number of compressed coefficients were needed to begin to accurately reconstruct the signal of interest. In structural health monitoring, the main concern is generally the detection of damage and not necessarily the collection of accurate time series. With this in mind it was decided to investigate alternate techniques to try and detect the presence of damage while using a relatively small number of compressed coefficients. An extension of the matched filter to the compressed domain known as the "smashed filter" seemed an appropriate technique to evaluate. [7]. To summarize, the smashed filter is implemented in basically the same manner as the conventional matched filter. The main difference is that the smashed filters are generated by taking phase shifted versions of the signals of interest, and then subjecting them to the measurement process Φ. The matched filters h_m are generated from the signals of interest h as:

$$h_m = \Phi h \tag{4}$$

For this work, training signals from both the damaged and undamaged states were collected. The damage case training signal was collected in such a manner that it featured a small mismatch with the operational damage case signals used to evaluate the performance of the smashed filter. By allowing a small mismatch between the training damage case, and the operational damage case we can get an initial sense of the robustness of the smashed filter to small perturbations. This mismatch was

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

achieved by perturbing the location of the rubber bumper. Once the smashed filters were generated, 100 experiments resulting in 128 compressed coefficients per experiment were collected from the structure in both the damaged and undamaged cases. Subsets of the 128 compressed coefficients were then used to evaluate the performance of the smashed filter for various numbers of compressed coefficients. To implement the smashed filter, the inner product of the compressed coefficients and the smashed filter vectors was calculated for each experiment. The smashed filter with the largest inner product was then selected and the experiment was classified as damaged or undamaged based on whether or not the corresponding smashed filter came from the damaged or undamaged case. Table 1 and Table 2 illustrate the results of applying the smashed filter to the experiments for the undamaged and damaged cases respectively. From this data we see that once at least 32 compressed coefficients are used to calculate smashed filters, the probability of misclassifications become very low. It is noteworthy that the number of compressed coefficients needed to achieve accurate classification with the smashed filter is only 1/8 the number of data points in the original time series measurement. Based on these results, the smashed filter has the potential to significantly reduce the number of measurements needed to classify whether or not a structure is damaged.

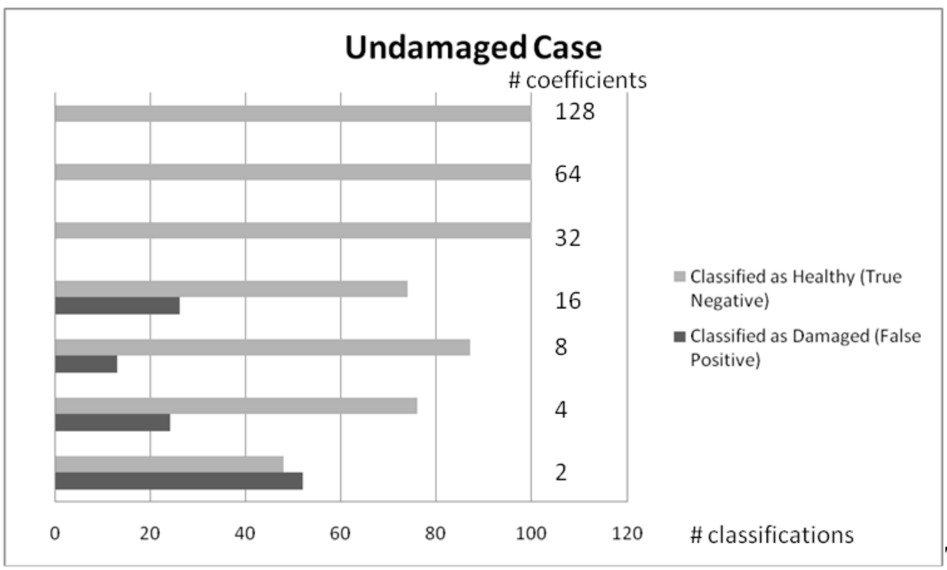

Table 1.

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544

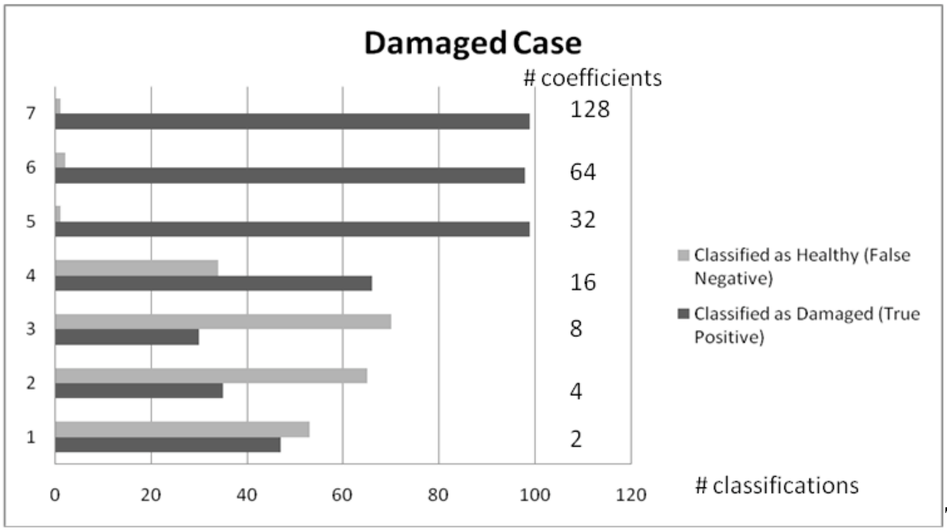

Table 2.

CONCLUSIONS

In this work the performance of compressed sensing techniques for embedded structural health monitoring sensor nodes have been investigated. A preliminary look at the ability of ℓ_1 norm regularization to reconstruct time series measurements of acceleration collected from structures has been presented. The suitability of the smashed filter for damage classification in structures has also been presented. It was found that the smashed filter has potential to significantly reduce the number of measurements required to classify the state of health of a structure. Although not discussed in detail, results obtained in this work suggest that compressed sensing has the potential to conserve energy and bandwidth in embedded wireless sensor nodes for structural health monitoring applications. Future work will focus on better quantifying these savings.

REFERENCES

1. Mascarenas, D., Flynn, E., Farrar, C., Park, G., Todd, M., "A Mobile Host Approach for Wireless Powering and Interrogation of Structural Health Monitoring Sensor Networks," Sensors Journal, IEEE, Vol 9, Issue 12, pp 1719-1726, Dec 2009.
2. Emmanuel Candès, Compressive sampling. (Int. Congress of Mathematics, 3, pp. 1433-1452, Madrid, Spain, 2006)
3. Richard Baraniuk, Compressive sensing. (IEEE Signal Processing Magazine, 24(4), pp. 118-121, July 2007)
4. Jason Laska, Sami Kirolos, Yehia Massoud, Richard Baraniuk, Anna Gilbert, Mark Iwen, and Martin Strauss, Random sampling for analog-to-information conversion of wideband signals. (IEEE Dallas Circuits and Systems Workshop (DCAS), Dallas, Texas, 2006)
5. Candès, E. J., Romberg, J., Tao, T., Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory* **52** (2006), 489–509.
6. Mattingley, J., Boyd, S., "CVXMOD-Convex optimization software in python," http://cvxmod.net/, August 2008.
7. Mark Davenport, Marco Duarte, Michael Wakin, Jason Laska, Dharmpal Takhar, Kevin Kelly, and Richard Baraniuk, The smashed filter for compressive classification and target recognition. (Computational Imaging V at SPIE Electronic Imaging, San Jose, California, January 2007)

Mascarenas, D., Hush, D., Theiler, J., Farrar, C., Los Alamos National Laboratory, PO Box 1663 MS T001, Los Alamos, NM 87544