

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 12-02-2016		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 5-Nov-2013 - 4-Nov-2015	
4. TITLE AND SUBTITLE Final Report: Rapid anomaly detection and tracking via compressive time-spectra measurement			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER W911NF-14-C-0006		
			5c. PROGRAM ELEMENT NUMBER 665502		
6. AUTHORS Lenore McMackin, Matthew A. Herman, Kevin Kelly			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES InView Technology Corporation 8900 Shoal Creek Blvd, Suite 125 Austin, TX 78757 -6840			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 64505-CS-ST2.2		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Report developed under Topic #A12a-T007, contract W911NF-12-C-0106. This report describes the results of a two-year Phase II STTR program undertaken by InView in partnership with Rice University to develop a compressive sensing based imager with high speed anomaly detection capability. During this period, novel multi-channel compressive sensing measurement patterns and detector geometries were developed along with statistical methods for performing anomaly detection on data in the compressed domain. Detection performance was analyzed using receiver operating characteristic and precision-recall curves. Novel					
15. SUBJECT TERMS Compressive sensing, anomaly detection, Hadamard domain processing, machine vision, single-pixel camera					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Lenore McMackin
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 512-656-9153

Report Title

Final Report: Rapid anomaly detection and tracking via compressive time-spectra measurement

ABSTRACT

Report developed under Topic #A12a-T007, contract W911NF-12-C-0106.

This report describes the results of a two-year Phase II STTR program undertaken by InView in partnership with Rice University to develop a compressive sensing based imager with high speed anomaly detection capability. During this period, novel multi-channel compressive sensing measurement patterns and detector geometries were developed along with statistical methods for performing anomaly detection on data in the compressed domain. Detection performance was analyzed using receiver operating characteristic and precision recall curves. New concepts, algorithms and techniques were simulated and tested experimentally using specially modified compressive sensing cameras under laboratory conditions. It was found that the number of measurements needed for anomaly detection was far less than the number of measurements needed for compressive imaging, validating high speed system operation.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

02/12/2016 1.00 Matthew A. Herman, Tyler Weston, Lenore McMackin, Yun Li, Jianbo Chen, Kevin F. Kelly. Recent results in single-pixel compressive imaging using selective measurement strategies, SPIE Sensing Technology + Applications. 20-APR-15, Baltimore, Maryland, United States. : ,

TOTAL: 1

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

<u>NAME</u>	<u>PERCENT_SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT_SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Names of Under Graduate students supported

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

1 a. William Marsh Rice University

1 b. 6100 Main Street

Houston TX 770051827

Sub Contractor Numbers (c):

Patent Clause Number (d-1):

Patent Date (d-2):

Work Description (e): Develop compressed domain anomaly detection algorithms and experiments

Sub Contract Award Date (f-1): 11/5/13 12:00AM

Sub Contract Est Completion Date(f-2): 11/4/15 12:00AM

1 a. William Marsh Rice University

1 b. PO Box 1892, MS 16

Houston TX 772511892

Sub Contractor Numbers (c):

Patent Clause Number (d-1):

Patent Date (d-2):

Work Description (e): Develop compressed domain anomaly detection algorithms and experiments

Sub Contract Award Date (f-1): 11/5/13 12:00AM

Sub Contract Est Completion Date(f-2): 11/4/15 12:00AM

Inventions (DD882)

Scientific Progress

Technology Transfer

InView Technology Corporation plans to commercialize an anomaly-detection-capable camera built using the multi-pixel architecture developed during this project.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 02-12-2016		2. REPORT TYPE Final Technical		3. DATES COVERED (From - To) 05 Nov 2013- 04 Nov 2015	
4. TITLE AND SUBTITLE Rapid anomaly detection and tracking via compressive time-spectra measurement				5a. CONTRACT NUMBER W911NF-14-C-0006	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) McMackin, Lenore Herman, Matthew A. Weston, Tyler Kelly, Kevin				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) InView Technology Corporation 6201 East Oltorf Street Suite 400 Austin, TX 78741				8. PERFORMING ORGANIZATION REPORT NUMBER 1	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) ARMY RESEARCH OFFICE ARO-RTP 4300 S. MIAMI BLVD. DURHAM NC 27703				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; Distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Report developed under Topic #A12a-T007, contract W911NF-12-C-0106. This report describes the results of a two-year Phase II STTR program undertaken by InView in partnership with Rice University to develop a compressive sensing based imager with high speed anomaly detection capability. During this period, novel multi-channel compressive sensing measurement patterns and detector geometries were developed along with statistical methods for performing anomaly detection on data in the compressed domain. Detection performance was analyzed using receiver operating characteristic and precision recall curves. New concepts, algorithms and techniques were simulated and tested experimentally using specially modified compressive sensing cameras under laboratory conditions. It was found that the number of measurements needed for anomaly detection was far less than the number of measurements needed for compressive imaging, validating high speed system operation.					
15. SUBJECT TERMS STTR Report, compressive sensing, anomaly detection, Hadamard imaging, sparsity, dimensional reduction, statistical pattern recognition, receiver operating characteristic curves, precision- recall					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)
U	U	U	UU		

Reset

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18



Contract no. W911NF-14-C-0006

Proposal no. A2-5466

FINAL REPORT

Item # 0002AM

Title: Rapid anomaly detection and tracking via compressive time-spectra measurement

Contract performance period: 05 Nov 2013 - 04 Nov 2015

**Submitted:
12 Feb 2016**

**InView Technology Corporation
6201 E. Oltorf St., Ste. 400
Austin, TX 78741**

PROJECT SUMMARY

InView, in partnership with Rice University, completed a two-year Phase II plan whose objectives defined the development of a compressive sensing-based high-speed detection and tracking sensor. The Phase II work was built on the unique “single-pixel camera” architecture and algorithms invented at Rice University and developed commercially by InView Technology Corporation. In this architecture the compressive sensing sampling strategy is implemented optically using a high-resolution spatial light modulator that encodes the scene to be imaged with a set of patterns mathematically constructed to exploit redundancy in the structure of the scene to reduce the number of measurements required for faithful reconstruction. While randomized patterns are effective in general imaging tasks, we show that patterns developed based on selected components of the Hadamard spectrum—that is, the Partial Complete strategy-- can be used for anomaly detection and have the advantage of requiring many fewer measurements than required for imaging, enabling high-speed operation. This final report summarizes our approach and achievement of five main program objectives:

- 1: Development of lab metrics and accurate targets for change detection system testing
- 2: Optimizing Real Time Anomaly Detection
- 3: Development and implementation of Multiplexed Spatial-Spectral Hybrid Subspace sampling patterns for high-speed scenario specific anomaly detection
- 4: Development of real-time change detection for moving object tracking with a static background
- 5: Assemble and test CS hardware prototype for high-speed detection

Along the way we developed new concepts in Hadamard imaging and compressive sensing including the development of an improved multi-pixel compressive camera that operates on a multiplicity of spatial regions of the scene in parallel to speed imaging frame rate and also provide localization of anomalies detected in within the field of view. We quantified the imaging and anomaly detection algorithms and architectures in simulation that utilized statistical methods for analyzing data in the compressed domain. We then experimentally demonstrated operation on bench top hardware prototypes built at Rice and at InView based on both single-pixel and multi-pixel compressive camera designs.

Compressed sensing approaches for rapid detection and tracking has direct applications in lower-cost, higher-performance sensors particularly in the shortwave infrared where focal plane array solutions are expensive and low resolution. The systems developed here have direct application in autonomous vision systems for military, law enforcement, and border security uses. Commercial applications of compressed domain anomaly detection and tracking abound in machine vision, inspection and process control and image-based internet searching.

FOREWORD

InView Technology Corporation and Rice University have had a long and fruitful collaboration on developing applications of the mathematical theory of compressive sensing and bringing them to market. In 2013 InView released the world's first compressive-sensing based camera built on the "single-pixel camera" architecture first conceived at Rice. Since then we have collaborated on the development of the multi-pixel video camera architecture, multi-spectral systems, and advanced compressed domain image processing techniques. In particular, InView would like to acknowledge the work of Rice professor Kevin Kelly and his growing research group for their many contributions to this project, as well as InView scientists Dr. Matthew A. Herman and Tyler Weston for their tireless efforts.

Lenore McMackin, PhD,
President and CTO InView Technology Corporation

TABLE OF CONTENTS

Project Summary	3
Foreword	4
Table of Contents	5
Executive summary	7
 DESCRIPTION OF TECHNICAL RESULTS	 9
 Section 1. Introduction.....	 9
 Section 2. Compressed Domain Anomaly detection Algorithms, Measurement Strategies and Simulations	 11
2.1 Sum-to-one and Walsh Hadamard Matrices	11
2.2 Sylvester Type Hadamard Matrices	12
2.3 Selective Measurement Strategies Enabled by the Block Structure of Kronecker product matrices	13
2.4 Partial Complete.....	17
2.5 PC Block selection strategy	18
2.6 Local global specifically designed for multi-detector architectures	21
2.7 Mean-to-deviation versus SNR analysis of compressive measurements.....	23
2.8 Z-Score and MMD	24
2.9 Anomaly detection statistics testing simulations with Partial-Complete Block Sets	25
2.10 ROC and PR curve development and calculation of thresholds	28
2.11 Background on ROC curves	28
2.12 Development of Precision Recall curves: Better than ROC	29
 Section 3. Implementing adaptive compressed sensing methods on a focal plane array	 30
3.1 Anomaly Detection Table Top Demonstrator and Laser Anomaly Design.....	31
3.2 Hardware Synchronization.....	31
3.3 Laser Trigger Anomaly Control Experiment Results	33
3.4 “Virtual channel” geometry for imaging and anomaly detection	34
3.5 PSF Development for Multi-pixel Calibration	39
3.5.1 Raster-scan for Point Spread Function Measurement.....	39
3.6 Focal Plane Array (FPA) anomaly detection	42
Appendix 3-1 PSF-aware reconstruction algorithms for imaging:	49
 Section 4. Anomaly Detection Demonstrations and Results.....	 51
4.1 Simulation scenarios	51
4.2 Statistical inference	52

4.3 Simulation with different detection matrix	53
4.4 ROC and PR curve.....	57
4.5 Anomaly detection experiment with SPC.....	58
4.5.1 Laser triggered anomaly control system	59
4.5.2 Experiment results and discussion.....	61
4.6 Anomaly detection with FPA.....	64
4.6.1 Virtual channel.....	65
4.6.2 Point spread function correction to the imaging system.....	67
4.6.3 Anomaly detection results with FPA and discussion.....	67
Appendix 4-1: Improving the PR curve with a ground truth signal broadening technique	71
Section 5. Conclusion and future directions	73

EXECUTIVE SUMMARY

InView, in partnership with Rice University, completed a two-year Phase II plan whose objectives defined the development of a compressive sensing-based high-speed detection and tracking sensor. The Phase II work was built on the unique “single-pixel camera” architecture and algorithms invented at Rice University and developed commercially by InView Technology Corporation. In this architecture the compressive sensing sampling strategy is implemented optically using a high-resolution spatial light modulator that encodes the scene to be imaged with a set of patterns mathematically constructed to exploit redundancy in the structure of the scene to reduce the number of measurements required for faithful reconstruction. While randomized patterns are effective in general imaging tasks, we show that patterns developed based on selected components of the Hadamard spectrum can be used for anomaly detection and have the advantage of requiring many fewer measurements than required for imaging.

Compressed sensing approaches for rapid detection and tracking has direct applications in lower-cost, higher-performance sensors particularly in the shortwave infrared where focal plane array solutions are expensive and low resolution. The systems developed here have direct application in autonomous vision systems for military, law enforcement, and border security uses. Commercial applications of compressed domain anomaly detection and tracking abound in machine vision, inspection and process control and image-based internet searching.

The final report includes the following results in detail

- Developed new algorithms and measurement strategies to perform and optimize compressed domain anomaly detection. Throughout the project we compared STOne (Sum-to-One), and Walsh-Hadamard techniques in the Partial Complete data acquisition strategy.
 - Implemented the statistical methods of z-score and MMD algorithms as anomaly detection methods in the compressed domain
 - Developed Partial Complete measurement strategy used previously for imaging, into a strategy for anomaly detection that maximizes detectable anomaly energy in the first few measurements so that as few measurements as possible need to be used for detection
 - Demonstrated that Partial Complete block subsets can be chosen that simultaneously cover the entire field of view without nulls and provide a strong signal for imaging and anomaly detection. In particular, the canonical block set derived from 4 components of the Hadamard spectrum has shown optimal operation over a range of feature and anomaly sizes.
 - Demonstrated that only a fraction of the number of patterns within each selected block need to be used for high probability of detection
 - Developed ROC and PR curves as performance metrics

- Developed line broadening technique to improve detection capability
- Used PR curves derived from simulations and experimental measurements to determine operating thresholds and operating limits for an anomaly detection camera as a function of noise and anomaly brightness and duration
- Developed a re-programmable “virtual detector” method to implement an adaptable multi-pixel compressive sensing measurement strategy on an available 64 x 64 InGaAs pixel array
 - Demonstrated that multi-pixel geometry operates at high speed compared to the single pixel CS geometry not only because multiple pixels are working in parallel, but also because only a few blocks of the Partial-Complete spectrum need to be measured
 - Implemented point spread function (PSF) method to improve the aggregation of measurements across multiple detectors in either a Focal Plane Array geometry or virtual array geometry to improve multi-pixel image quality and provide a more realistic anomaly detection simulation environment
- Built demonstration systems at Rice and InView on which compressive multi-pixel imaging and anomaly detection has been implemented using the virtual array geometry, partial complete measurement strategy, MMD and Z-score anomaly detectors in the compressed data domain, and PSF for image reconstruction.
 - Demonstrated the detection of laser anomalies using all of these techniques on detecting laser anomalies in real video data in the presence of moving objects and background clutter
 - Completed experiments demonstrating operational range of the camera
- Future work
 - Use of joint information across sensors

DESCRIPTION OF TECHNICAL RESULTS

Section 1. Introduction

The results of this Phase 2 STTR presented here represent the follow on to Phase 1 program A12a-T007 in which a new class of compressive sensing matrices were developed that can reconstruct images at multiple scales while also possessing corresponding fast transforms enabling high-speed low resolution image previews and full-scale image reconstructions. These matrices were dubbed Sum-to-One, or STOne pattern.¹ New modes for change detection were explored for operation on CS-matrix encoded image data by predicting regional statistical outliers in the data by applying a combination of z-score test and the *Maximum Mean Discrepancy* (MMD) statistic in the compressive domain. The product of the Z-score and the MMD displayed the sensitivity of MMD while encoding the same qualitative information as the Z-statistic. It was the goal of this Phase 2 to compare and optimize real-time anomaly detection methods developed in Phase I, implement them in simulation and experiment, and construct a prototype compressive sensing anomaly detection camera system.

Early in the Phase 2 we found that the Partial-Complete sensing strategy^{2,3} developed at InView had several advantages in optimizing compressive domain feature encoding measurement patterns. Like STOne patterns, they could be recursively generated from a Kronecker product, which also endows them with a fast transform. In addition, the “sequency” structure partitions the Hadamard domain into sequency-based blocks of patterns related by scale size and spatial features that can be chosen to facilitate anomaly detection with a minimal number of measurements. The point of our simulation effort was to determine the minimum number of blocks and the minimum number of patterns from within each block necessary to achieve a certain level of anomaly detection precision and recall. Higher operating speeds are associated with fewer patterns. It was found that 4 blocks are needed at a minimum that are carefully selected to observe the full field of view without nulls. It was also found that a surprisingly low number of patterns from within each block can produce effective anomaly detection signals even in the presence of noise, facilitating the detection of short duration events.

Using the partial-complete strategy also facilitated implementation of compressive sensing and anomaly detection on a multi-pixel compressive sensing camera architecture. A multi-pixel architecture effectively parallelizes the single-pixel camera measurement path into multiple, simultaneously operating cameras each of which examines only a portion of the field of view. The Phase 2 began under the assumption that the multi-pixel camera already developed at InView⁴ would be the platform for experimental verification of the anomaly detection and tracking procedures. However, the large amount of optical crosstalk between the channels

¹ Tom Goldstein, et al., “The STONE Transform: Multi-Resolution Image Enhancement and Real-Time Compressive Video,” arXiv:1311.3405v2.

² Matthew A. Herman, “Compressive Sensing with Partial-Complete, Multiscale Hadamard Waveforms,” in Imaging and Applied Optics, OSA Technical Digest (online) (Optical Society of America, 2013), paper CM4C.3.

³ Matthew A. Herman, et al., “Recent results in single-pixel compressive imaging using selective measurement strategies,” Proc. SPIE 9484, 948409 (2015).

⁴ Matthew A. Herman, et al., “A higher-speed compressive sensing camera through multi-diode design,” Proc. SPIE 8717, 871706 (2013)

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

inherent in that system produced specialized requirements for the modulation patterns that were not compatible with the partial complete strategy. In response, InView and Rice designed and built a new multi-pixel architecture with greatly reduced crosstalk, a more direct optical path to the sensor and implemented the architecture using a commercially available 64 x64 InGaAs sensor array. This architecture had the added benefit of having 4096 sensors compared to just 32 sensors on the original system, and the ability to reconfigure these sensors digitally into “virtual” detectors with reconfigurable geometries. The following sections describe our results in more detail.

Section 2. Compressed Domain Anomaly detection Algorithms, Measurement Strategies and Simulations

In compressive sensing (CS) of images, measurements of a scene x are observed via a measurement or sensing matrix A as $y = Ax$. In this discretized model, x is assumed to have N pixels arranged in a column vector, and A is an underdetermined $M \times N$ matrix, where typically $M \ll N$. Images are computationally reconstructed using information in the length- M vector y along with knowledge of the matrix A used to make the measurements.

The oft used graphical representation of the matrix equation describing measurement vector y , is shown below:




Fig. 2-1. Matrix representation of $y = Ax$.

Ideally, the rows of A are instances of random noise (e.g., drawn from a Gaussian or Bernoulli process). However, this requires significant memory storage when $M \cdot N$ is large. Moreover, the computational complexity of the reconstruction algorithm cannot be reduced when A is purely random. Instead it is advantageous to use structured and deterministic waveforms, which can then be pseudo-randomized to achieve the property of “incoherence” desired in CS applications. These matrix types are often much more computationally cost effective and have additional advantages. We now introduce the matrix structures that were investigated for use in compressive anomaly detection in this project.

2.1 Sum-to-one and Walsh Hadamard Matrices

A new class of compressive sensing matrices has been developed that are generated in an efficient implementation based on the Sum-to-One (STOne) transform and have been shown to reconstruct images at multiple scales. The STOne matrix, S_n , of order $n = 4^k$ is recursively constructed via the Kronecker product⁵

$$S_{4^{k+1}} := S_4 \otimes S_{4^k} \quad (2-1)$$

for $k = 0, 1, 2, \dots$, where $S_1 = [1]$ and

⁵ T. Goldstein, et al., “The STONE transform: multi-resolution image enhancement and real-time compression video,” arXiv:1311.3405

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

$$S_4 := \frac{1}{\sqrt{4}} \begin{bmatrix} -1 & +1 & +1 & +1 \\ +1 & -1 & +1 & +1 \\ +1 & +1 & -1 & +1 \\ +1 & +1 & +1 & -1 \end{bmatrix}. \quad (2-2)$$

STOne matrices are orthogonal, and their rows and columns all add up to +1. When data is acquired from a full-resolution $N \times N$ image using the STO transform, a block of n^2 measurements can be used to create a low-resolution “preview,” in the form of an $n \times n$ image, where $n < N$, that approximates the full-scale $N \times N$ image. The low resolution preview is reconstructed using a simple fast transform, and is very inexpensive computationally.

Throughout this project, the results of anomaly detection procedures using both STOne and Hadamard matrices were compared. We found the behavior of the change detection algorithm to identify anomalies is essentially identical between STOne measurement patterns and scrambled or permuted versions of well known Walsh-Hadamard (WH) patterns. Subsequent investigation showed that STOne matrices are equivalent to Sylvester Hadamard matrices due, in part, to its recursive Kronecker product construction. Based on this mathematical similarity and our experimental results, we focused our anomaly detection studies on patterns from permuted WH matrices and a special ordering of quasi-Sylvester Hadamard matrices.

2.2 Sylvester Type Hadamard Matrices

A general Hadamard matrix H_N of order N is defined as having $\{\pm 1\}$ entries, with rows and columns that are orthogonal. The Kronecker product of two Hadamard matrices H_F and H_B is also a Hadamard matrix: $H_{F \cdot B} := H_F \otimes H_B$. We can characterize a given row (or column) of H_N in terms of its “sequency,” which simply counts the number of transitions from +1 to -1, or vice-versa. This is similar to the familiar notion of “frequency” that we associate with sinusoids.

The most common Hadamard matrices used in practice have orders that are powers of two. The Sylvester-type Hadamard matrix of order $N = 2^K$, for some integer K , is recursively defined for all $1 \leq m \leq K - 1$, according to

$$H_{2^K} := H_{2^m} \otimes H_{2^{K-m}} \quad (2-3)$$

Where

$$H_2 = \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix} \quad (2-4)$$

is the fundamental building block.

It is the K -fold application of the Kronecker product that endows the Sylvester-type matrix with a fast $O(N \log N)$ implementation, similar to the fast Fourier transform (FFT). Thus, when referring to the fast Hadamard transform (FHT), we imply the form defined in Eq. (2-3).

For example, a Sylvester-Hadamard matrix of order $N = 8$ is shown in the left panel of Figure 2-2. The first row of a Sylvester-Hadamard matrix always consists of only plus ones, which corresponds to a sequence of zero, while the second row always consists of alternating plus ones and minus ones, which corresponds to the highest frequency: $N - 1$. Sylvester-Hadamard matrices contain rows that span all of the frequencies from 0 to $N - 1$. However, due to its recursive Kronecker construction, the ordering of these rows is not in terms of frequency, but rather in the “natural” order for power-of-two Hadamard matrices. If it is desired, the rows (or columns) can be reordered according to increasing frequency and this form is called the Walsh-type Hadamard matrix:

$$W_N := P_W H_N$$

where P_W is the necessary permutation matrix. A Walsh-Hadamard matrix of order $N = 8$ is shown in the right panel of Fig. 2-2, where the rows increase in frequency from 0 to 7, downward from the top row.




Fig. 2-2. (a) Sylvester-type Hadamard matrix of order $N=8$. Permuting the rows of this matrix leads to the frequency-ordered Walsh-Hadamard matrix in (b).

2.3 Selective Measurement Strategies Enabled by the Block Structure of Kronecker product matrices

The characteristics of Sylvester-type Hadamard matrices can be summarized as follows:

- Efficient recursively generated patterns due to the Kronecker product: $H_{2^k} = H_2 \otimes H_{2^{k-1}}$
- Kronecker product also endows a fast transform and access to information at local scales and global scales which is useful for multiple-pixel CS
- Permuting rows of Sylvester-type Hadamard matrices creates the *Sequence* ordering of Walsh-type Hadamard matrices: $W_{2^k} = P_W H_{2^k}$
- Sequence facilitates the grouping of patterns used for CS into blocks with similar local underlying features called “signatures” which are helpful in efficiently choosing the fewest and most effective patterns for compressive anomaly detection

Our selective measurement strategy utilizes patterns obtained from a specially permuted Hadamard matrix that contains blocks of rows, of which each has a common local signature pattern. The sequence-based blocks partition the Hadamard spectrum, thus permitting analysis of the scene in terms of these local signature patterns. Note that Hadamard patterns are typically

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

described in terms of their sequency, which is a global property of each individual row. Instead, we only are only concerned with the local sequency structure.

The local-signature, row-block point of view can be beneficial since it permits us to adaptively select the best blocks with which to sense the signal/scene of interest, or to select the best blocks based on a priori information. As a result, in imaging applications more fine-scale detail appears in the scene, and in detection applications fewer false positives can result.

Note, this signature row-block partitioning is a general mathematical technique that can be applied to the Kronecker product of any two matrices, of any size (i.e., not just powers of two). As an example, we illustrate the technique by converting a Hadamard matrix to its Row-Block form.

Assume $N = F \cdot B$ and that there exist Hadamard matrices of order F and B such that

$$H_N = H_F \otimes H_B = \begin{bmatrix} h_0 \otimes H_B \\ h_1 \otimes H_B \\ \vdots \\ h_{F-1} \otimes H_B \end{bmatrix}$$

where we have denoted the rows of H_F by h_i . There exists a permutation P such that

$$H_N P = \begin{bmatrix} H_B \otimes h_0 \\ H_B \otimes h_1 \\ \vdots \\ H_B \otimes h_{F-1} \end{bmatrix} = \begin{bmatrix} \mathcal{B}_0 \\ \mathcal{B}_1 \\ \vdots \\ \mathcal{B}_{F-1} \end{bmatrix}$$

Here $\mathcal{B}_i = H_B \otimes h_i$ represents the i^{th} matrix block, with B rows and N elements per row, of the signature row-block Kronecker product. The indexing of the $[\mathcal{B}_i]$ blocks is with respect to the rows in H_F and all B rows in each $\mathcal{B}_i$ are made up solely from the unique signature corresponding to the pattern in row h_i .

As an example, we show the case of $H_N = H_F \otimes H_B = H_8 \otimes H_4$ seen in Fig. 2-3(a). . In Fig. 2-3(b), the columns of H_N are permuted; the resulting signature row-block matrix $H_N P$ has blocks of B rows, each with a unique length- F signature. In Fig. 2-3(c) the row-blocks of $H_N P$ are permuted to put the row-block local signatures in their sequency order. Note in this example F and B were both powers of two, and so the result is similar to a typical Walsh-Hadamard matrix. When B is not a power of two, then Figures (a)–(c) do not coincide with the familiar Sylvester- or Walsh-Hadamard matrices. Yet, at the same time, when F is constrained to be a power of two, we can exploit its spanning sequency property to examine local structures. This is illustrated in the next example.






Fig. 2-3 (a) Hadamard matrix $H_N = H_F \otimes H_B$; (b) the columns of H_N are permuted and each block of $B = 4$ rows has a unique length- $F=8$ signature; (c) the row-blocks of $H_N P$ are then permuted to put the block signatures in their sequence order.

In the previous example the unique signatures are rows of matrix H_F , and are one-dimensional in form. To better visualize the “sequency” content of each signature, we may re-format the signatures into two-dimensional tiles. While we still assume our Hadamard matrix to be of the form $H_N = H_F \otimes H_B$, we also require $F = 4^K$ so that the tiles will be squares of size 2^K per side. Figure 2-4 shows 2-D local signature tiles for the case of $F = 64$, arranged in Walsh sequence

order. Evident from this tile ordering is that the content of the signatures goes from low frequency to high as you move from the upper left corner to the lower right.





Fig. 2-4. Two-dimensional local signature blocks arranged in Walsh order

The value of the signature row-block structure is that we can use rows with known local signatures to extract information about an observed signal or scene. For many applications, it is the local information that is the most salient, e.g., the details, textures, or anomalies within an observed scene. For CS, we can choose to use all, or only some, of the patterns within a block according to our sampling budget.

Under our row-block signature structure, the CS measurement process is modeled by $y = Ax + e$, where the measurement matrix A is now given by

$$A = R(P_W H_N P)D$$

The term in brackets is the row- and column-permuted Hadamard, D is a diagonal matrix that acts as a coarse scale random modulator with solid $\{\pm 1\}$ over tiles the size of a signature tile, and R chooses signature blocks of interest and samples individual rows within those blocks.



Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).




Fig. 2-5. (a) Hadamard spectrum of image, x , where Hadamard encoded measurements are arranged in sequence blocks based on the signature tiles indicated in the graph; (b) Image, x .

Consider the image x in Figure 2-5(b), which contains $N = 768 \times 1024 = 64 \cdot 12288$ pixels. By choosing $F = 64$ blocks with $B = 12288$ patterns per block we can implement the desired signature row-block structure (where in this case B is not a power of two). The Hadamard spectrum $P_W H_N P_D$ of image x is shown in Fig. 2-5(a), where signals from selected signature blocks are notated. Note that the 2-D spectrum has been vectorized into 1-D, but that the measurements within each signature block (shown in Fig. 2-4) are still grouped together. The amplitude of the signal indicates the relative strength of the signature pattern in the image. Through experiments and simulations we have demonstrated that choosing a set of measurements from a complementary set of four (4) blocks is sufficient to perform anomaly detection in a process we call *Partial Complete*.^{6,7}

2.4 Partial Complete

The Partial-Complete (PC) schema partitions the Hadamard spectrum into separate blocks, where each block contains data with a specific scale of spatial frequency. By subsampling each of the blocks (e.g., 1%, or less) we obtain a sufficient statistic that determines the best blocks to use, thereby *adapting the sensing strategy* to use the waveforms that are best matched to the scene.

We also applied the PC sensing strategy to the simulations and by carefully selecting the measurement kernels, observed a concentration of the anomaly signal in well-defined blocks of the PC Hadamard spectrum. Figure 2-6 illustrates the motivation for choosing only these blocks. A direct visualization of signal in each transform block during movement of the anomaly (in red) shows that a majority of the energy is concentrated in this set when compared to background video sequence (in blue). While the amplitude within each block varies when the anomaly is spatially convolved with different regions of the background, the total energy between all blocks is nearly constant.

⁶ Herman, M. A., "Compressive Sensing with Partial-Complete, Multiscale Hadamard Waveforms," Proc. Optical Society of America, Imaging and Applied Optics (Arlington, VA, June 2013).

⁷ Matthew A. Herman, Tyler Weston, Lenore McMackin, Yun Li, Jianbo Chen, Kevin F. Kelly, "Recent results in single-pixel compressive imaging using selective measurement strategies," Proc. SPIE 9484, Compressive Sensing IV, 94840A (May 14, 2015);

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).




Fig. 2-6. Snapshots of Partial-Complete Hadamard spectrum as the anomaly crosses the image. Components in red show significant change between anomaly and non-anomaly cases and can be used for anomaly detection.

2.5 PC Block selection strategy

Recall for a scene with N pixels that we split the Hadamard domain into $F = 4^K$ non-overlapping blocks, with $B = N/F$ measurements per block. An important *dual feature* is that the *pixel* domain is partitioned at the same time into B non-overlapping tiles, each of size $2^K \times 2^K$ pixels. There is a strong periodicity induced in the pixel domain by the structure of the Hadamard waveforms as a result of their Kronecker product structure. So while we can choose to take measurements only from blocks selected from the Hadamard spectrum based on signal strength, one drawback is that certain periodic nodal patterns appear in the dual pixel (spatial) domain. The nulls in these nodal patterns record minimal, or in some cases zero energy, effectively creating blind spots in the spatial domain where anomalies would be overlooked. Therefore, in addition to selecting blocks based on spectral energy content we also want to select measurements from complementary blocks whose pixel-domain patterns contain no nulls when they are taken together.

We focus now on the case of scenes with $N = 256 \times 256 = 2^{16}$ pixels and $F = 64$ blocks (labeled as $B_0, \dots, B_{63}$), which results in $B = 2^{10}$ measurements per block in the Hadamard domain; in the pixel domain this results in $B = 2^{10}$ tiles partitioning the FOV, where each tile contains 8×8 pixels. As previously discussed, this can be easily generalized to the case of B that is not a power of two.

The energy patterns resulting from measurements from different Hadamard row-blocks are illustrated in Fig. 2-7, where we show an example image frame and highlight the region of the image in which we shift the anomaly pixel by pixel. Each of the single-block cases show nulls represented by dark-colored lines in the patterns. When the patterns are averaged together all of the nulls are removed, showing that these four blocks provide a complement to each other.




Fig. 2-7. Anomaly energy as a function of a 2x2 anomaly positioned within the red square of the example frame for four selected blocks of the Hadamard spectrum. Dark areas corresponding to low anomaly energy are eliminated when the blocks are combined.

The graphs above depict the *average absolute block energy* E_k of the anomaly in the Hadamard domain

$$E_k = \frac{1}{B} \sum_{i=0}^{B-1} |E_{k,i}|$$

where $E_{k,i}$ is the spectral energy of the anomaly in the i th measurement from block B_k

$$E_{k,i} = H_{B_k,i}^1 - H_{B_k,i}^0$$

and $H_{B_k,i}^1$ is the i th Hadamard measurement from block B_k with the anomaly in the scene, and $H_{B_k,i}^0$ corresponds to the scene when the anomaly is not present. Specifically, we denote the scene of interest *without* the anomaly as x^0 , and the scene x^1 *with anomaly* z as

$$x^1 = x^0 + z.$$


Note, since the Hadamard transform is a linear operator, their difference yields the Hadamard transform of the anomaly. In particular, for the i th mode (or measurement) of block $\mathbf{B}_k$ of the Hadamard transform of x^1 we have

$$H_{\mathbf{B}_k,i}^1 = H_{\mathbf{B}_k,i}(x^1) = H_{\mathbf{B}_k,i}(x^0 + z) = H_{\mathbf{B}_k,i}(x^0) + H_{\mathbf{B}_k,i}(z) = H_{\mathbf{B}_k,i}^0 + E_{k,i}.$$

Solving for $E_{k,i}$ yields the expression above. Note, other ways of determining $E_{k,i}$ are also possible. Other complementary groups include:

$$\{\mathbf{B}_4, \mathbf{B}_6, \mathbf{B}_{32}, \mathbf{B}_{48}\}, \{\mathbf{B}_{36}, \mathbf{B}_{38}, \mathbf{B}_{52}, \mathbf{B}_{54}\}, \{\mathbf{B}_2, \mathbf{B}_3, \mathbf{B}_{16}, \mathbf{B}_{24}\}, \\ \{\mathbf{B}_{34}, \mathbf{B}_{35}, \mathbf{B}_{50}, \mathbf{B}_{51}\}, \{\mathbf{B}_{20}, \mathbf{B}_{22}, \mathbf{B}_{28}, \mathbf{B}_{30}\}, \{\mathbf{B}_{18}, \mathbf{B}_{19}, \mathbf{B}_{26}, \mathbf{B}_{27}\}.$$

In addition to having no energy nulls when taken as a group, the members of each of these groups share a certain scale level, as shown in the tiles below. This makes each set a good choice for detecting anomalies of a specific size corresponding to that sequency level. In general, we found that the measurements from the group $\{\mathbf{B}_{36}, \mathbf{B}_{38}, \mathbf{B}_{52}, \mathbf{B}_{54}\}$ yielded the highest energy overall. Thus, without *a priori* knowledge of size of the anomaly this group seems to be wisest choice to focus on.






Fig. 2-8. Signature tiles of complementary Hadamard block sets corresponding to different anomaly sizes (e.g., for size 2×2, we would choose measurements from group { $B_{18}, B_{19}, B_{26}, B_{27}$ }).

2.6 Local global specifically designed for multi-detector architectures

In this project we transitioned from a single-pixel architecture to a multiple pixel architecture to increase the speed of acquisition of compressive sensing measurements for imaging and for anomaly detection. The hardware development part of the program implemented a multi-pixel architecture using a small array of InGaAs pixels (64 x 64) in place of a single photodetector. The hardware is described in more detail in a later section. Here we note that the highest efficiency is gained when compressive measurements are not only acquired in parallel by multiple detectors, but also that the field of view corresponding to each of the detectors is encoded in a way that relates it to its neighbors.

We completed a description of theory and Matlab implementation of dual Local-Global measurements, which is a measurement technique that is specifically designed for multi-detector architectures. The key enabling idea is that an $N \times N$ transform matrix H_N that can be decomposed as the Kronecker product of two smaller transforms possesses tremendous structure of which we can take advantage for efficient imaging and anomaly detection. The mathematical analysis below shows that we have access to both local and global transforms. This means that we can compare the statistics of transform data associated with different sized scenes providing

access to multiple scales of data from one set of measurements. This can help to detect anomalies efficiently.

We developed a process based on Local-Global measurement strategy to relate the local data (i.e., taken in the compressive measurement domain) to the global domain without having to reconstruct the scene in the pixel domain. Based on careful mathematical analysis, we have developed a technique that allows us to choose just a few detectors (e.g., any contiguous 2×2 , 4×4 , 8×8 , etc.) and groom their measurements together as if just a single detector was viewing the respective larger FOV.

This new method can be of benefit since we can now look at the statistics of multiple scales of a scene of interest in the compressive measurement domain. In fact, we can now have access to all of the local and global data, as well as all scales in between all at the same time. Examining the statistics at different scales should provide another tool with which to detect anomalies.

The key enabling idea to Local-Global is that an $N \times N$ transform matrix H_N that can be decomposed as the Kronecker product of two smaller transforms

$$H_N = H_F \otimes H_B$$

possesses tremendous structure that we can take advantage of. For example, the well known Sylvester Hadamard and STOne transform matrices satisfy this Kronecker product decomposition. In particular, we can use the mixed-product property of Kronecker products to show that

$$H_N = (H_F \otimes I_B) (I_F \otimes H_B) \quad (2-5)$$

where I_n denotes an $n \times n$ identity matrix. Here, $(H_F \otimes I_B)$ is a sparse matrix and $(I_F \otimes H_B)$ is block diagonal. Then the H_N -transform of a length- N vector x

$$X = H_N x$$

can be viewed as a two-stage process associated with the mappings $(H_F \otimes I_B)$ and $(I_F \otimes H_B)$ in (2-5). The first stage consists of *optically* acquiring the local transform data of smaller non-overlapping blocks of the signal. Since $N = BF$, we split the signal x into F non-overlapping groups of length B

$$x = [x_0, x_1, \dots, x_{F-1}]$$

and denote X_k as the H_B -transform of the k th group x_k , for $k = 0, 1, \dots, F-1$. The set $\{X_k\}$ contains all the local or micro information of the scene, which is obtained in parallel in the first stage of the process via the *intermediate transform*

$$X_{\text{Int}} = (I_F \otimes H_B) x = [X_0, X_1, \dots, X_{F-1}]. \quad (2-6)$$

In the case of our proposed focal-plane array (FPA) transform-camera, the k th detector acquires the H_B -transform data from the small scene x_k .

The second stage of the process, implemented in hardware, consists of mathematically multiplexing the local transform data using the macro weighting patterns contained in H_F via the mapping

$$X = (H_F \otimes I_B) X_{\text{Int}}. \quad (2-7)$$

Thus we have access to both local and global transform data, essentially simultaneously. Moreover, assuming that a matrix $H_{F'}$ exists with $F' < F$, we can obtain smaller global transforms of size $N' = BF'$

$$X' = (H_{F'} \otimes I_B) X'_{\text{Int}}$$

where X'_{Int} contains just F' of the X_k (note that X_{Int} in Eq.(2-6) contains all F of the X_k), which do not have to be contiguous groups. This means that we can compare the statistics of various transform data $\{X_k\}$, $\{X'\}$, X associated with different sized scenes; essentially, we have access to multiple scales of data. This can help to detect anomalies.

2.7 Mean-to-deviation versus SNR analysis of compressive measurements

Compressive sensing measurements are fundamentally different from pixel-based imaging and detection methods. Previous studies⁸ have examined theoretical aspects of signal-to-noise (SNR) comparing single-pixel based imaging to focal plane array imaging. While instructive, these studies are difficult to interpret in experimental systems. We have we examined how to better assess the signal-to-noise ratio (SNR) of our camera system. Specifically, we have changed our perspective to now be with respect to the *optical signal*, since this is the domain in which we actually make our measurements.

First, we defined a new metric, which we have named the *mean-to-deviation ratio* (MDR). For a given set of ideal, noiseless samples in the optical domain, $y_{\text{opt, clean}}$, denote

$$\mu_{\text{clean}} = \text{mean}(y_{\text{opt, clean}}) \quad \text{and} \quad \sigma_{\text{clean}} = \text{std}(y_{\text{opt, clean}})$$

as their mean and standard deviation, respectively. Note that the mean μ_{clean} provides no salient information of the observed scene except its overall brightness. Instead, the information of the scene is encoded into the **deviations around the mean**. Thus, the standard deviation σ_{clean} provides a metric describing the average energy of this encoded information. The mean-to-deviation ratio (MDR) defined as

$$\text{MDR} = \mu_{\text{clean}} / \sigma_{\text{clean}},$$

⁸ Joel A. Tropp, Jason N. Laska, Marco F. Duarte, Justin K. Romberg, and Richard G. Baraniuk, "Beyond Nyquist: Efficient Sampling of Sparse Bandlimited Signals," in Information Theory, IEEE Transactions on , vol.56, no.1, pp.520-544, Jan. 2010.

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

essentially quantifies how much larger the brightness of the scene is as compared to the energy of its encoded deviations, providing a metric analogous to signal-to-noise ratio specifically describing signal levels in a compressive sensing measurement system. In terms of decibels, MDR can be expressed as

$$\text{MDR}_{\text{dB}} = 20 \cdot \log_{10}(\text{MDR}).$$

The MDR is a good metric to assess the limits of the dynamic range of an observed scene's measurements, but it is only with respect to an ideal, noiseless scenario. In practice, we observe noisy measurements in the form of

$$y_{\text{opt, meas}} = y_{\text{opt, clean}} + e,$$

where e represents additive noise. Denote $\mu_{\text{meas}} = \text{mean}(y_{\text{opt, meas}})$ as the mean of the noisy signal, and $\sigma_{\text{noise}} = \text{std}(e)$ as the standard deviation of the additive noise. Assuming the noise is of zero mean, we can claim $\mu_{\text{meas}} \approx \mu_{\text{clean}}$. Then, with the SNR defined as

$$\text{SNR} = \mu_{\text{meas}} / \sigma_{\text{noise}} \approx \mu_{\text{clean}} / \sigma_{\text{noise}},$$

we can examine the ratio of SNR to MDR:

$$\text{SNR/MDR} \approx (\mu_{\text{clean}} / \sigma_{\text{noise}}) / (\mu_{\text{clean}} / \sigma_{\text{clean}}) = \sigma_{\text{clean}} / \sigma_{\text{noise}}.$$

Clearly, this means that the SNR of CS measurements needs to be better than the MDR, i.e., the energy of deviations due to the encoded scene must be greater than the energy of deviations contributed by the noise. We are currently in the process of quantifying experimental signal requirements based on this analysis. For example, for images viewed with a 256×256 portion of the DMD, we found most scenes to have an MDR_{dB} in the range of 45–48 dB. We found a good rule of thumb is for the SNR to be at least 10 dB greater than the MDR.

2.8 Z-Score and MMD

During Phase I, we explored new modes for change detection based on Eq. (2-5) observing the statistics of groups of compressive measurements in the time domain. Anomalies can be detected because they create compressive measurements that are statistical outliers.

The “Z-score”, also known as standard score in statistics, is the number of standard deviations an observation or data is away from the mean. Thus, a positive standard score represents a datum above the mean, while a negative standard score represents a datum below the mean.

The “MMD”, also known as maximum mean discrepancy⁹, is a criterion that was used in Phase 1 to determine whether two sets of data are from the same or different probability distributions.

We found that the product of these two criteria created a sensitive statistic for anomaly detection in the compressed domain.

⁹ Gretton, A., Borgwardt, K. M., Rasch, M., Schölkopf, B., and Smola, A. J., “A kernel approach to comparing distributions,” in [Proc. of the 22nd AAAI Conf. on Artificial Intelligence], 1637-1641, AAAI Press (2007). Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

2.9 Anomaly detection statistics testing simulations with Partial-Complete Block Sets

Using the MMD-Z-score product as our statistical metric, we investigated the performance of compressed domain anomaly detection in simulation. The point of these simulations was to determine whether measurements were optimized by selecting them from partial-complete block sets of the Hadamard spectrum on the basis of both 1) similarity of the sequency pattern to the anomaly size and 2) complementarity to other patterns for full field of view coverage without nulls. The results of these simulations, shown next, confirm that this selection process performs well.

The simulation setting is identical to the cases used throughout this project using the same set of video frames to simulate a changing scene. As illustrated in Fig. 2-9, two cars are pulled towards each other in front of a color-checker board at frame rate of 250 fps at a resolution of 256×256 . We let the anomaly appear once in the sequence of frames, starting at 1/3 of the simulation time and lasting for approximately 1/3 of the total time. As shown in Figure 1(b), the anomaly was a 4×4 block of pixels with an intensity of approximately 8 times the scene average. To better simulate the problem, we gave the anomaly intensity a Gaussian profile.




Fig. 2-9. From left to right are three frames of scene images (256×256) showing two toy cars moving towards each other. In (b), a bright pixel-block of size 4×4 in the background represents an anomaly point.

In the first simulation we set the anomaly at pixel location (76, 76) and tested all candidate block sets. We added 5dB image noise to the movie. Since the scene has $N = 256 \times 256 = 64 \cdot 1024$ pixels, if we use $F = 64$ signature blocks, then each block will contain $B = 1024$ patterns to observe the scene with. We begin with using four full blocks in each set as the detection patterns. Thus, the 100% detection cycle is 4096 measurements in length. Then, to reduce data acquisition requirements, we only use a randomly selected quarter of the measurements within each block so the total cycle length for results shown in Fig. 2-10 is 1024. We repeat the same strategy for each of the test block sets, and plot the results below. From Figure 2-10(a) we see that the use of Block Set 1 is encouraging since the raw data with and without the anomaly look so similar, yet the algorithm is able to detect the anomaly. Results from Block Set 2 in Figure 2-10(b) also contain good detection results as expected. Block Set 3 in Figure 2-10(c) does not detect the anomaly. This was not surprising since the signature contained in the Set 3 Hadamard block corresponds to patterns with 2-pixel-wide stripes (see Fig. 2-8), which does not complement the 4×4 -pixel structure of the anomaly. Similarly, results for trial Block Sets 4, 5, and 6, shown in Figure 2-11, do not show anomaly detection.




Figure 2-10. Detection results via (a) Set1, (b) Set2, (c) Set 3. In each case the detection window is data 256 points. The first row is the synthetic data; the second row is the change (anomaly indicator) plotted on a semilog scale; the third row is the change plot; fourth row is the Z-score; and the last row shows the MMD score.




Fig. 2-11. Detection results via (a) Set 4, (b) Set 5, (c) Set 6. In each case the detection window is 256 data points show no change is detected. This is expected due to the mismatch between selected pattern blocks and the characteristics of the anomaly.

We also conducted laboratory experiments on the Partial-Complete block detection scheme. We used a laser pointer to represent an anomaly in a scene imaged by the single-pixel camera (SPC). For a fixed FOV, the scene with and without an anomaly was acquired at two different resolutions: 128x128 pixels and 256x256 pixels. During the acquisition, the car moved across the FOV horizontally. The dominant part of anomaly point (laser point) was verified as having a FWHM support of approximately 4x4 pixels in the finer scale 256x256-pixel scene, as shown in the inset to Figure 2-12. When coarsified to the 128x128-pixel scene, this same anomaly should have a support of approximately 2x2 pixels, enabling us to test two different anomaly scale sizes.




Fig. 2-12. Left image is the example frame captured by the SPC (256x256); Right image is a zoom in of the anomaly.

2.10 ROC and PR curve development and calculation of thresholds

While it is instructive to display the results of simulations and experiments in a format that includes the data, the MMD and Z-score and the combined detection metric in four separate graphs for each measurement, we have generated receiver operating characteristic (ROC) curves¹⁰ as a way to evaluate and present the performance of our anomaly detection algorithm for different scenarios. ROC curves are a convenient and succinct way to illustrate the behavior of a binary classifier. Our change detection algorithm, at any given moment of time, gives a “score” or a probability as to whether or not an anomaly is present. When this score is above or below a certain threshold, the algorithm acts as a binary classifier.

2.11 Background on ROC curves

ROC curves are based on a so-called “*confusion matrix*,” which consists of the total number of true positives (*TP*), false negatives (*FN*), false positives (*FP*), and true negatives (*TN*) that were determined by the binary classifier for a given threshold. Table 1 shows the canonical confusion matrix, where the columns indicate whether a particular event was **actually** a positive or negative

¹⁰ Fawcett, Tom (2004); [ROC Graphs: Notes and Practical Considerations for Researchers](#), Pattern Recognition Letters, 27(8):882–891.

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

occurrence, and the rows indicate whether that event was *classified* as a positive or negative occurrence. The *TP*, *FP*, *FN*, *TN* labels then follow directly from their respective definitions.

Table 1. The confusion matrix associated with a binary classifier.

	Actual Positive	Actual Negative
Positive Classification	<i>TP</i>	<i>FP</i>
Negative Classification	<i>FN</i>	<i>TN</i>

Given the confusion matrix for a certain situation, the ROC curve is simply a graph of the *true positive rate*:

$$TPR = \frac{TP}{TP+FN} \quad (2-8)$$

versus the *false positive rate*:

$$FPR = \frac{FP}{FP+TN} \quad (2-9)$$

2.12 Development of Precision Recall curves: Better than ROC

The well known metric of a ROC curve simply illustrates the *true positive rate* (TPR) *versus the false positive rate* (FPR). However, we have noticed that certain characteristics of the statistical change detection methods used with compressive data acquisition prevent ideal ROC performance. Specifically, for a single positive event (i.e., an anomaly appearing or disappearing) the change detection outputs a series of large spikes in a window of approximately 4096 (each point corresponds to a moment of time). The location of the spikes within this window is not uniform, and is always different. This is due to the different times as well as different locations of the anomaly appearing and disappearing from the scene. In order to assign whether the algorithm has properly classified a moment of time as a *true positive* (TP) or a *false positive* (FP), we must be able to correlate that moment of time with whether the original moment of time was actually a positive or negative event. In the current algorithm, this means that although we are correctly classifying positive events, we are unable to properly quantify it in the context of a TPR and a FPR, and to display their relationship in an ROC curve.

Fortunately, the ROC curve is just one possible metric that can be derived from the confusion matrix. When dealing highly skewed events (e.g., a very brief anomaly) a different metric, called the *Precision-Recall* (PR) curve, can better capture the behavior of binary classification algorithms.¹¹ “Recall” is the same as the TPR in Eq. (2-8). It only uses information from the **first column of the confusion matrix**, and answers the question:

*“For all of the **actual positive events**, what percentage of them were **correctly classified**?”*

¹¹ J. Davis and M. Goodrich, “The Relationship Between Precision-Recall and ROC Curves,” Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, PA, 2006.

The new measured rate, *Precision*, is defined as

$$\text{Precision} = \frac{TP}{TP+FP}. \quad (2-10)$$

Notice that it only uses information from the **first row of the confusion matrix**: these are the events that were classified as positive. Hence, Precision answers the question:

*“For all the events that were **classified as positive**, what percentage of them were **actually positive**?”*

Observing brief anomalies present a scenario whose positive and negative events are highly skewed, say by a ratio of 4–6 orders of magnitude. Since the ROC curve is insensitive to the skew of classes, we instead constructed and analyzed the PR curve, which can tell much more about the behavior of algorithm.

In the following section we summarize the results of our experimental investigation of high speed event detection in the compressed domain.

Section 3. Implementing adaptive compressed sensing methods on a focal plane array

The single pixel model of compressive sensing drove the imaging and anomaly detection simulations of the previous section. Extending the design of the imaging and detection system to include a multi-pixel sensor improves our range of possible applications. The multi-pixel model allows us to multiplex many measurements simultaneously, allowing for higher speeds. The multi-pixel model is also suited to the Partial-Complete and Local-Global measurement strategies which can add further efficiency to anomaly detection, foreground/background subtraction and library learning.

In Section 2, we outlined several measurement selection strategies, including Partial-Complete measurements, which can be implemented in either single- or multiple-pixel systems, and Local-Global measurements, which are taken from multiple detectors. The key enabling idea in both theories is that an $N \times N$ transform matrix H_N that can be decomposed as the Kronecker product of two smaller transforms. The recursive structure of the Kronecker product is a fundamental enabler to efficient measurement acquisition. Our mathematical analysis showed that we have access to both local and global transforms, meaning we can compare the statistics of transform data associated with different sized scenes (that is, multiple scales of data) from one set of measurements. This can help to detect anomalies efficiently.

In this section we first describe the multi-pixel compressive camera for imaging and anomaly detection, and virtual channel geometries for multi-pixel compressive measurements. We then

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

describe development and testing of point-spread function calibration matrix for correction of optical path distortions and image stitching.

3.1 Anomaly Detection Table Top Demonstrator and Laser Anomaly Design

To support anomaly detection simulations and experimental implementations, a multi-pixel imaging system has been developed at InView. In this imaging system, the InView camera is focused on a target, the DMD modulates the target scene, and the 64 x 64 sensor array captures a sequence of images of the modulated image with a frame rate of up to 1 kHz. After the array is read out, the pixels of the image sequence are electronically binned into separate virtual channel readings. The system is shown in Fig. 3-1 below.




Figure 3-1. Compressive camera with multi-pixel optics and electronics modifications used in our experiments. Parts labeled are: (1) pattern generator BNC connection; (2) pattern generator PCIe connection; (3) sensor array BNC connection; (4) CameraLink connection to the Hamamatsu 64 x 64 sensor array.

3.2 Hardware Synchronization

The system required optical and electrical modifications to the original single-pixel camera to support the multi-pixel sensor. The main optical modification replaces a condensing lens system located after the DMD with an imaging system that relays the modulated DMD image to the detector plane. Electrical modifications were required to synchronize the patterns shown on the DMD with the measurements taken by the multi-sensor Hamamatsu detector array. The BNC cable shown in Fig. 3-1 connects the DMD pattern generator's sync signal with the Hamamatsu detector's trigger line.

The system utilizes two computers to drive the hardware. One PC connects via PCIe to the DMD pattern generator. Another PC connects to the sensor via CameraLink. The synchronization of DMD patterns and sensor acquisition is made possible with the BNC trigger line. Controlling sensor features such as exposure time and trigger source require use of the Hamamatsu custom API rather than the standard Camera Link API. To enable image data acquisition, we bridge the two controlling PCs by running a network daemon on each PC.

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

At the start of an acquisition, the pattern generation and sensor computers perform a network synchronization. First the machines negotiate the number of buffers they are each able to record for a single time-continuous experiment. Then the sensor readies itself to receive the triggers. Once the sensor is ready and listening for the first trigger, pattern generation starts. This process is shown in the diagram of Fig. 3-2.




Figure 3-2. Synchronization diagram showing communication links between the Pattern Generator, DMD controller and Hamamatsu multi-pixel sensor for acquiring CS data.

Some example Hadamard data acquired using this process is included below in Fig. 3-3. The chart graphs ten sensors from the middle of the Hamamatsu array as they recorded 40 modulation patterns from the DMD. The PCIe transfer maximum speed is currently more than half of the maximum frame rate of the Hamamatsu sensor (1 kHz), so we are currently not oversampling the DMD modulation.




Figure 3-3. CS data from Hamamatsu multi-pixel based camera taken in sync with patterns generated on the DMD and data acquisition control signals.

3.3 Laser Trigger Anomaly Control Experiment Results

We introduced a software system to control the laser and car movement mechanism in the anomaly detection experiment set, in order to obtain the ground truth of the anomaly that match with the raw data measurements, and then use the data sets for the ROC and PR curve analysis. We show below the experimental results of the laser trigger system that verify the synchronization between the laser trigger ground truth and the anomaly measurement data.




Figure 3-4. 2-channel Relay set used to precisely control the appearance of a laser anomaly in anomaly detection experiments.

The trigger system is realized using a two-channel relay set like the one shown in Fig. 3-4. A relay is an electrically operated switch and is used where it is necessary to control a circuit by a low-power signal (with complete electrical isolation between control and controlled circuits such as TTL or ADC voltage output), or where several circuits must be controlled by one signal. In our experiment the relay set has two channels: one controls the laser as an anomaly and one controls the motor that pulls the car moving across the scene. Fig. 3-5 shows the results of the measurement data overlay with the laser trigger ground truth that was obtained both from the National Instrument analog and digital convert device. We can see that the trigger edges of the laser match well with the light intensity increase in the measurement data.




Figure 3-5, A comparison between the laser trigger ground truth (red) and the measurement data (blue), the DMD is running random Walsh-Hadamard pattern at 1 kHz.

The relay system allows us to control the desired event precisely, providing confidence that consistency with measurements can be achieved.

3.4 “Virtual channel” geometry for imaging and anomaly detection

In this section, we describe how we choose image sub-regions for imaging and anomaly detection by grouping FPA pixels on a 64×64 Hamamatsu detector array into *virtual channels*. Of all possible configurations of sensors and mirrors the 8×8 array of virtual sensors was found most convenient for testing the partial complete (PC) measurement selection strategy. This “canonical” configuration of grouped FPA detectors and corresponding DMD mirrors is set up in the following way:

An area of the DMD encompassing a square group of 512×512 mirrors is subdivided into an 8×8 array of 64×64 -mirror groups. The Partial-Complete (PC) patterns we put on DMD have a resolution of 64×64 mirrors and are duplicated across each of these 64 mirror groups on the DMD to fill the 512×512 -mirror area. In the ideal case, the FPA pixels are binned into an 8×8 array of virtual detectors. As a result, each of these 64 virtual channels (of FPA pixels) receives a signal from a block of 64×64 mirrors on the DMD. In reality, the mapping from DMD mirror groups to the virtual detector channels is not clearly delineated because optical aberrations are introduced by the DMD and the relay optics. In finding a mapping from the DMD to the channels of the FPA in multi-pixel compressive sensing, an important issue is how to deal with image distortions. In this project we investigated using an experimentally determined point spread function (PSF) to describe the actual mapping from the DMD to the

detectors, and implemented it as a matrix inserted into the optimization process, as described next.

We used a square 512×512 portion of the DMD to sense an observed scene. Due to the optics, the image of scene on the FPA was skewed. A snapshot of a “flat” (i.e., empty) scene imaged with all of the DMD mirrors set to the ON position is shown in Fig. 3-6. We can clearly see skewing of the square FOV.




Figure 3-6. Distortion and skew can be seen by observing a flat (empty) scene and turning the active 512×512 portion of the DMD mirrors to the ON position.

As mentioned previously, we created an 8×8 array of *virtual channels*, where each channel is comprised of a portion of the FPA’s active 4096 detectors yet treated as a distinct “single-pixel” camera. Ideally, each virtual channel should only receive light from a 64×64 region of the DMD’s 512×512 active mirror area. However, due to the skewing, distortion and optical crosstalk, the mapping of DMD mirrors to virtual channels is non-trivial. In order to determine the mapping we rastered a 64×64 region across the DMD’s 512×512 active area and recorded which of the FPA’s detectors received light. The resulting 8×8 array of snapshots from this raster scan can be seen in Fig. 3-7.




Figure 3-7. The 8×8 array of virtual channels obtained by rastering a 64×64 section of the active 512×512 portion of the DMD. This provides the mapping from the DMD mirrors to the 64 virtual channels.

Fig. 3-8 shows what happens when we superimpose images all of the virtual channels. The left panel shows the result from simply summing up the raw snapshot frame data from all 64 raster shifts; we also see a significant amount of background noise, which was found to be due to detector dark current.




Figure 3-8. Superimposing the 8×8 array of virtual channels into one image. (Left) Result using raw frame data, f_i . (Center) Dark noise frame, f_{dark} . (Right) Result after each data frame has been adjusted for dark current, $f_{i,\text{adj}}$, as in Eq. (1). Note that the scale is on the order of 10^4 for frame data and on the order of 500 for the dark frame.

The background noise due to dark current is quantified by taking measurements with the sensor array when there is no light present. We collected multiple snapshot frames (i.e., a *frame* is a full set of 64×64 measurements from the FPA from one moment in time), and then averaged them together to generate an estimate “dark” frame, f_{dark} , as seen center in Fig. 3-8. With the average dark frame, we can now approximately remove the noise due to dark current by subtracting f_{dark} from each collected frame of data. Denoting f_i as the i th data frame, the *adjusted data frames* are simply calculated as

$$f_{i,\text{adj}} = f_i - f_{\text{dark}}. \quad (1)$$

The right panel of Fig. 3-8 shows that summing of the adjusted data frames $f_{i,\text{adj}}$ effectively removes the noise due to dark current.

Using the virtual detector geometry we have been describing and the method for manually determining the optical boundaries of each virtual channel we began some experiments in multi-channel anomaly detection. Figure 3-8.1 below show 8 x 8 virtual channel reconstructions of a star target scene and that same scene in which have been injected a handful of bright laser spots. Fig. 3-8.1(a) contains the background image, while Fig. 3-8.1(b) contains the bright anomalies. Looking at the top row of sensors in the anomaly image, counting from the left, the first ROI is “Detector 1”, the second ROI, “Detector 2”, has a long diagonal streak, the third ROI, “Detector 3” has a speckle pattern with 3 small Gaussian objects, and the fourth ROI, “Detector 4” has a single small Gaussian object recovered. We used data from Detector 3 to exercise our anomaly detection algorithms.






Fig. 3-8.1. Recovered images from multi-pixel experimental data. There are 64 regions of interest. (a) The no-anomaly background scene anomalies, and (b) the scene with anomalies.

From the Detector 3 data we have extracted the measurements corresponding to Sets 1, 2 and 6 of the complementary blocks shown in Fig. 2-8 of Section 2. Using these measurement blocks, we have run the anomaly detection algorithm on the data with and without anomalies. Fig. 3-8.2 shows the signal, z-score, MMD and Anomaly Indicator graphs created from the actual DVC camera measurements. The results show that all Sets 1, 2, and 6 of complementary PC signature blocks detect the 4x4 and 2x2 anomalies appearing in the ROI.





(c)

Fig. 3-8.2. Anomalies appearing in Detector 3's ROI (shown as a picture in Fig. 3-8.1(b)) are detected using Partial Complete block sets 1, 2 and 6 of complementary signature blocks (shown in Fig. 2-8).

3.5 PSF Development for Multi-pixel Calibration

3.5.1 Raster-scan for Point Spread Function Measurement

In order to get rid of optical distortion and enhance the joint reconstruction of data from multiple sensors without resorting to tedious manual mirror-to-FPA mapping, we have implemented a point spread function (PSF) that mathematically describes system pixel geometry and aberrations in a matrix formulation that can be experimentally calibrated and used directly in the optimization procedure. A mathematical model of the PSF-aware measurement and optimization process is given in Appendix 3-1. The reconstruction algorithm fuses an array of smaller images together by making use of the optical system's PSF. The mathematical model of the blurring PSF operator is separated from the sensing patterns as an independent matrix so that measurement patterns maintain an associated fast transform method and the proposed reconstruction algorithm can be implemented in a fast and efficient way.

It is straightforward to experimentally obtain the PSF matrix by raster-scanning groups of mirrors on the DMD and acquiring an image of each "point" on the detector array. In this way, each image contains local PSF information. This local information arranged in corresponding rows of a large matrix, where each row represents a raster pattern on the DMD. To make the process faster, a multi-pixel raster is achieved by dividing the whole DMD into number of blocks where each block has its own raster-scan operating in parallel with the others. One critical thing is that the blocks should be big enough so that light blur between any two raster points not overlap with each other. Fig. 3-9 depicts the single and parallel raster methods.




Figure 3-9. (a) Conventional raster-scan method, the red line represents how the PSF of the DMD pixel moves on the detector array.(b) Multi-pixel raster-scan method, one pixel within multiple blocks are lighted to perform the parallel raster-scan

We have collected calibration and imaging data using these procedures and have processed the results using the PSF methodology along with other methods to create a calibration matrix that is used in the image optimization process, as described next. The results have shown reduction in image distortion obtained within the image optimization process without added post-processing steps.

Following the notation in Appendix 3-1, the overall PSF matrix, Q , is the product of Q_{virt} and Q_{FPA} , where Q_{virt} is the response of the FPA from coarse rastering (e.g., groups of 64×64 mirrors) of the DMD, and Q_{FPA} is the response of the FPA from fine single-mirror rastering of the DMD. The challenge with rastering just *a single mirror* is that not much light is received, and thus the SNR is very low. We have worked to optimize several aspects of our measurement process to ensure a high photon count, and thus a relatively good SNR.

Below are figures with the results from the single-mirror rastering. Fig. 3-10 shows the full raster data in its 1-D vectorized form. The active portion of the DMD is $512 \times 512 = 262,144$ mirrors; thus we have 262,144 total raster measurements.




Figure 3-10. (a) Single-mirror raster data made by turning only 1 mirror to the on position on the DMD and measuring the response of the single-mirror illumination on the detector array. (b) Single-mirror raster data reshaped into its 2-D image form reconstructs a view of the DMD from the point of view of the detectors.

Fig. 3-10(b) shows the raster data reshaped into its 2-D image form. Here we can see that illumination of the FOV is not uniform, with the brightest region in the center. Thus, the middle section of measurements in Fig. 3-10 (a) has larger values than the beginning and end. Further, the dark spot in the center of Fig. 3-10 (b) clearly corresponds to the dip in intensity in Fig. 3-10 (a). It is interesting to observe the apparent periodic structure in Fig. 3-10 (b), which looks somewhat like graph paper. Note that this image is actually a view of the DMD as seen through the “eyes” of the FPA’s 4096 detectors. Therefore, we believe that the dark lines are due to the light that is lost in between the detectors.

Improvements in experimental imaging results using the PSF-aware algorithm are shown in the images below. Fig. 3-11 shows the results of a simulation where a noiseless ground truth image was “corrupted” using the optical distortion information from the PSF and then used as the scene for PSF-aware measurement and reconstruction. This simulation shows the potential for correcting the distortion incurred by the optical system’s PSF.

In Fig. 3-12 we see image reconstructions taken from our 64×64 multi-pixel camera partitioned into 8×8 array of 64 virtual channels. In panel (a) we see each virtual channel independently reconstructed and stitched together to form the original FOV. However, panel (b) shows the result of using the same exact measurements with our PSF-aware reconstruction algorithm.




Figure 3-11. Simulated PSF-aware reconstruction (a) ground truth, (b) optically distorted image created from PSF data showing skew and virtual boundary artifacts, and (c) improved PSF aware reconstruction




Figure 3-12. A total of 64 virtual channels were reconstructed (a) without and (b) with the PSF-aware partial-complete algorithm. Considerable improvement to artifacts at the boundaries of the virtual channels was achieved.

We next describe the PSF approach for reconstructing data sets using PC with a goal of optimizing experimental performance, and also begin implementation of MMD and Z-Score algorithms for event detection.

3.6 Focal Plane Array (FPA) anomaly detection

As proposed, we are implementing anomaly detection on a focal plane array (FPA) to address the limitations of SPC system. We are using a 64 x 64 Hamamatsu InGaAs FPA, which is sensitive to short wave infrared (SWIR) wavelengths. With the FPA each detector in the FPA acts as a SPC system, but sees a smaller region of interest (ROI). Smaller ROIs have several advantages including the ability to detect fainter anomalies and also be able to detect and locate multiple anomalies appearing simultaneously. By working in parallel, multiple smaller ROI's also produce faster operation. Using virtual channels allows us to control the ROI size on the DMD algorithmically, and without hardware changes.

During our preliminary test of the FPA system, we found that because of imperfections in the relay optics and the Scheimpflug effect of the digital micro mirror (DMD), the projected image from DMD to the FPA is blurred and skewed; these are conditions that will affect the detection performance of the system. Although we cannot currently correct for the effect of these optical distortions in the detection step as we did in the reconstructed images described above, we can still incorporate the PSF into our simulations in order to generate more realistic results. We implemented the PSF in anomaly detection using the FPA virtual channel geometry and investigated detection performance in the ideal case and the non-ideal measured PSF case.

In this simulation, the PC patterns we put on the DMD for each virtual channel have a local resolution of 64 x 64 mirrors and are duplicated across the DMD resulting in a 512 x 512 total

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

Page 12 of 74

resolution. In the ideal case, the virtual channels are groups of 8×8 FPA pixels that are binned into a single virtual detector. This results in an ideal virtual FPA comprising a total of 64 channels arranged as an 8×8 array of 64 virtual detectors, where each of the channels receives a signal from a block of 64×64 mirrors on the DMD. Fig. 3-13 shows which pixels on our 64×64 FPA constitute virtual channel 12. In panel (a) we see the exact group of 8×8 pixels that correspond to the ideal case, while in panel (b) we see that in reality, the PSF causes some distortion and skew. Note that the non-ideal PSF of the system was carefully measured using the scanning procedure described above.




Figure 3-13. Virtual channel #12 with ideal PSF and measured PSF

We conducted an anomaly simulation using a video dataset of resolution 512×512 pixels that is consistent with the real multi-pixel camera. The scene is an urban highway as background with vehicles moving across (Fig. 3-14). Ideally, the image on the FPA is an exact replica of the original image scaled down by demagnification, as seen in Fig. 3-15(a). However, Fig. 3-15(b) shows that in reality the image on the FPA is skewed and blurred. Fig. 3-15(c) shows how the virtual channels appear after grouping the pixels of the FPA together.




Figure 3-14. Ground truth of one frame from the test movie set.




Figure 3-15. Image on the FPA, with different PSF. (a) has the ideal PSF with no distortion and blur, (b) is measured image from real world FPA, (c) is the 8 x 8 virtual channels that grouped from the measured FPA image.

As discussed in previous reports, the PC patterns have the advantage over random permuted WH patterns, as it can be categorized into blocks of different modes that share a unique signature, some of which will have a bigger response to anomaly of particular size. We ran simulations with a 4x4 anomaly appearing in the middle of virtual channel 1, in the top left corner of the FOV, as seen in Fig. 3-16. We used 100% of the PC patterns in the previously selected “optimal” group of complementary blocks $B_{36}, B_{38}, B_{52}, B_{54}$. The intensity of the anomaly was 2000 versus the image’s maximum intensity of 255; roughly 8 times brighter. Noise was added to the test data set to simulate the real world experiment, and here we start with a SNR level of 15.




Figure 3-16. Ground truth of the anomaly (left) and the ROI perspective from virtual channel 1 (right).




Figure 3-17. Anomaly detection results with the ideal PSF (left) and measured PSF (right) in virtual channel 1. The distorted image gives less signal of the change detection.

In simulation, we found that the detection signal of the anomaly is slightly higher in the ideal PSF case than the measured PSF one (Fig. 3-17). We also ran other simulations with different anomaly intensity, different SNR and different virtual channels, all of which have similar results (Figs. 3-18 through 3-20).




Figure 3-18. Detection result of anomaly with different SNR in virtual channel 1, the anomaly intensity is also set to 1000




Figure 3-19. Ground truth of the anomaly (left) and the ROI perspective from virtual channel 20 (right).




Figure -3-20. Anomaly detection result from another virtual channel (channel 20), the detection is slightly different compare to other channel, but the blur and skew affects the detection performance the same way.

We also simulated the scenario where an anomaly appears at the edge of a virtual channel ROI. Fig. 3-21 shows just the ground thruth ROI for virtual channel 20 with the anomaly located in its top-left corner. The crosstalk between FPA detector and image distortion will significantly affect the detection result as shown in Fig. 3-22.




Figure 3-21. Ground truth of an anomaly locates on the upper left corner of virtual channel 20's ROI.




Figure 3-22. Due to the cross talk between virtual channels in the real world experiment, the simulated anomaly detection with measured PSF shows a significant low SNR, almost 5 times less than the ideal PSF case.

As mentioned earlier, one merit of using a FPA to detect an anomaly is that each detector “sees” a smaller ROI and thus an anomaly of a given size will have a bigger impact on detection as compared to the SPC. The next simulation confirms this expectation.




Figure 3-23. The detection signal of using different methods, SPC and FPA. The FPA virtual channel has a significant improvement on SNR as each detector collects signal from a smaller ROI and the anomaly will have a bigger impact on the measurement.




Figure 3-24. ROC and PR curve for SPC and FPA anomaly detection results. The FPA method elevates the performance under these metrics, especially for the PR curve.

A comparison between our previous SPC simulation result and this FPA implementation is made and illustrated in Figs. 3-23 and 3-24. Fig. 3-23 shows the detection signal increases significantly with our FPA strategy. Even with the non-ideal imaging conditions, the detection SNR is still better than the SPC result. ROC and PR curves in Fig. 3-24 are included to illustrate the significant improvement under different detection metrics, especially for the PR curve.

The FPA anomaly detection method also has the ability to handle scene with high noise level, the next simulation shows the method is able to detect anomalies at a SNR level of 5, but fail to do so at the level of 2, as shown in Fig. 3-25.




Figure 3-25. Simulation results to test when the anomaly detection using FPA fails under different noise level

Appendix 3-1 PSF-aware reconstruction algorithms for imaging:

I. INTRODUCTION

Suppose that we use an $f = f_m \times f_n$ detector element focal-plane array (FPA) to sense a scene x that consists of $N = m \times n$ pixels. Suppose that we further group the detector elements into $v = v_m \times v_n$ “virtual channels,” where each channel ideally views an $N_{\text{ch}} = m_{\text{ch}} \times n_{\text{ch}}$ local portion of the scene, with $m_{\text{ch}} = m/v_m$ and $n_{\text{ch}} = n/v_n$. For example, we may observe a 512×512 scene with a 64×64 FPA, and then map its detector elements into an 8×8 array of virtual channels, so that each local area is 64×64 .

Suppose that we modulate the scene with an $m \times n$ spatial light modulator (SLM), such as a digital micromirror device (DMD). Because of optical skew and other alignment issues the DMD does not always line up perfectly with the FPA, and so each virtual channel does not necessarily precisely consist of $f_m/v_m \times f_n/v_n$ detector elements. Further, non-negligible optical crosstalk occurs between detector elements from the individual mirrors on the DMD. We can represent the overall “point-spread” or crosstalk/blur transfer function of the optical system in a vectorized form as a $v \times N$ matrix Q :

$$Q = Q_{\text{virt}} \cdot Q_{\text{FPA}}.$$

Here, Q_{virt} is a $v \times f$ matrix that models how the FPA’s f detectors contribute to each of the v virtual channels, and Q_{FPA} is an $f \times N$ matrix that models how the FPA’s detectors respond to the the DMD’s N mirrors. Note that Q_{virt} and Q_{FPA} , and therefore Q , are extremely sparse.

In practice, we obtain Q_{virt} by rastering an $m_{\text{ch}} \times n_{\text{ch}}$ area across the DMD. There are v sets of these non-overlapping local regions; each region corresponds to a particular row of Q_{virt} , and the non-zero elements of that row corresponds to which of the FPA’s detectors received light from the rastered region. Similarly, we obtain Q_{FPA} by rastering individual mirrors, one at a time, through each $m_{\text{ch}} \times n_{\text{ch}}$ area of the DMD. In total, there are N different raster locations, and these correspond to the columns of Q_{FPA} . The rows of Q_{FPA} correspond to the f detectors of the FPA; each row will have non-zero elements indicating which of the DMD’s mirrors it receives light from.

II. THE SENSING STEP

In what follows, we model the physical setup in its 1-dimensional vectorized form. We modulate the scene x with a pattern h_k on the DMD. From here on it is assumed that the same N_{ch} -element pattern is displayed in each of the v local regions;¹ denote this smaller pattern as the column vector $h_{\text{ch},k}$. Thus, the full concatenated pattern can simply be expressed as

$$h_k = \mathbf{1}_v \otimes h_{\text{ch},k},$$

where $\mathbf{1}_v$ denotes the unit column vector of length v , and $\otimes$ denotes the (“right-side”) Kronecker product.

Let y_k denote a column vector consisting of the v measurements from the virtual channels sensed by the FPA when the scene x is modulated by pattern h_k :

$$y_k = \begin{bmatrix} y_k^0 \\ \vdots \\ y_k^{v-1} \end{bmatrix} = Q \cdot (h_k \odot x), \quad (1)$$

where the superscript of element y_k^j indicates the j th channel. Here $\odot$ denotes the *pointwise product*, which mathematically represents the process of modulation. It is essential to recognize that the vectorization of x is with respect to the v local regions mentioned previously. (Explain more of this.)

A. Modeling the full measurements process

Suppose that $h_{\text{ch},k}$ is the k th column of some matrix $H_{\text{ch}}^\dagger$, where † is the (conjugate) transpose operation. Further, suppose for simplicity that $H_{\text{ch}}^\dagger$ is $N_{\text{ch}} \times N_{\text{ch}}$, although this is not crucial. Then, since the pointwise product is commutative, it can be shown that after collecting N_{ch} sets of measurements $\{y_0, \dots, y_{N_{\text{ch}}-1}\}$, and *regrouping them into their respective channels*, we have

$$\begin{bmatrix} \begin{bmatrix} y_0^0 \\ \vdots \\ y_{N_{\text{ch}}-1}^0 \end{bmatrix} \\ \vdots \\ \begin{bmatrix} y_0^{v-1} \\ \vdots \\ y_{N_{\text{ch}}-1}^{v-1} \end{bmatrix} \end{bmatrix} = H x_{\text{PSF}} \quad (2)$$

where

$$H = I_v \otimes H_{\text{ch}},^2 \quad (3)$$

with I_v denoting the $v \times v$ identity matrix, and where the *PSF-affected image*

$$x_{\text{PSF}} := \text{vect}(Q \cdot \text{diag}(x) \cdot \Lambda) \quad (4)$$

is a length- N column vector. Here, $\text{vect}(\cdot)$ is the **row-majorized** vectorization operation, $\text{diag}(x)$ is a diagonal matrix with x on its main diagonal, and $\Lambda = 1_v \otimes I_{N_{\text{ch}}}$. The purpose of this formulation is that it separates the *full measurement process* into the operator H in (3), which is assumed to have a **fast implementation** (e.g., when H_{ch} is a Sylvester-type Hadamard matrix), and the PSF-affected image x_{PSF} in (4), which is solely a function of the scene of interest and the PSF of the optical system, both of which are assumed to be static during the observation period. It is interesting to note that (4) reduces to $x_{\text{PSF}} = x$ when Q is ideal, i.e., when $Q = I_v \otimes 1_{N_{\text{ch}}}^\dagger$.

III. COMPRESSIVE SENSING AND RECONSTRUCTION

In compressive sensing (CS), the goal is to take fewer measurements than some full set, which in this case is associated with N_{ch} possible patterns. Denote R_{ch} as the $M \times N_{\text{ch}}$ channel subsampling operator, with $M \leq N_{\text{ch}}$. It is a matrix primarily of zeros, but with a single “1” in each row: a “1” in entry (i, k) means that the i th measurement selects the k th row (i.e., pattern) from H_{ch} . Since we assume that DMD displays the same pattern for each channel, we can define the $vM \times N$ subsampling operator for all channels to simply be

$$R = I_v \otimes R_{\text{ch}}.$$

Assume that we use R to make vM observations of the scene x through v virtual channels that are distorted by some PSF function. Using (2) we can model this as

$$y = RHx_{\text{PSF}} + e \quad (5)$$

where e is a general additive noise term with vM elements.

Then, given y , R , and Q , we can recover an approximation of the observed scene by solving the *PSF-aware optimization problem*

$$\min_{\hat{x}} \|\hat{x}\|_{\text{TV}} + \frac{\lambda}{2} \|RH\hat{x}_{\text{PSF}} - y\|_2^2 \quad (6)$$

where $\hat{x}_{\text{PSF}}$ is the PSF-affected version of $\hat{x}$ according to (4). In words, (6) finds the candidate $\hat{x}$ with smallest TV-norm and whose PSF-affected transform coefficients best match the observations vector y .

Section 4. Anomaly Detection Demonstrations and Results

Sections 2 and 3 discuss the theory of the algorithms and mathematics we have developed for use in our single-pixel- and FPA cameras, as well as some of their implementations. In this section we describe in more detail the wide range of computer simulations and laboratory experiments we have conducted and interpret their results.

The key to designing a robust detection system is finding its optimal operating point. Toward this end, we conducted extensive experiments to discover the behavior of our system. Detecting an anomaly is a binary classification process. Therefore, in essence, these experiments were *calibration tests* to find the optimal operating threshold for different scenarios with which to make binary decisions. Finding the optimal operating threshold is the main theme of this section.

4.1 Simulation scenarios

We begin with a series of simulations of anomaly detection on a video set. The video frames consisted of two toy cars moving toward each other in the foreground. The background consisted of a static colored checkerboard and Newton's cradle, which contains swinging metal balls; this represented the "steady-state" motion of the background.

The simulations tried to match a real-world camera setup: the scene was modulated with binary DMD patterns. In the case of the SPC, the modulated frame was then summed to a single value; in the case of the multi-pixel camera, the modulated scene was first convolved with the measured PSF function and then the ROI corresponding to each virtual channel was summed to a single value. Examples of testing scene are shown in Fig. 4-1.




Figure 4-1. From left to right are three frames of scene images (256×256) showing two toy cars moving towards each other. In (b), a bright pixel-block of size 4×4 in the background represents an anomaly point.

The simulation parameters used are also matched to the experiment. For example, we assumed that the DMD was running at a speed of 15 kHz. We used frames from a video sequence with an assumed rate of 30 frames per second. Thus each frame of the scene is modulated and measured $15000/30 = 500$ times before the next frame. Hence for a 128×128 resolution scene, the compression ratio will be 3%, and for a 256×256 resolution the ratio is 0.7% which represents a high compression ratio.

4.2 Statistical inference

In the simulation, we repeatedly modulated the scene with patterns designed for anomaly detection and formed the data by summing up the post-modulated image. When the anomaly was present, the measurements showed different statistics as compared with the case without anomaly; this was quantified, in part, by the MMD part of our change detection algorithm. Given the acquired time series measurements $\{b_j\}$ from the selected modulation patterns, the metrics were calculated through the following steps:

1. We calculate the change, $z_j = \sum_{k=1}^M |b_{j+M+k} - b_{j+k}|$, where M is the number of patterns in this selection. This purely records the difference of measurements between the current set of patterns and the previous set of patterns.
2. Then we use a window of size W to slide through the change series with one step at a time. For each window, we calculate the standard score (Z-score) from the two nearest non-overlapping window. The ratio of the mean value of the difference between those two windows and the smaller variance of the changes within two windows forms the Z-value.
3. Based on the changes, we calculate the maximum mean discrepancy (MMD), which is used to identify whether a data are from the same or different probability distributions.
4. The product of Z-score and MMD score generates our change detection score.
5. The change detection scores are also a time series of data points. At every given moment of time we say that an anomaly has occurred if the change detection score is above a predetermined threshold.

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

We made an interesting observation of the time series of change detection scores: There is always a secondary negative spike that follows an anomaly signal spike. This phenomenon is due to the sliding window that is part of the calculation process: we compare a section of data of length W to the same length of data that are k steps before it. Therefore, a negative spike always occurs W points after true signal spike. This can be used as extra information for us to determine whether the signal is a true positive or not by looking for the secondary spike and pair spikes.

4.3 Simulation with different detection matrix

There are many types of measurement patterns that can be used in conjunction with the DMD to detect anomalies. We studied patterns from the following three types of matrices and compared their detection capabilities:

- Sum-To-One transform (STOne),
- Permuted Walsh-Hadamard transform (WH),
- Partial-Complete Hadamard transform (PC).

Each transform matrix possesses unique properties. The permuted WH matrix contains columns that have been randomly scrambled, which destroys its strong sequency structure; the result is that the pattern contained in each row has been reduced to a pseudo-random binary noise waveform. Further, the rows are randomly chosen so that there will be minimum coherence between modulation patterns.

In terms of the STOne matrix, we conducted a mathematical study that showed how it is actually equivalent to a Sylvester- or Walsh-Hadamard matrix. Yet at the same time, due to a deterministic sequence of permutations and inversions of its rows and columns, the STOne matrix has the appearance of being random, or incoherent with respect to the bases that sparsify images. The STOne matrix is also able to generate low-resolution previews instantly from very few measurements, which can be of benefit in compressive image and video reconstruction.

However, both the permuted WH and STOne patterns “democratically” modulate an observed scene and do not possess any bias or preference on the prior information of an anomaly. That is, their measurements have near-constant magnitude with little deviation. Fig. 4-2 shows the magnitude of change detection plots from a simulation using both WH and STOne patterns on the DMD. The spikes in the change plot indicate an anomaly appearing or disappearing. It is clear that their change detection performance is virtually identical.




Figure 4-2. Anomaly detection results using the change plot for both WH and STOne, the spikes in the plots indicate the appearance and disappearance of the anomaly. Note, that these graphs show the absolute value of change, which rectifies the negative spikes to be positive.

Due to the sliding window in calculating the change value, each spike will be accompanied by a secondary spike delayed by the length of the detection window size W . In most case the secondary spike is negative. The magnitude of the spikes can be also be used to indicate the effectiveness of the detection respect to the anomaly intensity to background illumination ratio: as the intensity ratio increases, intensities of the spikes become larger. When the anomaly becomes dimmer, at some point it will not be able to be detected by the algorithm and the number of false positive spikes will be dominate. We found that the detection capability of WH and STOne to be very similar as illustrated in Fig. 4-2. Further, they sometimes had poor performance when the signal-to-background illumination ratio (SBR) was low.

The SBR is a metric we developed to quantify how bright an anomaly is relative to its background:

$$SBR = \frac{\sum \text{pixels in region with anomaly}}{\sum \text{pixels in same region without anomaly}} \quad (4-1)$$

Now we examine the Partial Complete algorithm and analyze groups of signature blocks in the Partial-Complete Hadamard spectrum. For a scene modulated by DMD with N pixels, we split the Hadamard domain into $F = 4^K$ non-overlapping blocks, with $B = N/F$ measurements per block. An important dual feature is that the pixel domain is partitioned at the same time into B non-overlapping tiles, each of size $2^K \times 2^K$ pixels. In this case the test scene has pixels of $N = 256 \times 256 = 2^{16}$ pixels and $F = 64$ blocks (labeled as $B_0, \dots, B_{63}$)¹², which results in $B = 2^{10} =$

¹² Similar to STOne, the Partial-Complete schema can also generate coarse previews based on a limited set of measurements. For example, in the case of $F = 64$ a coarse preview with 8×8 tiles can be constructed by obtaining the measurements of block B_0 . Finer resolution previews can also be generated by gathering measurements from blocks B_4, B_{32}, B_{36} , etc.

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

1024 measurements per block in the Hadamard domain; in the pixel domain this results in $B = 2^{10}$ tiles partitioning the FOV, where each tile contains 8×8 pixels.

In previous studies we showed that the Hadamard spectrum of an anomaly appearing in an arbitrary location within an 8×8 tile of pixels was essentially the same regardless of which tile in the FOV. That is, there is a strong periodicity induced in the pixel domain by the structure of the Hadamard waveforms as a result of their Kronecker product structure. Thus it suffices to examine any 8×8 tile.

We noticed that certain groups of four blocks B_k in the Hadamard domain were complementary in the sense that they covered up nulls in the measurements corresponding to specific locations within an 8×8 tile of pixels. Nulls in the measurements are undesirable since an anomaly corresponding to that location could result in it not being detected, i.e., a false negative.

It is helpful to view the energy of the PC Hadamard spectrum with and without an anomaly. For example in Fig. 4-3, 64 blocks of coefficients are plotted, where it is clear that some blocks show a bigger response when an anomaly is present. If one chooses blocks that have small response to an anomaly, it will probably be missed.




Figure 4-3. Examples of spectrum when varying the location of the anomaly, with x-axis being the block number and y-axis being the coefficients. The anomaly locations are also shown in the top of each graph.

This motivated us to investigate which blocks fully covered, or spanned, the possible anomaly locations within an arbitrary partitioned 8×8 patch in the FOV. Our analysis discovered several sets of complementary groups-of-four signature blocks, which are shown in Fig. 2-8 of Section 2. The patterns from these block sets can maximize the anomaly response, providing better detection results.

For example, let us consider the group Set 2 consisting of blocks $B_{36}, B_{38}, B_{52}, B_{54}$. The signature of block B_{36} shows a checkerboard pattern composed of solid 4×4 square areas of the same size as our test anomaly. The signatures for blocks B_{38}, B_{52}, B_{54} are simple 2-pixel *cyclic shifts* of the signature for B_{36} in the horizontal and vertical direction. We observed that the nulls locations for the patterns of block B_{36} are fully covered by the union of the patterns in blocks B_{38}, B_{52}, B_{54} . To better visualize this, we simulated a 4×4 anomaly in different locations and then averaged the energy of the measurements from each block. Fig. 4-4(a) shows the observed scene, where the red dotted rectangle indicates the locations that we rastered the anomaly through. Figs. 4-4(b)–

(e) show how the average measurement energy for blocks $B_{36}, B_{38}, B_{52}, B_{54}$ change as a function of the location of the anomaly. The dark areas are nodes, or minima, that are undesirable. However, comparing the subfigures, we see that for every minima, there exists a maxima in that same location in a different block. To be sure that we cover all locations we averaged Figs. 4-4(b)–(e). The result is shown in Fig. 4-4(f), where the blue nodes no longer exist.




Figure 4-4. Examples of nodal patterns (darkest regions) as a function of anomaly location indicated by the red dashed region of the scene in (a) for different signature blocks. (b) Block B_{36} ; (c) Block B_{38} ; (d) Block B_{52} ; (e) Block B_{54} ; (f) The average of blocks $B_{36}, B_{38}, B_{52}, B_{54}$ shows the nodal locations have been covered.

Next we wanted to see how the change detection algorithm performed when using PC measurements as compared to permuted WH measurements. In the following simulation, we fixed a 4×4 anomaly at an arbitrary location in a 256×256 scene, and then compared the change detection algorithm using measurements from PC signature blocks $B_{36}, B_{38}, B_{52}, B_{54}$ versus measurements from randomly selected permuted WH patterns. Fig. 4-5 shows that the chosen PC signature block sets result in a significant improvement over permuted WH in detecting a 4×4 anomaly at this arbitrary location. Combining this result with the previous simulation, suggests that we can safely generalize this to conclude that this should be true for any location. Fig. 4-5 also shows the detection performance under different SNR conditions. This demonstrates that the proposed approach has a better chance to detect a low-intensity anomaly.




Figure 4-5. The change detection using permuted WH and PC block sets with different noise add to the simulation frames. In all scenarios the PC approach outperforms permuted WH.

4.4 ROC and PR curve

Evaluating the detection performance by merely looking at the change graph is insufficient and possibly misleading. As such, we decided to use the well known *receiver operating characteristic* (ROC) and *precision-recall* (PR) graphs as our comparison criteria. These metrics gave us the ability make more quantitative comparisons between the different methods. They were defined in Eqs. (2-8)–(2-10).

As discussed in Section 2, PR is a better metric when dealing with strongly skewed positive and negative events such as anomalies. We can plot a *PR curve*, which is a collection of operating points, each with an associated threshold to make a binary decision. However, determining the best operating point is a difficult decision since it is not always clear how to balance the variables of precision and recall.

There exist many ways to determine an optimal operating point such as area under the curve (AUC), the F-score and the G-measure. The AUC simply finds the area of the rectangle from the origin to a particular point on the PR curve. The rectangle with the largest area associated with a

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

particular point on the PR curve maximizes the combined contribution of the precision and recall. Similar to this is the G-measure, which is the geometric mean of precision and recall. The F-score is harmonic mean of precision and recall, and is more conservative than the AUC or the G-score. We decided to use a more general form: the F_β -score is just the β -weighted-harmonic mean of precision and recall:

$$F_\beta = (1 + \beta^2) \frac{\text{precision} * \text{recall}}{\beta^2 * \text{precision} + \text{recall}}. \quad (4-2)$$

A larger value of β puts more importance on recall, while a smaller value deems precision as more essential. We evaluated our system for the common values of $\beta = 0.5, 1, 2$. Note that the regular F-score occurs when $\beta = 1$, where the precision and recall are evenly weighted.

4.5 Anomaly detection experiment with SPC

With the simulation results showing the advantage of using PC signature block sets as detection patterns we then constructed a physical tests. A big merit of the anomaly detection system is that it can be implemented with the original SPC setup without significant modification. This also allows the anomaly detection function to be realized on the InView SPC camera devices that are compact and portable with all the optics and analog to digital conversion (ADC) electronics integrated as a whole system.

The InView camera uses a Texas Instrument DMD chip (DLP7000) with a resolution of 1024×768 to modulate the scene. The DMD has a maximum operation speed of 18 kilohertz and has a wide spectrum range of reflection, from 400 nm to near Infrared (NIR) of 2000 nm. The micromirrors are 13.6 μm on the diagonal and rotate on an axis to two angles. The DMD is put at the location where the image plane of the scene is, and a total reflection (TIR) prism is used to channel one direction of the reflected light into a path where other optics elements are set to project the scene into light collection devices, in this case, a single detector. The scene we tested was a toy car being dragged through a printed background with a constant speed; the background is a gray scale image of a parking lot and buildings. The photon collection device was a Thorlab (PDA36A) silicon detector with a maximum readout speed of 5 kilohertz at 70 dB noise reduction option, and an effective detecting spectrum range from 400 nm to 1000 nm (Fig. 4-6), which covers the whole visible wavelength.




Figure 4-6. Spectrum response of Thorlab PDA36A single detector

We used a red laser pointer as the anomaly, which has a wavelength of 700 nm. And has a maximum intensity of 5mW/mm^2 . The scene with and without the anomaly is shown in Fig. 4-7. The intensity curve in Fig 4-7(c) is plotted from the image intensity values along the line drawn across the scene in Fig 4-7(b). The intensity curve shows anomaly intensity respect to the background and foreground illumination intensity.




Figure 4-7. Ground truth of the testing scene with (b) and without (a) the anomaly. (c) shows intensity profile across the line in (b) that pass through the anomaly.

4.5.1 Laser triggered anomaly control system

In order to obtain the ROC and PR curves we need to synchronize the measurements with the ground truth of the anomaly. We also developed a laser triggering system that verifies the synchronization between the laser trigger ground truth and the anomaly measurement data.




Figure 4-8. 2-channel Relay set used to precisely control the appearance of a laser anomaly in anomaly detection experiments.

The trigger system is realized using a two-channel relay set (SunFounder 2 Channel 5V Relay) like the one shown in Fig. 4-8. A relay is an electrically operated switch and is used where it is necessary to control a circuit by a low-power signal (with complete electrical isolation between control and controlled circuits such as TTL or ADC voltage output), or where several circuits must be controlled by one signal. In our experiment, the relay set has two channels: one controls the laser as anomaly and the other controls the motor that pulls the car moving across the scene. Fig. 4-9 shows the measurement data and the laser trigger ground truth obtained from the National Instrument analog and digital conversion device. We can see that the trigger edges of the laser match well with the light intensity increase in the measurement data.




Figure 4-9. A comparison between the laser trigger ground truth (red) and the measurement data (blue), the DMD is running random Walsh-Hadamard pattern at 1 kHz.

The relay system allows us to control the anomaly event precisely, providing confidence that consistency with measurements can be achieved even with a small SNR. Fig. 4-10 displays a

repetition of two experiments with the exactly same parameters, and the measured triggers shows reliable consistency.




Figure 4-10. Measurements of two laser trigger signals with the same parameters.

With the functional laser trigger system, we are able to run the anomaly detection experiment with known trigger edge and repeatable experiment conditions, which allows us to generate the ROC and PR curves.

4.5.2 Experiment results and discussion

As discussed previously, we have found that using the PC Hadamard patterns with particular signatures will provide higher SNR in detecting anomalies. In this experiment we use the signature block combination of B_{36} , B_{38} , B_{52} , B_{54} , with 100% patterns in each block, so the total number of patterns is 4096. The DMD is running at 2000 frames per second and the patterns are repeated 10 times. The toy car is also moving across the scene, as shown in Fig. 4-11, where the anomaly size is around 4×4 in a 256×256 scene.




Figure 4-11. An example image of the scene, the anomaly is zoomed in and the toy car is moving across the scene during the experiment.

The laser anomaly is turned on and off three times in each measurement. The experiment is run with the anomaly in 10 different locations to ensure that it will be covered by all four signatures. Fig. 4-12 shows the data from the fifth trial along with the laser trigger curve.




Figure 4-12. Raw measurement data of anomaly location #5, the anomaly appeared three times (red), which can also be seen from the intensity increase in the measurement (blue).

Fig. 4-13 shows the change detection results of anomaly locations #5 and #8. As usual, the change detection algorithm is the product of the Z-score and MMD. The red spikes indicate when the laser anomaly was actually turned on or off; these are the positive events to be detected. The blue spikes indicate high scores output from the change detection algorithm, i.e., what it has calculated as a positive event. When a blue spike coincides with a red spike we register a true positive, and when it does not we register a false positive.




Figure 4-13. the change detection result of two different anomaly location (location 5 and 8). Each detection has the anomaly turned on and off three times (red). The change detection is combination of Z-score and MMD.

With the laser trigger ground truth we obtained confusion matrices using the `perfcurve` function in MATLAB at different (but common) thresholds for the different 10 change detection trials. For each threshold we then computed an average the confusion matrix for the 10 different anomaly locations. The resulting average ROC and PR curves are shown in Fig. 4-14.




Figure 4-14. Average ROC curve and associated PR curve.

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

We are now able to compare the anomaly detection performance between the PC method and WH method through ROC and PR curves. Fig. 4-15 shows the detection result of the two methods under different noise level. As expected the PC method outperforms the WH approach in each case.




Figure 4-15. ROC and PR curve plots for PC and WHT patterns under different SNR circumstances, the SNR is labeled at the top of each graph.

4.6 Anomaly detection with FPA

The previous experiments proved that the SPC system is feasible for high-speed anomaly detection. However it has some limitations too. The single detector collects all the light reflected from the DMD and it can be difficult to detect an anomaly that is very small or weak.. Even with our proposed PC patterns that are tuned to size of an anomaly, the SPC system's performance suffers when the anomaly become too dim or the background illumination increases. Another disadvantage of the SPC anomaly detection system is its incapability of detecting the *location* of the anomaly in real time since the measurements are global, i.e., they record the full FOV.

The measured data can be used to reconstruct the scene for the purpose of locating the anomaly, yet this process is computationally demanding and the reconstruction may take seconds for a high-resolution scene. In order to improve the anomaly detection system and overcome the limitations of SPC imager, we replaced the single detector with a Hamamatsu focal plane array (FPA) that consists of 64×64 Indium gallium arsenide (InGaAs) detectors. The FPA has a spectrum respond region of 700 nm to 2000 nm, with a maximum data readout speed of 1000 hertz. Fig. 4-16 shows some sample images taken by the FPA.

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).




Figure 4-16. Image captured by the Hamamatsu 64×64 Indium gallium arsenide (InGaAs) FPA. Note, there exist a few defective pixels at the bottom of the FPA.

The FPA is placed at the image plane of the DMD, and the optical system maps a 512×512 DMD region to the 64×64 pixel array. In the ideal situation we would like to use the data collected by the total 4096 detectors, where each detector on the FPA collects light from a block of 8×8 DMD mirrors. Smaller regions of interest (ROIs) will significantly increase the SNR of the anomaly compared to the SPC system where the single detector ‘sees’ the whole DMD area.

4.6.1 Virtual channel

As described in Section 2 we designed and created so-called “virtual channels” from the FPA pixels. There are many possible sizes for these channels, and they do not need to be square. However, the channel size must correspond to a portion of the DMD that can support a Hadamard pattern. We decided on ROIs that contained 64×64 mirrors—in part because it also has nice imaging capabilities—but also reported on 32×32, 16×16, and 8×8 ROIs in previous progress reports (note, factors of 12 are also permissible since a Hadamard matrix of that size exists).

Smaller ROIs means fewer measurements are necessary. To be consistent with our simulation process we need at least a DMD ROI resolution of 64×64 for each detector in the 64×64 Hamamatsu FPA, so a minimum of 4096×4096 resolution DMD is required to reproduce the results in the simulation which is unavailable with the hardware we have. With a 64×64 ROI, the FPA views a FOV with an array of $8 \times 8 = 64$ virtual channels.

The virtual channel strategy has another benefit of fast processing speed. As currently we have limited computation power and data transmission bandwidth, it will take tremendous time to process the data from 4096 detectors without the help of parallel computing. An area of future research is to use all the detectors from FPA as their own channel. A comparison of data collected using the SPC and one particular FPA virtual channel (channel 20) is illustrated in Fig. 4-17. The corresponding ROC and PR curves are shown in Fig. 4-18 to show the improvement of the FPA system. Note the significant improvement in the PR curve.




Figure 4-17. The detection signal of using different methods, SPC and FPA. The FPA virtual channel have a significant improvement on SNR as each detector collects signal from a smaller ROI and the anomaly will have a bigger impact on the measurement.




Figure 4-18. ROC and PR curve for SPC and FPA anomaly detection results. The FPA method improves the performance under these metrics, especially for the PR curve.

4.6.2 Point spread function correction to the imaging system

As discussed in Section 3 we also measured the point-spread function (PSF) of our optical system and used it to correct the distortion of the image on the FPA. The PSF describes the response of an imaging system to a point source or point object. We were able to use the information contained in the measured PSF to correct for errors in the recovered image generated by our reconstruction algorithm. However, using the PSF to remove or “un-mix” the distortion directly from the compressive measurements is not straightforward or entirely obvious, and adds a significant computational burden that is undesirable. This is a good area of future research.

4.6.3 Anomaly detection results with FPA and discussion

We now discuss the final performance of our anomaly detection system. Recall that the goal was to find an ideal operating point for a given observed scene. We performed a series of laboratory experiments to discover the behavior of our system. The same steps would be done in whatever environment the camera is operating in. The overall flow of data was

- Take compressive PC measurements of the scene with and without anomalies
- Feed them into the detection algorithm
- Generate confusion matrices and PR curves
- Find ideal operating point and its associated threshold using the F_β -score in Eq. (4-2) for a user-defined value of β .

We considered several testing situations that consisted of different anomaly intensities and durations. Specifically, we considered anomaly intensities of 980, 480, 280 μW , and durations of 150, 60, 25 ms. Therefore, there were nine (9) different cases to judge. For each case we performed the anomaly detection process in ten (10) arbitrary locations within the ROI of one particular virtual channel and generated average ROC and PR curves as described in Section 4.5.

We used the F_β -score to find the optimal operating point for each case. Fig. 4-19 shows the summary of the calculated maximum F_β -score for each of the 9 cases for values of $\beta = 0.5, 1, 2$. Fig. 4-20 shows the associated threshold for each of these operating points. We see from the plots in Fig. 4-19 that the maximum F_β -score decreases as the anomaly intensity decreases. This makes sense since our precision and recall are directly tied to the SNR of the anomaly. It also makes sense that the associated thresholds also decrease with anomaly intensity since resulting detection scores will be smaller for relatively weaker anomalies.

It is interesting to note that the detection results summarized by the maximum F_β -scores in Fig. 4-19 were not significantly affected by the anomaly’s duration. For each intensity, the shortest duration (25 ms) was 6-fold briefer than the longest (150 ms), yet the maximum F_β -scores did not change very much. This may be an indication of robustness of our method. A future area of research is to investigate how system to push our system to a breaking point. Doing so can free

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

Page 67 of 74

up other factors and constraints in our system that we can trade for the ability to detect very brief events.

In terms of the effect of the β weight, comparing thresholds for $\beta = 1$ in Fig. 4-20, we see resulting lower thresholds for $\beta = 2$. A lower threshold means more events will pass the binary test at the expense of precision. In the other direction, for $\beta = 0.5$, we see higher thresholds, which results in more conservative detection of the anomalous events.




Figure 4-19. Maximum F_{β} -score calculated for $\beta = 0.5, 1, 2$ from the PR curve results. Each color bar corresponds to an anomaly of different intensity and duration indicated by the x and y axes.




Figure 4-20. Detection thresholds determined by the maximum F_{β} -score in Fig. 4-19.


Given that we have determined the ideal operating point and threshold from the F_{β} -score, we are interested to see how well it performs for anomalies that are not necessarily in the 9 cases used for calibration. We tested the camera's detection ability for anomalies with arbitrary intensities and durations. Fig. 4-21 shows the raw compressive data stream and the change detection scores

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

for 9 different test cases. The change detection graphs also contain a red line which is a threshold gleaned from the calibration steps of Fig. 4-20. Instead of the anomaly's absolute intensity, we measured its SBR (see Eq. (4-1) above). This is a relative metric that can be helpful to extend the results to different operating scenarios.

In general, we discovered that the detection algorithm usually correctly classified an anomaly with different intensities and durations as long as the SBR was above 1.02. Further for the given thresholds most of the false positives were avoided. However, as the anomaly's intensity got weaker and duration briefer we found the detection change contained mostly noise. Therefore, no matter what threshold we chose either the false positive rate was 100% or we had a true positive rate or 0%. However this is not due to the incapability of the detection algorithm or threshold determination method. Rather it is an SNR problem that defines the limit of our system. Yet there lies the potential to improve the detection performance by using even smaller sized virtual channels or possibly even the native resolution of the 64×64 FPA. These will increase the SNR of a given anomaly compared to larger sized virtual channels.



Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

Page 69 of 74

ARMY STTR-A2-5466 Rapid anomaly detection and tracking via compressive time-spectra measurement

Figure 4-21. Test of arbitrary anomaly detection using pre-determined thresholds. Each row corresponds to an anomaly of a particular intensity but different durations. Brighter anomalies are at the top, and longer duration anomalies are on the left.

Appendix 4-1: Improving the PR curve with a ground truth signal broadening technique

We noticed that many of our average PR curves, such as in Fig A-1, had a surprisingly low precision rate, which gives the impression that the detection algorithm is not very effective. However the real data indicates otherwise, and so we looked closer at the detection results in comparison with the ground truth.




Figure A-1. Average ROC and PR curves before the improvement.

We found that the change detection scores always contained an ascending and descending slope when the anomaly is detected; on the other hand, the *ground truth* of the anomaly is more like a delta spike with very large slope, as shown in Fig. A-2. The difference between the real data and the ground truth leads to the low precision rate in the PR curve.




Figure A-2. The comparison between the real data and the ground truth. The blue line is the change detection result, and the red line indicates the ground truth anomaly event.

Our approach to improve the PR curve is to attempt to match the change detection score data with the ground truth. We used a data broadening method to make the indication of the anomaly in the ground truth rise slower as seen in Fig. A-3. In essence, we have slowed down its “slew-rate” so that it is more comparable to the real data. Compared with Fig. A-1, it is clear from Fig. A-4 that this broadening technique improves the PR results significantly. This then permits us to calculate better the F_β -scores for locating the optimal threshold.




Figure A-3. The broadened ground truth that matches the detection data.




Figure A-4. The ROC and PR curves after the improvement, where we see the precision raises to above 50%.

Section 5. Conclusion and future directions

During this project we expanded on the utility and efficiency of the original short-wave infrared single-pixel compressive imaging system by converting that camera architecture into a high-speed anomaly detection system. This enhanced capability is an improvement as compared to performing the same task with a traditional infrared focal plane array. It is also an important demonstration in the capability of compressive optical systems as sensor platforms for machine vision applications. By operating directly on the measured coefficients, a great deal of time and computation is saved. We successfully demonstrated the anomaly detection capability with a single sensor and further increased the speed and efficiency by moving to a 64×64 multi-pixel architecture. We also correspondingly designed a set of modulation patterns to be displayed on the DMD that allow the algorithms to be exploited across multiple length scales by enabling individual or joint operation of pixels in this task.

While this project demonstrated promising results, there are still other areas that can be pursued for future research and development. We discuss several of these next.

1. As mentioned in Section 4, we noticed that a secondary negative spike always follows the true signal spike that is associated with an anomaly. This secondary negative spike is always delayed by W data points, where W is the size of the sliding window. Therefore, a more sophisticated algorithm can use both the negativity and delay of the secondary spike to strengthen its ability to detect anomalies. Further, since most anomalies are brief, we can possibly use the set of detection spikes that occur when the anomaly disappears.
2. We briefly studied the crosstalk between FPA detectors or virtual channels both simulations and laboratory experiments. Crosstalk becomes an important issue when an anomaly is on the edge of two or more detector elements or virtual channels. This cause the signal of one particular region to decrease as the energy is spread into other regions, and detection performance may suffer from this. By monitoring the size and timing of signal changes in the adjacent pixels of virtual channels, we can use the crosstalk information to boost our detection confidence and also more accurately localize the anomaly. Information from the measured PSF of the optical system should be able to help with this problem.
3. Another area is to incorporate the Local-Global measuring and grooming technique into the proposed detection system. Recall that this is a two-step process: the first step is optical (see Eq (2-6)), and the second is computational (see Eq. (2-7)). As explained in Section 2, we use can obtain local measurements from the individual detectors in a FPA, and then directly groom (i.e., multiplex) them together into larger groups that correspond to global measurements in the compressive domain. This will permit us to compare the statistics of compressive domain data over different ROIs and larger areas. However, the theory presented in Section 2 does not take into account the real-world effect of the PSF function. That is, the optical crosstalk experienced by our FPA measurement system corrupts or blurs the local measurements obtained in Eq. (2-6). As a result, the global measurements that result from the grooming step of Eq. (2-7) propagate these errors. A

Approved for public release; distribution unlimited.

Report developed under Topic #A2-5466, contract W911NF-14-C-0006 DATA RIGHTS: IAW DFARS 252.227-7018(f).

valuable avenue of future research is to incorporate the PSF function into the Local-Global technique to determine its effect and, if possible, how to ameliorate it.

4. There still remains a great deal of hardware expandability with this new short-wave infrared detection system. This is possible since the DMD is capable of steering the modulated light in two complementary left and right directions from the micromirror surface normal. Throughout this project we employed our detection on only one of the two optical branches. As such, our high-speed detection system can also be expanded into simultaneous detection and imaging by adding an FPA to this second branch. Another alternative is to add a spectrometer to the second branch allowing simultaneous spatial and spectral anomaly detection. Finally it is worth mentioning that the micromirrors are broadband optical modulators expanding their potential out beyond 10 microns in wavelength with an appropriate window such as ZnSe on the DMD chip. By adding a mid-wave infrared detector to the second optical branch, this system would be capable of performing dual-band high-speed anomaly detection across a large portion of the spectrum further increasing its capability.