A Thin-Film, Large-Area Sensing and Compression System for Image Detection

Tiffany Moy, *Student Member, IEEE*, Warren Rieutort-Louis, *Member, IEEE*, Sigurd Wagner, *Fellow, IEEE*, James C. Sturm, *Fellow, IEEE*, and Naveen Verma, *Member, IEEE*

***Abstract*—This paper presents a sensing and compression system for image detection, based on large-area electronics (LAE). LAE allows us to create expansive, yet highly-dense arrays of sensors, enabling integration of millions of pixels. However, the thin-film transistors (TFTs) available in LAE have low performance and high variability, requiring the sensor data to be fed to CMOS ICs for processing. This results in a large number of interconnections, which raises system cost, and limits system scalability and robustness. To overcome this, the presented system employs random projection, a method from statistical signal processing, to compress the pixel data from a large array of image sensors in the LAE domain using TFTs. Random projection preserves the information required for subsequent classification, and, as we show, is highly tolerant to device-level variabilities and amenable to parallelized implementation. The system integrates an amorphous-silicon (a-Si) TFT compression circuit with an array of a-Si photoconductors, representing an 80 $\times$ 80 active matrix, performing up to 80$\times$ compression of the 80 signal interfaces. For demonstration, image classification of handwritten digits from the MNIST database is performed, achieving average error rates of 2–25% for 8–80$\times$ compression (e.g., 7% at 20$\times$ compression).**

***Index Terms*—Amorphous silicon, compression, image classification, thin film sensors, thin film transistors, variability.**

I. Introduction

LARGE-AREA electronics (LAE) is based on processing semiconductor thin-films at low temperatures, making it compatible with a wide range of materials. This enables diverse types of sensors [1]–[8], formed on a variety of substrates, such as glass or plastic. These substrates can be physically large (on the order of square meters) and conformal, enabling the deployment of a large number of spatially-distributed sensors. Thus, LAE is a compelling technology for embedded sensing on a large scale.

However, processing and analysis over the large amount of sensor data requires complex functions, which are not easily implemented in LAE. This is because the characteristic low-temperature processing of LAE results in active thin-film devices, such as thin-film transistors (TFTs), that suffer from low performance and high variability. Instead, CMOS ICs are much better suited to perform such tasks, offering transistors exhibiting orders-of-magnitude better performance and greater reliability. This motivates the need for hybrid LAE-CMOS systems [9], which leverage the benefits of both technologies. However, research in hybrid systems design [8], [10], [11] shows that the key challenge that emerges in adopting such an architecture, is the need for a large number of interfaces between LAE and CMOS. This limits scalability in the number of sensors and thus severely restricts the potential of hybrid systems.

One method used to address this interfacing challenge is the active-matrix design, as employed in flat panel X-ray imagers [12]. This results in an approximately square-root reduction in the number of interfaces. However, as an example, taking X-ray imagers of today, which are approaching tens of megapixels, we still require thousands of interfaces. This poses a dominating limitation in systems, impeding further scaling of such systems in the future. In this paper, we present an approach that substantially reduces the number of interfaces beyond the level achieved with an active matrix, by performing image compression via an approach called random projection. As we show, random projection is highly tolerant to variations and can achieve fast operation despite low-speed devices, thus making it highly suitable for TFT implementation.

The major contributions of this work are as follows:

1) We present the concept, as well as circuit and architecture designs, of a system that performs compression via random projection in the LAE domain using low-speed, variation-prone TFTs, for subsequent classification.
2) We develop a prototype and demonstrate the system experimentally, to evaluate the feasibility and performance of the system.
3) We analyze the impact of various non-idealities in the TFT implementation (i.e., device variations, nonlinearities) on system performance. Evaluating the potential of the random-projection approach will enable future device-level optimizations and relaxations.

The remainder of this paper is organized as follows. In Section II, we provide an overview of the system and the key principles involved, specifically random projections and their property of preserving the inner product between vectors. In Section III, we describe the thin-film implementation of the system, including the TFT circuits which make up the compression block. In Section IV, we examine the effects of

Manuscript received May 4, 2016; revised July 21, 2016; accepted August 9, 2016. Date of publication September 30, 2016; date of current version October 25, 2016. This work was supported by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA, and NSF (grants ECCS-1202168 and CCF-1218206). This paper was recommended by Associate Editor T. Serrano-Gotarredona.

The authors are with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: tmoy@princeton.edu; w@rieutortlouis.com; wagner@princeton.edu; sturm@princeton.edu; nverma@princeton.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCSI.2016.2600498

1549-8328 © 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

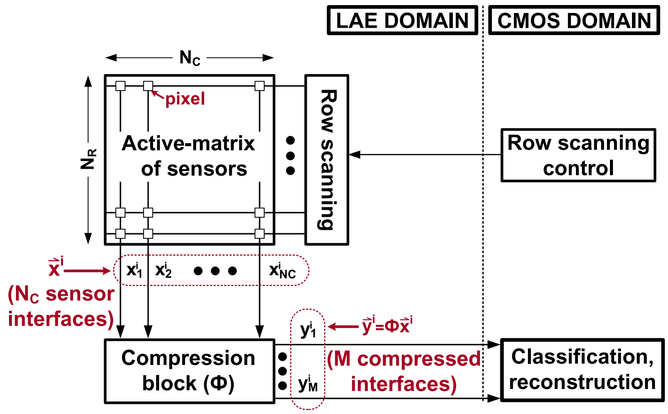

Fig. 1. Proposed architecture of image sensing and compression system. A reduction in interfaces from an active-matrix array of sensors is achieved using a TFT-based compression matrix Φ (i.e., N_C to M).

TFT non-idealities in the compression block, by modeling TFT nonlinearity and variation to analyze precisely how these result in deviations from an ideal compression matrix. In Section V, we demonstrate the performance of the image-compression system through classification and reconstruction on images from the MNIST database of handwritten digits [13]. Finally, in Section VI, we provide conclusions.

II. Overview of the Approach

Fig. 1 shows the block diagram of the proposed system. In an $N_R \times N_C$ active-matrix array of sensors, scanning row-by-row reduces the sensor interfaces from $N_R \times N_C$ to $\sim N_C$ (plus a few row-scanning control signals). The N_C-interface signal, designated as the vector $\vec{x}^i$, is then fed into a compression block. Compression is achieved by performing multiplication of $\vec{x}^i$ with an $M \times N_C$ matrix Φ, where $M < N_C$. Thus, the N_C-interface signal is now reduced to an M-interface signal $\vec{y}^i$, resulting in a compression factor of N_C/M. Transmitted to the CMOS domain, the compressed $\vec{y}^i$, can then be used to reconstruct the image; though, in this work our primary interest is in applications requiring classification of the sensed image, directly from the M-interface output signal $\vec{y}^i$ without reconstruction.

Previously, we presented a system [8] that directly applies TFT-implemented machine-learning classifiers to the sensor data. However, that embedded-classifier approach requires additional, specialized circuitry to program and store analog voltages in the TFT classifiers. In contrast, the system demonstrated in this paper, which performs feature extraction, rather than classification, in the LAE domain, needs only very simple, variation-tolerant TFT circuits with no programming. Furthermore, when applied in conjunction with an active matrix, this random-projection-based compression approach achieves greater reduction in the number of interfaces.

A. Random Projections

Many image compression algorithms utilize a transform domain, where the image information is known to be sparse (i.e., has a small number of non-zero transform coefficients). For example in JPEG, the two-dimensional (2-D) discrete-cosine transform (DCT) is used [14]. However, domain transformations, such as DCT, are generally too complex to compute using TFT circuits. Instead, the proposed system performs compression using TFT-based random projection.

Using the approach of random projections, the compressed output $\vec{y}^i$ (of length M) is generated by taking M linear measurements of the input signal (of length N_C). That is, $\vec{y}^i$ is derived from linear combinations corresponding to multiplication of the signal $\vec{x}^i$ with an $M \times N_C$ matrix, Φ, such that $\vec{y}^i = \Phi \times \vec{x}^i$. While in general the original signal vector $\vec{x}^i$ cannot be reconstructed, theoretical work shows that for certain Φ, the inner product between two compressed output vectors $\vec{y}^j, \vec{y}^k$ statistically preserves the inner product between the two corresponding original vectors $\vec{x}^j, \vec{x}^k$. For instance, a relevant mathematical result that provides bounds on the inner-product error for a specified set of vectors is the Johnson-Lindenstrauss Lemma [15]. However, more generally, as we describe in the next subsection, a matrix Φ can be chosen that yields some level of statistical inner-product preservation for all vectors.

Inner-product preservation is an important result because inner products are used as the similarity metric in a range of classification algorithms, such as support-vector machines (SVMs) [16]. Therefore, by using an appropriately chosen Φ, image detection can be performed *directly* on the compressed outputs $\vec{y}^i$ without the need for prior reconstruction of the original signals $\vec{x}^i$ [17]. The aim of our approach is to employ a Φ that is easily implemented using the low-performance and high-variability TFTs available in LAE. In this way the interfaces within a hybrid system to the CMOS domain can be significantly reduced, while retaining the ability to perform accurate classification over LAE sensor data.

B. Inner-Product Preservation via Random Φ

Random matrices, whose elements are draw from a zero-mean random variable exhibit inner-product preservation. An example of such a matrix Φ, which we use in this work, is one whose elements are drawn from a zero-mean Bernoulli random variable (i.e., one whose elements are ± 1 with probability 1/2) [17]. This particular Φ is selected, because with only ± 1 entries, compression is now reduced to simple add/subtract operations over the signal samples, and the random structure implies tolerance to device-level variations. These two results make implementation of compression via TFT circuits possible.

As illustrated in Fig. 2, one general reason that preservation of the inner products arises is because, for sufficiently large values of M (Φ is $M \times N_C$), $\Phi^{\mathsf{T}}\Phi$ approaches a scaled version of the identity matrix, MI_{NC}. That is, since Φ has elements which are randomly chosen to be ± 1, when $\Phi^{\mathsf{T}}\Phi$ is scaled (i.e., divided by M), diagonal entries are exactly 1, while off-diagonal entries have zero mean and normalized variance scaling with $1/M$. In particular, this scaling is illustrated in Fig. 3(a) for various chosen matrices Φ. We see that larger M improves convergence with the identity matrix for any chosen Φ, thanks to smaller variance of the off-diagonal entries.

Since the off-diagonal entries of $\Phi^{\mathsf{T}}\Phi$ are not precisely zero, causing deviation from a scaled identity matrix, MI_{NC},

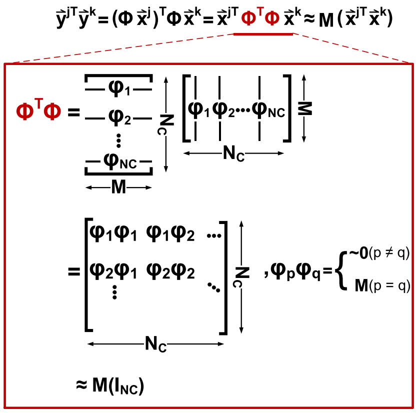

Fig. 2. Mathematical derivation of inner-product preservation after compression via matrix Φ.

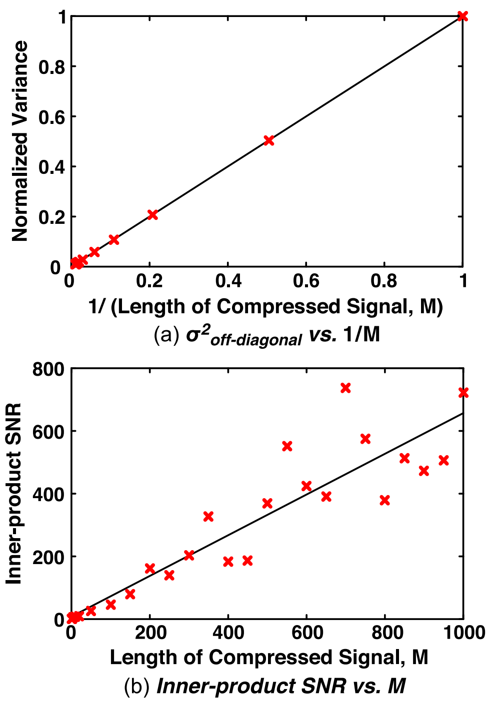

Fig. 3. (a) Normalized variance of the off-diagonal elements of $(1/M) \times \Phi^{\mathsf{T}}\Phi$, versus $1/M$. (b) Inner-product SNR of 20 images from the MNIST database of handwritten digits versus M, compressed using an ideal Φ. Both figures show data averaged over 10 different ideal Φ, with $N_C = 1000$.

we naturally expect the inner products, which we rely on for classification, to also exhibit error with a similar dependence. Since M sets the compression factor, in order to observe its effect on inner-product preservation, we compare the inner products of 1000×1000 (1 Mpixel) images (resized) from the MNIST database of handwritten numerical digits, before and after compression using an ideal $M \times N_C$ Φ, for various M. To quantify the quality of inner product preservation, we calculate the inner-product signal-to-noise ratio (SNR), a measure inversely related to the inner-product error:

$$\text{Inner-product SNR} = \frac{\sum_{j \neq k} \left(M(\vec{x}^{j\mathsf{T}}\vec{x}^k)\right)^2}{\sum_{j \neq k} \left(\vec{y}^{j\mathsf{T}}\vec{y}^k - M(\vec{x}^{j\mathsf{T}}\vec{x}^k)\right)^2} \quad (1)$$

where j, k refer to different rows of the images, as accessed out by an active matrix. Averaging over 10 different ideal Φ for each compression factor, we indeed observe that the inner-product SNR exhibits the expected linear trend with M, as shown in Fig. 3(b).

C. Implications of a TFT-Based Implementation

In our system, the required matrix multiplication corresponding to random projection is implemented using TFT circuits. As previously mentioned, this raises the concerns of large variations and low speed. With regards to large TFT variations, these effectively cause deviations from nominal multiplication by ± 1. However, as illustrated in Fig. 2, what is critical is not that the matrix elements are precisely ± 1, but rather that the off-diagonal elements of $\Phi^{\mathsf{T}}\Phi$ approach zero. This criterion may be met even in the presence of TFT variations, provided that the variations are uncorrelated. Accordingly, even with large TFT variations, we demonstrate (in Section V) that the prototype system maintains image classification performance out to high compression factors. With regards to low TFT speed, as will be described in the next section, multiplication by a matrix whose elements are nominally ± 1 can be implemented in a highly parallel manner, where the additions involved are trivially achieved by TFT current summation on shared nodes. Further, the time constants of the shared nodes can be made much lower than those set by the TFT capacitances.

D. Sparse Image Reconstruction

In the presented system, our interest is in the classification of images. However, we point out that the approach of compression using random projections also has a relationship to the reconstruction of $\vec{x}^i$. The theory of compressive sensing states that if there exists a transform basis, assigned as Ψ, in which a signal is sparse (i.e., the information of the signal is captured by k non-zero transform coefficients, where $k \ll N_C$), the transform coefficients can be determined from a small collection of measurements derived from linear combinations of the signal samples [18]. The only requirement is that $\Phi\Psi$ satisfies the Restricted Isometry Property (RIP) [19]. For typical transforms used for image compression (e.g., DCT), the requirements of RIP are met with high probability when the elements of Φ are drawn from a zero-mean Bernoulli random variable (i.e., ± 1 with probability 1/2) [20].

So indeed our choice of Φ also broadly permits reconstruction of images from the compressed vector $\vec{y}^i$. However, unlike classification, which does not require explicit knowledge of Φ, generally, reconstruction of the original signal does. This poses a limitation because though we have knowledge of the nominal Φ implemented by the system, TFT nonlinearities and variations make the effective Φ difficult to know precisely. Thus, reconstruction performance is limited by these non-idealities. Nonetheless, in Section V we demonstrate that image reconstruction from compressed data is possible.

III. Thin-Film Implementation

Having provided background on the algorithmic approach of the image compression system, we now describe the

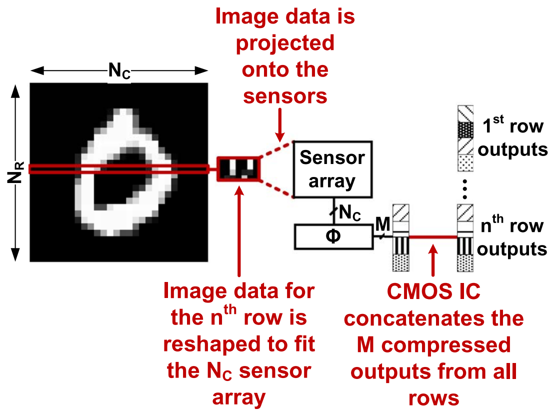

Fig. 4. Implementation of the system, which emulates an $N_R \times N_C$ active matrix with an array of N_C sensors. Image data is projected row-by-row onto the N_C-array of sensors and the corresponding sensor outputs are fed into the $M \times N_C$ compression block Φ. The resulting compressed output data from all rows are then concatenated in the CMOS domain.

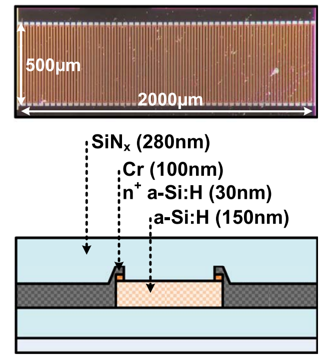

Fig. 5. An image of an a-Si photoconductor and its schematic cross section, similar to that in [8].

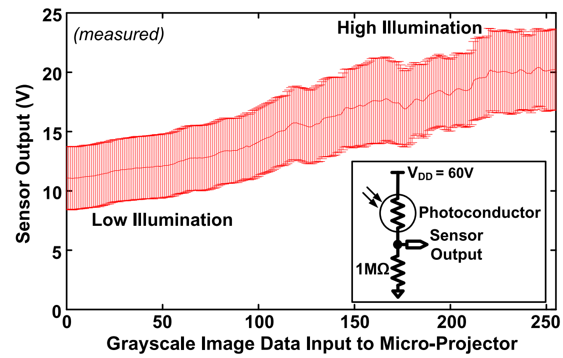

Fig. 6. Photoconductor sensor output voltage versus illumination. The error bars show standard deviation across 28 sensors. The inset shows the sensor voltage divider configuration with the supply voltage and fixed resistor values labeled.

thin-film implementation. The components required in the complete system are as follows: 1) the sensor array, meant to be accessed using an active matrix, and whose data are fed into 2) the TFT-based compression block.

Our aim is to represent an active matrix consisting of a large number of sensors (i.e., large N_R, N_C). However, as illustrated in Fig. 4, instead of fabricating $N_R \times N_C$ image sensors and their corresponding access TFTs, we fabricated only N_C image sensors. Images to be detected are projected one row at a time onto the fabricated sensors, where a uniform square of light is projected onto each sensor, corresponding to a single pixel in the row. Accordingly, pixel data from the N_C sensors are made available one row at a time, as they would be in a typical row-scanning active matrix.

Since sensor data is fed to the compression block row-by-row, compression is also performed one row at a time. Eventually, the compressed output data from all rows of the image are concatenated (Fig. 4). Classification (and reconstruction) is then performed as it would be on a standard CMOS IC.

A. Image Sensors

The image sensors (i.e., pixels) are implemented as shown in Fig. 5, which is similar to [8]. Each pixel corresponds to a voltage divider formed by a fixed 1 MΩ resistor and an island of undoped amorphous-silicon (a-Si) in an inter-digitated layout, serving as a photoconductor. Since such a photoconductor exhibits a suitable response to variations in lighting condition (i.e., light versus dark), the image data is presented to the sensor array via a micro-projector.

As shown in Fig. 6, the level of illumination is thus sensed by the pixel as an output voltage. The measured pixel response when varying the grayscale level of the image data inputted to the micro-projector (i.e., projecting different levels of illumination) exhibits a relatively linear relationship (we expect residual nonlinearity to be addressed through training of the classifier model within the system [8]). The error bars correspond to the standard deviation across 28 different characterized photoconductors (easily accessed via the sample layout). With a 60 V supply voltage, the resulting sensor output voltage range, allows us to drive the subsequent TFTs in the compression block (described below) in the above-threshold regime.

B. TFT-Based Compression Block

Fig. 7 shows the TFT-based implementation of the $M \times N_C$ compression matrix Φ, which reduces the N_C pixel-sensor signals to M output signals. Each of the N_C sensor outputs $(x_1^i \ldots x_{\mathrm{NC}}^i)$, corresponding to data from one row of the projected image, feeds the gates of M TFTs, which correspond to the M rows of Φ.

As previously mentioned, since the elements of Φ are randomly chosen to be ±1, compression simply involves addition/subtraction operations, as determined by the elements of Φ. To implement this, first a TFT is used to convert the pixel-sensor voltage into a current, determined by the transconductance of the TFT. Next, to perform addition or subtraction, the TFTs of each row (driven by $x_1^i \ldots x_{\mathrm{NC}}^i$) are connected together at either a positive or negative current-summing node, depending on whether a matrix entry of +1 or −1 is to be implemented. That is, by construction, the TFTs that represent +1 elements of the ideal Φ are connected to positive current-summing nodes, while the TFTs that represent −1 elements of the ideal Φ, are connected to negative current-summing nodes. This results in differential signals corresponding to the compressed output $y_1^i \ldots y_{\mathrm{M}}^i$. As previously mentioned, the sensors are biased such that their output voltages $(x_1^i \ldots x_{\mathrm{NC}}^i)$ that feed into the

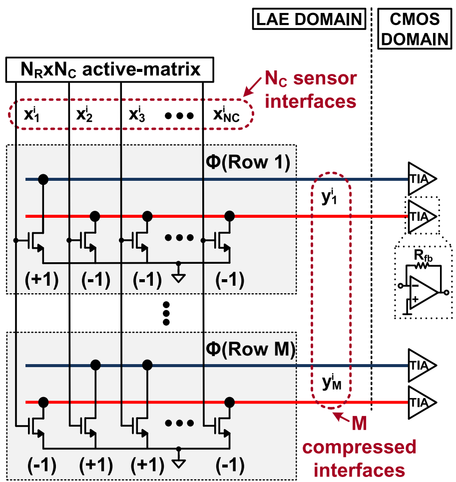

Fig. 7. TFT-based implementation of the $M \times N_C$ compression block Φ. The (+1)'s and (−1)'s correspond to the elements of the ideal Φ.

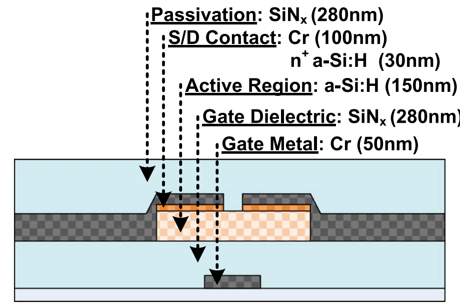

Fig. 8. The schematic cross section of the a-Si TFTs that are used to implement the compression block.

TFT Sizing (μm)	W = 500
	L = 10
Threshold Voltage V_T (V)	$\bar{V}_T$ = 1.1
	σ = 0.4
Mobility μ (cm²/Vs)	$\bar{\mu}$ = 0.63
	σ = 0.03

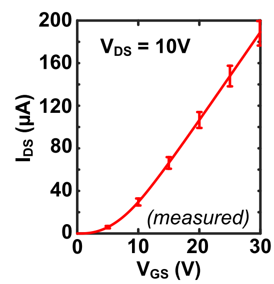

Fig. 9. The $I_{\rm DS}$ vs. $V_{\rm GS}$ curve for the fabricated a-Si TFTs, with error bars showing the standard deviation across 80 devices.

compression block, operate the TFTs in the above-threshold region. A fixed drain-source voltage of 10 V is maintained across the TFTs, by connecting each summing node to the virtual ground of a transimpedance amplifier (TIA), which can be implemented in the CMOS IC. We note that a 10 V bias is not actually required from the CMOS chip, since its ground can be suitably offset with respect to the TFT circuits.

Such an architecture for implementing compression is able to achieve fast operation despite the low speed of TFTs, and is able to do so with a small number of TFTs ($M \times N_C$) compared to the total number of TFTs in an active matrix (minimally $N_R \times N_C$). First, the small number of additional TFTs ensures that the compression block imposes minimal loading on the active-matrix data lines, thereby having small impact on their settling time. Second, despite a small number of additional TFTs, highly parallel operation is achieved, with all outputs of the compression block derived at once. Third, the dominant time constant in the compression block is set by TFTs driving current on the shared summing nodes, where the impedance is substantially reduced thanks to a virtual ground condition imposed by the TIA. Namely, implemented as an op-amp with feedback resistor $R_{\rm fb}$ (Fig. 7), the TIA has an input impedance $Z_{\rm in} \approx R_{\rm fb}/(1 + A)$, where A is the open-loop gain of the op-amp. Thus, despite a large total capacitance on this node set by the TFTs ($C_{\rm tot} \approx C_{\rm gd} \times N_C$), even modest op-amp gain results in a low input impedance and a small time constant $\tau_{\rm comp} = Z_{\rm in} \times C_{\rm tot}$. As an example, using our compression-block design with an X-ray imager having 2304 data lines and 3200 gate lines, performing row scanning at 1 kHz [12], yields a $C_{\rm tot}$ of approximately 2800 pF ($\approx$ 2304 TFTs × 1.25 pF/TFT). Our $N_C = 80$ implementation employs a TIA with $R_{\rm fb} =$ 50 kΩ (though an even smaller resistor would be used with $N_C = 2304$), giving $Z_{\rm in} < 50\ \Omega$, assuming an op-amp gain of just 1000. The resulting time constant $\tau_{\rm comp} \approx 1.4 \times 10^{-7}$ sec, which is much faster than the 1 kHz row-scanning rate.

In this system, the compression block is implemented using a-Si TFTs that are fabricated in-house with our standard low-temperature (< 180 °C) process [21]. The schematic cross section of the a-Si TFT layers is pictured in Fig. 8. Fig. 9 shows the measured $I_{\rm DS}$ vs. $V_{\rm GS}$ transfer curve for the fabricated a-Si TFTs, with error bars showing the standard deviation across 80 devices. With an average sensor output-voltage range of 11–20 V (Fig. 6), the TFT current levels are roughly 50–100 μA. This is suitable for our 80-column system, but TFT sizing and biasing (i.e., sensor output-voltage range) can be designed for active matrices of larger sizes. In addition to the non-linearity of TFT transfer curves, the substantial variation observed across the TFTs implies that the sensor signals fed to the TFT gates are not multiplied by exactly ± 1 (or, more accurately, a constant transconductance across the TFTs), as assumed for an ideal Φ. We analyze the effects of a non-ideal TFT-implemented Φ in the following section. We point out that while our interest in this work is in exploring the ability of the compression approach to overcome such non-idealities, circuit-level solutions may additionally be employed in the system (e.g., TFT source degeneration for enhancing linearity).

IV. Analysis of TFT Non-Idealities

In this section, we examine the consequences of using a TFT-based Φ for implementing compression via random projection. We do this by performing simulations in MATLAB. In particular, we are interested in the effects on inner-product preservation of variations and nonlinearity in the TFT transfer curves (i.e., $I_{\rm DS} - V_{\rm GS}$ relationship, with constant $V_{\rm DS}$). We analyze this by both modeling and measuring the actual TFT transfer curves (this is done for characterization and analysis only, not for actual operation in the system). With inner products serving as a similarity metric for classification, it is clear that inner-product preservation has correspondence

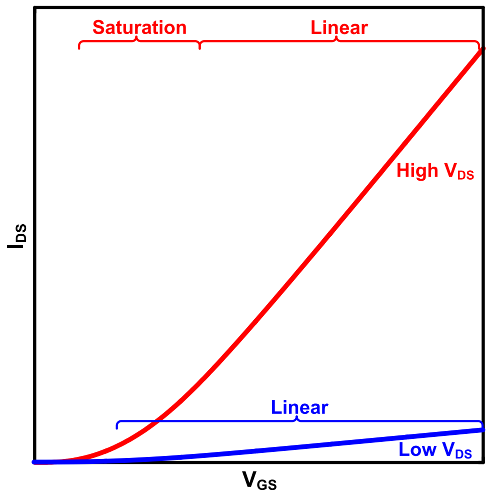

Fig. 10. Comparison of linear regions for above-threshold operation of TFTs with different $V_{\rm DS}$.

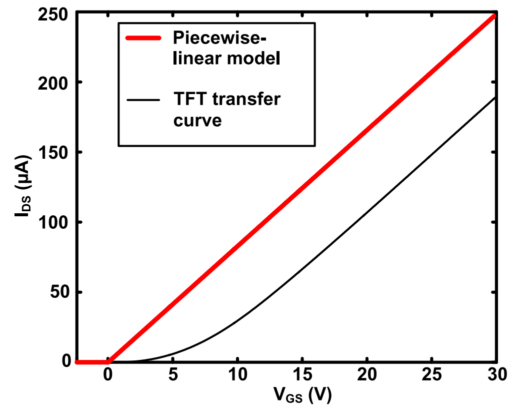

Fig. 11. Piecewise-linear model curve compared to a measured $I_{\rm DS}$ vs. $V_{\rm GS}$ curve averaged across 80 TFTs.

with classification performance. However, we note that classification via data-driven (machine-learning) training algorithms present significant opportunities to overcome errors in the inner-product-preserving compressed vectors. For instance, the approach of Data-Driven Hardware Resilience (DDHR) [22], exploits training *to* the error-affected data so that distortions to the data are learned during the classifier-training process; this substantially enhances error tolerance further. Thus, both inner-product preservation and the ultimate classification performance must be analyzed. We also point out that in addition to stationary variations in device parameters, the TFTs may possibly be subject to non-stationary variations (i.e., drifts). Though a high level of stability in the TFT parameters can be achieved through processing techniques [21], generally, in a machine-learning system classifier, retraining may also be employed to track any resulting changes in the data statistics. Further, specifically within the random-projection approach, if such variations impact the TFT currents in an uncorrelated manner (e.g., after simple mean subtraction), we expect them to be addressed similar to stationary variations.

We start by analyzing inner-product preservation in this section. As mentioned in Section II, the quality of inner-product preservation depends on the how well $\Phi^{\sf T}\Phi$ approximates a scaled version of the identity matrix, $MI_{\rm NC}$. To quantitatively measure the distance between these two matrices, we employ the l1-norm:

$$\text{Distance From Identity} = \left\| \frac{1}{M} \times \Phi^{\sf T}\Phi - I_{\rm NC} \right\|_1. \quad (2)$$

By introducing variations and nonlinearity in the TFT transfer curve and performing simulations, we aim to more precisely analyze the impact on system performance (i.e., image classification and reconstruction results). We first analyze the effects of variations by using a piecewise-linear model of the TFT transfer curve. As can be seen from Fig. 10, such a model well-represents the transfer curve of a TFT with a low $V_{\rm DS}$, since in this case, the TFT is almost entirely in the linear region, when operating above threshold. However, low $V_{\rm DS}$ values do not achieve large enough currents for practical system operation. As such, the TFTs are actually operated at a higher $V_{\rm DS}$ (i.e., $V_{\rm DS} = 10$ V), which introduces a substantial nonlinear (saturation) region (Fig. 10). Thus, we also model the effects of such nonlinearity by employing actual measured TFT transfer curves, representing the true shape. We note that, strictly speaking, a piecewise-linear model also implies nonlinearity; however, for the model described below, at least nominally, the nonlinearity can be negated by simply offsetting the input pixel voltage (as described below).

A. *Piecewise-Linear Model for Analyzing Variations*

The typical parameters of interest for TFT variations are threshold voltage and mobility. Indeed, as seen from measurements of a-Si TFTs in Fig. 9, these parameters exhibit high variability. Using MATLAB, we generate a statistical model, to independently simulate variations in threshold voltage and mobility. In this way, it is possible to observe their individual effects on $\Phi^{\sf T}\Phi$ and, in particular, its distance from identity.

In order to model the offset of a TFT transfer curve (i.e., the sub-threshold versus above-threshold regimes of the TFT), a piecewise-linear model is employed, as shown in Fig. 11. That is, for values below a cutoff voltage the output of the piecewise-linear transfer curve is 0. On the other hand for values above the cutoff voltage, the piecewise-linear model has non-zero values that are linearly dependent on the input sensor voltages. The cutoff voltage is selected to be 0 V, so that in the ideal case, the curve is essentially linear. Since the pixel voltages are shifted to fall into the above-threshold region of the TFT transfer curve, the slope of the linear model is selected to be the same as the slope of a TFT transfer curve (averaged across 80 measured TFT devices) in this region (Fig. 9).

To simulate the variations in threshold voltage and mobility, the cutoff voltage (V_T) and slope (γ) of the piecewise-linear model are varied, respectively (Fig. 12). The shifts in threshold voltage and slope are drawn from a normal distribution for various different standard deviations σV_T and $\sigma\gamma$.

In order to derive the corresponding compression matrix Φ, we remember that the magnitude of any given element of Φ, ϕ_a, is defined as the multiplier of the input sensor voltage, V_a, which results in the output current, I_a:

$$\phi_a \equiv \frac{I_a}{V_a}. \quad (3)$$

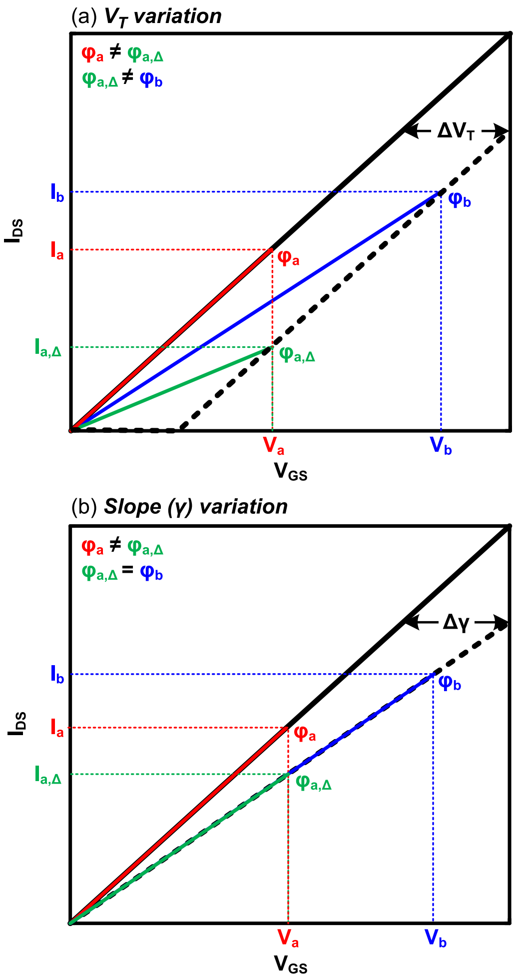

Fig. 12. Simulation of variations in (a) threshold voltage, V_T and (b) slope, γ, using the piecewise-linear model of the TFT transfer curve. Different effective ϕ arise due to these variations.

The shape of the piecewise-linear transfer curve results in a different effective ϕ_a for different input pixel values (i.e., the effective Φ we have *changes* with the input voltage level seen). This is because even though the ideal model is essentially linear, and thus has a fixed slope, variations in threshold voltage introduce an offset in the piecewise-linear model (Fig. 12(a)). The offset means that the ratio between the output current and its corresponding input voltage, ϕ_a, depends on the value of the input voltage itself, as shown in Fig. 12(a). Therefore, for a single row of an $N_R \times N_C$ image, there are $M \times N_C$ different ϕ_a, which make up a single Φ of dimension $M \times N_C$. Thus, for the entire image, there are N_R different Φ. Similarly, for different variations in slope, ϕ_a will also vary (Fig. 12(b)). However, since variations in slope do not introduce an offset to the piecewise-linear model, ϕ_a is the same across an entire image for a given device with a particular slope. Thus, for an entire $N_R \times N_C$ image, there are only $M \times N_C$ different ϕ_a corresponding to the different slopes.

This means that the distance of $\Phi^\mathsf{T}\Phi$ from MI_{NC} is modified from (2), to now be an average across all the different $\Phi^{j\mathsf{T}}\Phi^k$:

$$\text{Distance from Identity} = \frac{1}{\binom{N}{2}} \times \sum_{j\neq k} \left\| \frac{1}{M} \times \Phi^{j\mathsf{T}}\Phi^k - I_{\mathrm{NC}} \right\|_1 \tag{4}$$

where $j,\ k$ refer to the effective Φ for different image-pixel features (i.e., different rows of the active matrix) and N refers to the total number of effective Φ across all images. There are $\binom{N}{2}$ different combinations of $\Phi_i^\mathsf{T}\Phi_j$ to average over.

A further modification to (4) must be made, since a TFT, and thus the piecewise-linear model, does not multiply by 1, but rather introduces a scaling constant G, dependent on the transconductance of the device:

$$\text{Distance from Identity} = \frac{1}{\binom{N}{2}} \times \sum_{j\neq k} \left\| \frac{1}{M} \times G \times \Phi^{j\mathsf{T}}\Phi^k - I_{\mathrm{NC}} \right\|_1 \tag{5}$$

where $j,\ k$ refer to the effective Φ for different features (i.e., different rows of the image) and N refers to the total number of effective Φ across all images. For this piecewise-linear model, G is calculated based on the selected slope of the ideal model. More specifically, G is equal to the inverse of this slope squared. This is because, in calculating $\Phi^{j^\mathsf{T}}\Phi^k$, this slope, or transconductance, is included twice: once in Φ^j and once in Φ^k. Indeed, since introducing variations results in different ϕ, and thus effective transconductance, this G is still a good approximation, since it represents the average transconductance across all devices.

To observe the effect of these variations on $\Phi^\mathsf{T}\Phi$, with respect to compression factor, the distances are calculated for $N_C = 80$, for various $\sigma V_T/\sigma\gamma$ using (5) (averaged over 10 different sampling cases per $\sigma V_T/\sigma\gamma$). Since the effective Φ depend on the input voltages, we use 10 80 $\times$ 80 images (resized) from the MNIST database of handwritten digits, resulting in 800 different Φ. Moreover, the $\Phi^\mathsf{T}\Phi$ distances from MI_{NC} are averaged across 10 different ideal Φ (i.e., 10 different cases of matrix entries from Bernoulli sampling). Fig. 13 shows the simulation results for various compression factors. As expected, as σ increases, the distance of $\Phi^\mathsf{T}\Phi$ from identity increases, particularly at higher compression levels.

Indeed, the piecewise-linear model allows us to analyze the effect of device variations on Φ. However, we are also interested in measuring the effects of TFT transfer-curve shape (i.e., non-linearity) on Φ. This is explored next.

B. *TFT I_{DS} vs. V_{GS} Curve Lookup Table for Nonlinearity Analysis*

To measure the effect of the TFT transfer curve non-linearity, we use a lookup table of I_{DS} vs. V_{GS} curves derived from measurement. In order to also include the effect of threshold voltage and mobility variations, 80 different measured transfer curves are used in the lookup table (Fig. 14). Thus, since $N_C = 80$, there is a different TFT curve for each sensor signal for a given row of the image. Similar to the piecewise-linear model, non-linearity in the TFT transfer-curve leads to a different effective multiplier ϕ_a (i.e., entry in the Φ matrix) for different input pixel voltage values. Thus, the effect of non-linearity in the TFT transfer curve is quantitatively measured using (5). Since this is not an ideal, piecewise-linear curve, it is more difficult to calculate an appropriate G. This is because there is no single slope to base it on as in the piecewise-linear model. Thus, G is

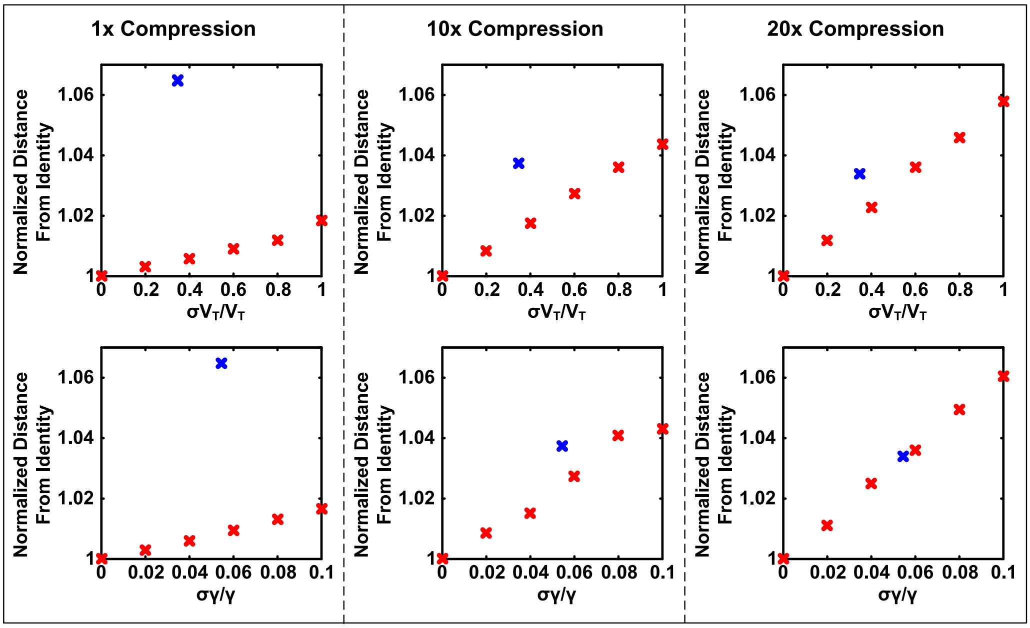

Fig. 13. Distance of $\Phi^{\mathsf{T}}\Phi$ from MI_{NC} versus standard deviation of the threshold voltage and slope for various compression factors, simulated using a piecewise-linear model (red markers) and measured TFT transfer-curves (blue marker). 10 images from the MNIST database of handwritten digits are used as sensor data. The data is averaged over 10 different ideal Φ, with $N_C = 80$.

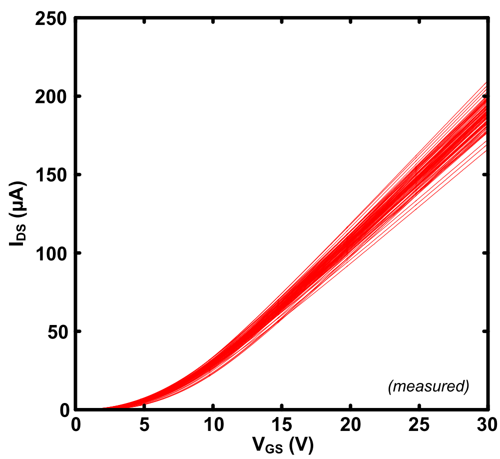

Fig. 14. The 80 TFT I_{DS} vs. V_{GS} curves used in the lookup table to measure distance of $\Phi^{\mathsf{T}}\Phi$ from MI_{NC}.

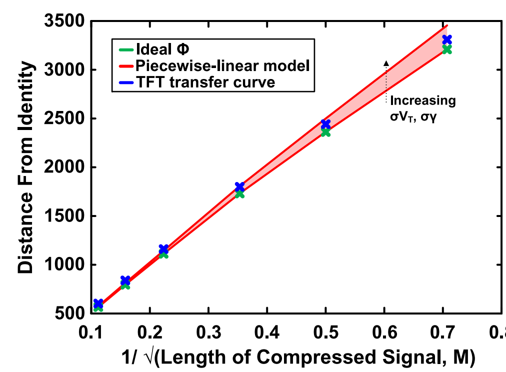

Fig. 15. Comparison of distance of $\Phi^{\mathsf{T}}\Phi$ from MI_{NC} versus $1/\sqrt{M}$ for ideal Φ, the piecewise-linear model and a TFT transfer curve lookup table. For the piecewise-linear model, the red band refers to the area covered by varying σV_T and $\sigma\gamma$. 10 images from the MNIST database of handwritten digits are used as the sensor data for both the piecewise-linear model and TFT transfer curve lookup table simulations. The data is averaged over 10 different ideal Φ, with $N_C = 80$.

calculated by averaging the ratios of the diagonal of the scaled identity matrix $(M \times I_{\mathrm{NC}})$ over the diagonals of $\Phi^{j^{\mathsf{T}}}\Phi^k$:

$$G = \frac{1}{\binom{N}{2}} \times \sum_{j \neq k} \frac{\sum M \times I_{\mathrm{NC,Diagonal}}}{\sum(\Phi^{j^{\mathsf{T}}}\Phi^k)_{,\mathrm{Diagonal}}}. \quad (6)$$

where $j,\ k$ refer to the effective Φ for different features (i.e., different rows of the image) and N refers to the total number of effective Φ across all images.

Since the 80 measured transfer curves have a σV_T and $\sigma\gamma$ shown in Fig. 9 (the measured $\sigma\mu$ is used to estimate $\sigma\gamma$), we can compare the effect of the nonlinear TFT transfer-curve shape to the piecewise-linear transfer-curve shape for these values of σ. This is shown as the blue marker in Fig. 13. We observe that the nonlinear TFT transfer curve indeed causes a larger distance of $\Phi^{\mathsf{T}}\Phi$ from MI_{NC}, implying that TFT nonlinearity has notable impact.

C. Comparison With Ideal Φ

We know that even for an ideal $M \times N_C$ Φ, $\Phi^{\mathsf{T}}\Phi$ does not exactly equal MI_{NC}, especially as M decreases (for instance, this can be seen in Fig. 3(a)). Thus, we require a baseline with which to compare the effects of TFT variations and nonlinearity. It is natural to use a baseline that corresponds to the distance from identity of $\Phi^{\mathsf{T}}\Phi$, for an ideal Φ, using (2).

Fig. 15 shows the distances from an identity matrix when using an ideal Φ, a Φ based on the piecewise-linear model of a TFT transfer curve (with simulated variations), and a Φ based on measured TFT transfer curves (with measured variations). As expected, the ideal Φ has the lowest distance from identity, while the TFT curve simulation has the largest distance from identity. From this comparison, the relative effects of TFT variations and nonlinearity can be clearly observed.

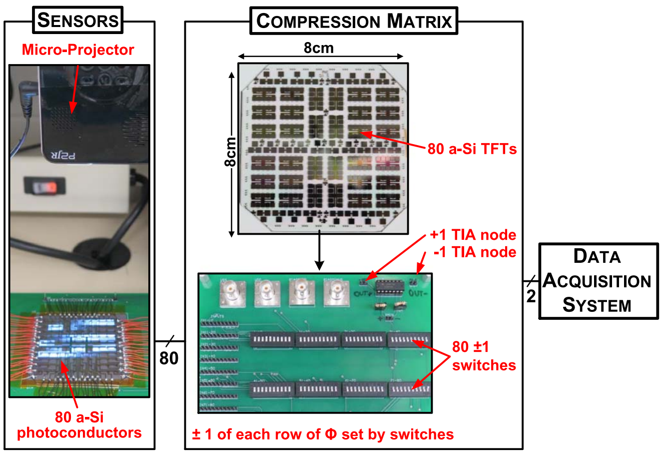

Fig. 16. Testing setup for image compression system prototype.

While distance from the identity matrix is a useful metric for analyzing the feasibility of implementing random-projection-based compression using TFTs, in an eventual system, classification and reconstruction performance is of ultimate interest. While we experimentally show the reconstruction performance achieved by our prototype system, we are particularly interested in demonstrating classification performance. In the following section, we show experimentally that with practical levels of TFT variation and nonlinearity, a classification system for recognizing images of numerical handwritten digits can achieve significant levels of compression.

V. System Demonstration

In this section, details of the image sensing and compression system prototype are presented. Fig. 16 shows the experimental setup. The system prototype consists of a-Si photoconductors and a-Si TFTs fabricated onto two separate 8 $\times$ 8 cm glass substrates at temperatures $< 180\,^{\circ}$C.

The system represents an 80 $\times$ 80 active-matrix array of photoconductor pixels (i.e., $N_R = N_C = 80$) using an array of 80 sensors (with $\sim$4 mm spacing) fabricated onto one glass substrate. A micro-projector is used to project the images, row-by-row, onto the photoconductors. Since the photoconductors are not positioned in a single row, but rather are distributed in a rectangular array, the projected image data is re-arranged from a single row to fit the arrangement of sensors in the array (as shown in Fig. 4). Thus, one pixel of image data is projected onto a single, corresponding photoconductor. The outputs from the sensor array are passed to the next substrate which contains the TFT-based compression matrix. To facilitate testing, a switch board is used, which enables 80 TFTs to be configured to any combination of ± 1 required to implement the different M rows of Φ. The TFT currents, thus combined, are fed to two TIAs (for the +1 and -1 summing nodes), whose output voltages are acquired by a PC using a DAQ system. By cycling through M different switch configurations (corresponding to the M rows of Φ), we can generate the M-element output $\vec{y}^i$ with just 80 TFTs (instead of $M \times 80$) and two TIAs (instead of 2 M). In this way, this system can implement any compression factor; indeed, in an actual system where only a single compression factor would be implemented, the $M \times 80$ TFTs implementing Φ would be hardwired to the corresponding +1 and -1 summing nodes. Finally, the PC concatenates the compressed outputs across all rows of the projected image, after which image classification and reconstruction can be performed.

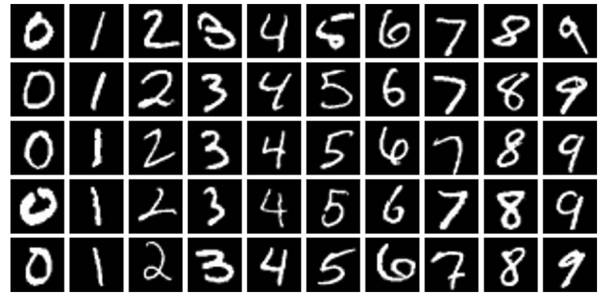

Fig. 17. A subset of the dataset derived from images from the MNIST database of handwritten digits [13], resized from the original 28 $\times$ 28 to 80 $\times$ 80 pixels, used for image classification and reconstruction.

To demonstrate the performance of the image compression system, image classification and reconstruction is performed on images from the MNIST database of handwritten digits. In total, the dataset consists of 1500 images with equal number of instances for each digit ("0" to "9"). Each image is resized from the original 28 $\times$ 28 pixels to 80 $\times$ 80 pixels since the prototype system emulates an 80 $\times$ 80 array of pixels. For illustration, a subset of the images in the dataset is shown in Fig. 17.

A. Inner-Product Preservation

In order to analyze the quality of inner-product preservation by the implemented TFT-based compression matrix, we are interested in the resulting effective Φ that corresponds to the prototype system (i.e., how closely $\Phi^{\mathsf{T}}\Phi$ represents MI_{NC}). As mentioned in Section II, this information can be extrapolated from the inner-product SNR. However, we use a slightly modified definition of (1), namely with the addition of a proportionality constant, α

$$\text{Inner-product SNR} = \frac{\sum_{j \neq k} \left(M\alpha \left(\vec{x}^{j\mathsf{T}} \vec{x}^{k} \right) \right)^2}{\sum_{j \neq k} \left(\vec{y}^{j\mathsf{T}} \vec{y}^{k} - M\alpha \left(\vec{x}^{j\mathsf{T}} \vec{x}^{k} \right) \right)^2} \quad (7)$$

where j, k refer to the different features. As previously mentioned, this is due to the fact that instead of actually multiplying by ± 1, a given TFT in the compression matrix outputs a current in response to an input pixel voltage, which is determined by the transconductance of the TFT. Thus, the scaling constant α is introduced and set using:

$$\alpha = \frac{\sum_{j \neq k} \vec{y}^{j\mathsf{T}} \vec{y}^{k}}{M \times \sum_{j \neq k} \left(\vec{x}^{j\mathsf{T}} \vec{x}^{k} \right)}. \quad (8)$$

Fig. 18 shows the inner-product SNR (as defined in (7)), versus length of the compressed signal, M, obtained from the thin-film compression system averaged over 20 images. Even for a shorter input signal length ($N_C = 80$), the data follows the expected trend (shown in Fig. 3(b)), though we observe a lower inner-product SNR. This is due to the addition of noise from variation in the sensors, as well as both the variation and

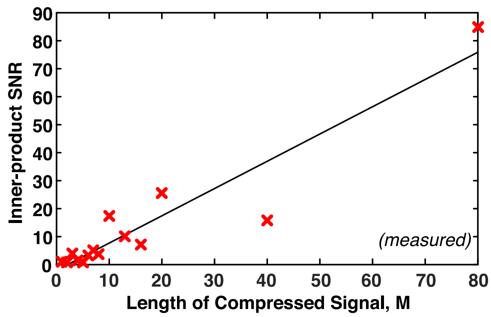

Fig. 18. Inner-product SNR collected from 20 of the images from the MNIST database of handwritten digits that were compressed using the TFT-based compression block Φ.

transfer-function nonlinearity of the compression-matrix TFTs (as discussed in Section IV).

Nevertheless, as we see from the following classification results, an adequately high inner-product SNR is achieved even for small lengths of the compressed signals.

B. Image Classification

As mentioned in Section IV, classification benefits from the fact that we do not need to know the precise compression matrix Φ utilized (indeed the precise Φ is difficult to know in the presence of variation and nonlinearity). Further, classification benefits from the ability to train the classifier to data resulting from the non-ideal compression. Thus a high level of classification performance and/or a high level of compression is possible.

To demonstrate this, one-versus-all classification is performed for each digit using a SVM with a radial-basis-function kernel. For ease of testing, a MATLAB-implemented SVM classifier is used; however, such a classifier can be readily integrated in a CMOS IC [23]. To divide the dataset into appropriate training and testing subsets, ten-fold cross validation is performed.

To characterize the classification performance we measure the true-positive (tp), true-negative (tn), and error rates, which are defined (for the case where we wish to classify the digit "0" vs. digits "1" to "9") as:

$$\text{tp rate} = \frac{\#\text{ of correctly classified ``0''s}}{\text{Total }\#\text{ of ``0''s in the test dataset}} \tag{9}$$

$$\text{tn rate} = \frac{\#\text{ of correctly classified other digits}}{\text{Total }\#\text{ of other digits in the test dataset}} \tag{10}$$

$$\text{error rate} = \frac{\#\text{ of incorrectly classified digits}}{\text{Total size of test dataset}}. \tag{11}$$

The measured tp, tn and error rates, for all digits, versus compression factors between $8\times$ to $80\times$, are shown in Fig. 19. High levels of classification performance are observed, even out to large compression factors. For instance, at $20\times$ compression, the average tp/tn/error rates are 90%/93%/7%. This is typical of the performance achieved with this dataset [13], yet with a substantially reduced number of interfaces (i.e., from 80 to 4 in this proof-of-concept demonstration).

C. Image Reconstruction

Though our primary interest for this system is classification, we also perform image reconstruction, assuming that $\Phi\Psi$ satisfies RIP (as required for compressive sensing). Here, Ψ represents the 1-D DCT basis, wherein images exhibit sparsity. In this case, using, for instance, the gradient projection for sparse reconstruction algorithm (GPSR) [24], it is possible to solve for the transform coefficients from the compressed outputs.

Fig. 20 shows representative results of a reconstructed image of a handwritten "3" for compression factors ranging from $1\times$ to $6\times$. Though our experimental approach is based on measuring the TFT transfer functions to precisely characterize the Φ transform applied to each input image, in a typical system, we expect that such measurements would not be performed. Thus, precisely knowing Φ, which is required for reconstruction, would not be possible in the presence of TFT variations and nonlinearity. Consequently, reconstruction performance is degraded compared to classification performance, especially as the compression rate increases. Furthermore, image compression algorithms such as JPEG use a 2-D DCT [14] for the sparsity basis Ψ, whereas we employ a 1-D DCT; this is because accessing the pixel array using an active matrix configuration precludes use of a 2-D DCT. That is, accessing only the column data for a given row results in loss of spatial information across rows. Finally, as a result of our experimental approach, where readout of an $N_R \times N_C$ active matrix is represented by projecting images row-by-row on a single set of N_C sensors, we observe that the sensed image (before TFT-based compression) exhibits a horizontal pattern (note that the image is rotated such that a horizontal pattern corresponds to a row of sensors) (Fig. 20). Such a pattern is also observed in the reconstructed images; but its impact on classification is believed to be minimal since training accounts for such variations. Nevertheless, it can be seen that for image reconstruction, some compression can be achieved.

VI. Conclusion

In this paper, we presented an image compression system based on LAE thin-film devices. LAE allows the formation of large area yet dense sensing arrays, which are suitable for use in large-scale systems. However, such systems require a LAE-CMOS hybrid architecture, as TFTs cannot compare to the high efficiency of CMOS ICs for processing and analysis of sensor data. Yet, difficulties in interfacing the two technologies, particularly due to the large number of interfaces, have made the implementation of such systems challenging.

A widely-used circuit-based approach for addressing sensor-interfacing challenges is the active matrix, through which sensor data is read out row-by-row, reducing the number of interfaces by approximately a square-root factor. Nevertheless, even when using active-matrix accessing, the number of interfaces scales with the total number of sensors. By incorporating random projections, a concept from statistical signal processing, the number of interfaces can be further reduced. Thus, we presented a thin-film image compression system, which employs a TFT-implemented block for multiplication by an $M \times N_C$ random-projection matrix, in conjunction with an

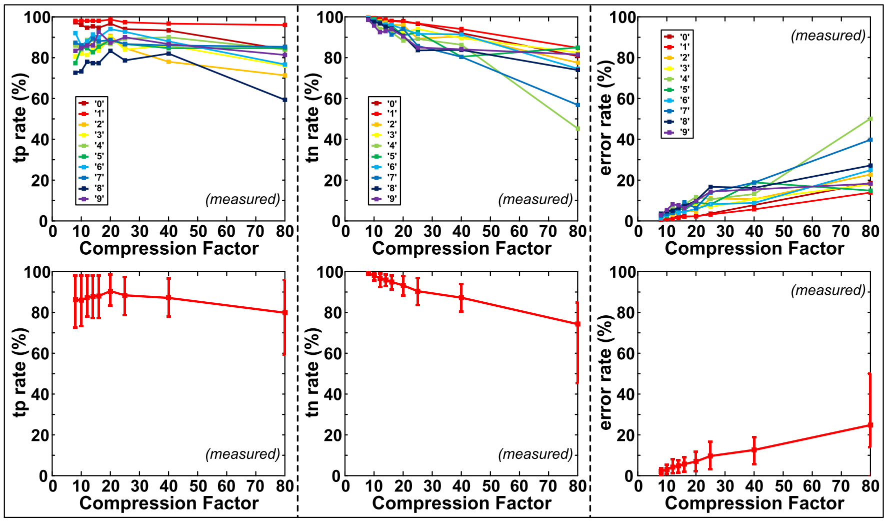

Fig. 19. One-versus-all classification performance (i.e., true-positive, true-negative, and error rates) of images from the MNIST database of handwritten digits versus compression factor for all digits ("0" to "9") (top) and averaged across all digits (error bars show min/max performance across digits) (bottom).

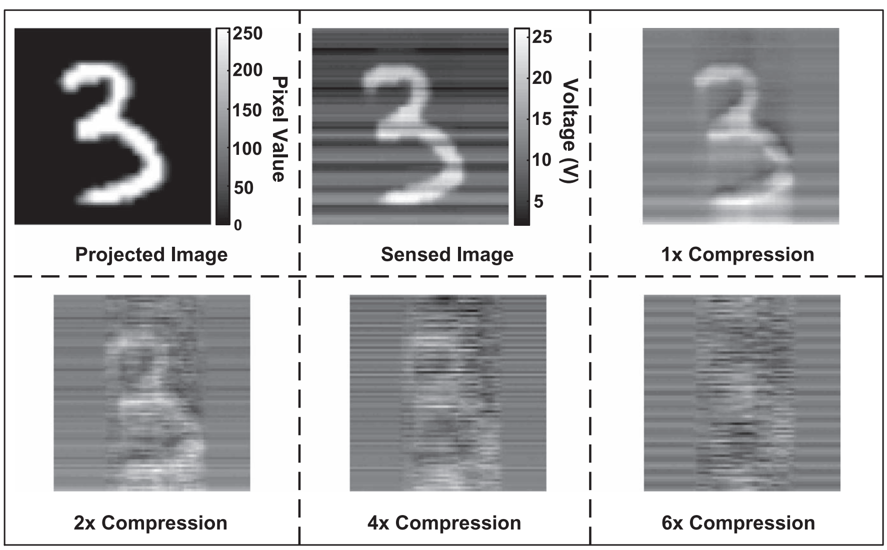

Fig. 20. Reconstruction of an image of a "3" from the MNIST database of handwritten digits with different compression factors using GPSR. Images are rotated 90 clockwise for visualization.

$N_R \times N_C$ active matrix, to substantially reduce the number interfaces by a factor of N_C/M.

The performance of the system was demonstrated by implementing classification and reconstruction of images from the MNIST database of handwritten digits. Emulating an 80×80 active-matrix array of photoconductor pixels, up to $80\times$ compression of the 80 interface-signals was demonstrated, with $20\times$ compression achieving average tp/tn/error rates of 90%/93%/7%.

REFERENCES

[1] D. He, V. Nausieda, K. Ryu, A. Akinwande, V. Bulovic, and C. Sodini, "An integrated organic circuit array for flexible large-area temperature sensing," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2010, pp. 142–143.

[2] H. Wang, L. Chen, J. Wang, Q. Sun, and Y. Zhao, "A micro oxygen sensor based on a nano sol-gel TiO_2 thin film," *Sensors*, vol. 14, no. 9, pp. 16 423–16 433, Sep. 2014.

[3] L. Zhuo, S. Jung, E. Brandon, and T. N. Jackson, "Flexible substrate micro-crystalline silicon and gated amorphous silicon strain sensors," *IEEE Trans. Electron Devices*, vol. 53, no. 2, pp. 380–385, Feb. 2006.

[4] S. Y. Han, K. T. Park, H. S. J. Jeon, Y. W. Heo, and B. S. Bae, "Optical properties of a-SiGe: H thin film transistor for infrared image sensors in touch sensing display," *J. Display Technol.*, vol. 8, no. 10, pp. 617–622, Oct. 2012.

[5] C. Yeom, K. Chen, D. Kiriya, Z. Yu, G. Cho, and A. Javey, "Large-area compliant tactile sensors using printed carbon nanotube active-matrix backplanes," *Adv. Mater.*, vol. 27, no. 9, pp. 1561–1566, Mar. 2015.

[6] N. D. Young, G. Harkin, R. M. Bunn, D. J. McCulloch, R. W. Wilks, and A. G. Knapp, "Novel fingerprint scanning arrays using polysilicon TFT's on glass and polymer substrates," *IEEE Electron Device Lett.*, vol. 18, no. 1, pp. 19–20, Jan. 1997.

[7] H. S. Kim and K. Y. Han, "High-linearity in-pixel thermal sensor using low-temperature poly-Si thin-film transistors," *IEEE Sensors J.*, vol. 15, no. 2, pp. 963–970, Feb. 2015.

[8] W. Rieutort-Louis, T. Moy, Z. Wang, S. Wagner, J. C. Sturm, and N. Verma, "A large-area image sensing and detection system based on embedded thin-film classifiers," *IEEE J. Solid-State Circuits*, vol. 51, no. 1, pp. 281–290, Jan. 2016.

[9] N. Verma, Y. Hu, L. Huang, W. Rieutort-Louis, J. Sanz-Robinson, T. Moy, B. Glisic, S. Wagner, and J. C. Sturm, "Enabling scalable hybrid systems: Architectures for exploiting large-area electronics in applications," *Proc. IEEE*, vol. 103, no. 4, pp. 690–712, Apr. 2015.

[10] Y. Hu, L. Huang, W. Rieutort-Louis, J. Sanz-Robinson, J. C. Sturm, S. Wagner, and N. Verma, "A self-powered system for large-scale strain sensing by combining CMOS ICs with large-area electronics," *IEEE J. Solid-State Circuits*, vol. 49, no. 4, pp. 838–850, Apr. 2014.

[11] J. Sanz-Robinson, L. Huang, T. Moy, W. Rieutort-Louis, Y. Hu, S. Wagner, J. C. Sturm, and N. Verma, "Large-area microphone array for audio source separation based on a hybrid architecture exploiting thin-film electronics and CMOS," *IEEE J. Solid-State Circuits*, vol. 51, no. 4, pp. 979–991, Apr. 2016.

[12] R. L. Weisfield, M. A. Hartney, R. A. Street, and R. B. Apte, "New amorphous-silicon image sensor for x-ray diagnostic medical imaging applications," in *Proc. SPIE*, vol. 3336, 1998, pp. 444–452.

[13] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.

[14] G. K. Wallace, "The JPEG still picture compression standard," *IEEE Trans. Consum. Electron.*, vol. 38, no. 1, pp. xviii–xxxiv, Feb. 1992.

[15] S. Dasgupta and A. Gupta, "An elementary proof of a theorem of Johnson and Lindenstrauss," *Random Struct. Algorithms*, vol. 22, no. 1, pp. 60–65, 2003.

[16] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines and Other Kernel Based Learning Methods*. Cambridge, U.K.: Cambridge Univ. Press, 2000.

[17] R. Calderbank, S. Jafarpour, and R. Schapire, "Compressed learning: Universal sparse dimensionality reduction and learning in the measurement domain," Tech. Rep., 2009. [Online]. Available: http://dsp.rice.edu/files/cs/cl.pdf

[18] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[19] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.

[20] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constr. Approx.*, vol. 28, no. 3, pp. 253–263, Dec. 2008.

[21] B. Hekmatshoar, K. Cherenack, A. Kattamis, K. Long, S. Wagner, and J. C. Sturm, "Highly stable amorphous-silicon thin-film transistors on clear plastic," *Appl. Phys. Lett.*, vol. 93, no. 032103, Jul. 2008.

[22] Z. Wang, K. H. Lee, and N. Verma, "Overcoming computational errors in sensing platforms through embedded machine-learning kernels," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 23, no. 8, pp. 1459–1470, Aug. 2015.

[23] K. H. Lee and N. Verma, "A low-power processor with configurable embedded machine-learning accelerators for high-order and adaptive analysis of medical-sensor signals," *IEEE J. Solid-State Circuits*, vol. 48, no. 7, pp. 1625–1637, Jul. 2013.

[24] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Topics Signal Process.*, vol. 1, no. 4, pp. 586–597, Dec. 2007.

Tiffany Moy (S'14) received the B.S.E. (*magna cum laude*) and M.A. degrees in electrical engineering from Princeton University, Princeton, NJ, USA, in 2012 and 2014, respectively. She is currently working toward the Ph.D. degree in electrical engineering at Princeton University.

Her research interests include thin-film circuits and algorithms for hybrid large-area electronics/CMOS system design.

Warren Rieutort-Louis (S'12–M'15) received the B.A. (Hons.) and M.Eng. degrees in electrical and information engineering from Trinity College, Cambridge University, Cambridge, U.K., in 2009, the M.A. degree in electrical engineering from Princeton University, Princeton, NJ, USA, in 2012, and the Ph.D. degree in Electrical Engineering at Princeton University in 2015.

He was a Graduate Teaching Fellow with Princeton McGraw Center for Teaching and Learning. His research interests include thin-film materials, processes, devices, and circuits for large-area electronic systems.

Dr. Rieutort-Louis was the recipient of the IBM Ph.D. Fellowship, the Andlinger Center Maeder Fellowship in Energy and the Environment, the Princeton Harold W. Dodds Honorific Fellowship.

Sigurd Wagner (SM'78–F'00) received the Ph.D. degree from the University of Vienna, Vienna, Austria.

Following a Post-Doctoral Fellowship at Ohio State University, he worked from 1970 to 1978 at the Bell Telephone Laboratories in Murray Hill and Holmdel, NJ, USA, on semiconductor memories and heterojunction solar cells. He then joined the Solar Energy Research Institute (now NREL), in Golden, CO, USA, as the founding Chief of the Photovoltaic Research Branch. Since 1980, he has been Professor of Electrical Engineering at Princeton University, Princeton, NJ, USA; in 2015 he became Professor Emeritus and Senior Scholar. He has been developing fundamentally new materials, processes, and components for flexible large-area electronics, electrotextiles, and electronic skin, and is considered the father of soft elastic electronics.

James C. Sturm (S'81–M'85–SM'95–F'01) was born in Berkeley Heights, NJ, USA, in 1957. He received the B.S.E. degree in electrical engineering and engineering physics from Princeton University, Princeton, NJ, USA, in 1979 and the M.S.E.E. and Ph.D. degrees in 1981 and 1985, respectively, both from Stanford University, Stanford, CA, USA.

In 1979, he joined Intel Corporation, Santa Clara, CA, USA, as a Microprocessor Design Engineer, and in 1981 he was a Visiting Engineer at Siemens, Munich, Germany. In 1986, he joined the faculty of Princeton University, where he is currently the Stephen R. Forrest Professor in Electrical Engineering. From 1998 to 2015, he was the director of the Princeton Photonics and Optoelectronic Materials Center (POEM) and its successor, the Princeton Institute for the Science and Technology of Materials (PRISM). In 1994–1995, he was a von Humboldt Fellow at the Institut fuer Halbleitertechnik at the University of Stuttgart, Germany. He has worked in the fields of silicon-based heterojunctions, thin-film and flexible electronics, photovoltaics, the nano-bio interface, three-dimensional (3D) integration, and silicon-on-insulator.

Dr. Sturm has won over ten awards for teaching excellence and was a National Science Foundation Presidential Young Investigator. In 1996 and 1997, he was the technical program chair and general chair of the IEEE Device Research Conference, respectively. He served on the organizing committee of IEDM (1988 to 1992 and 1998 to 1999), having chaired both the solid-state device and detectors/sensors/displays committees. He has served on the boards of directors of the Materials Research Society and the Device Research Conference, and co-founded Aegis Lightwave and SpaceTouch.

Naveen Verma (S'04–M'09) received the B.A.Sc. degree in electrical and computer engineering from the University of British Columbia, Vancouver, Canada in 2003, and the M.S. and Ph.D. degrees in electrical engineering from Massachusetts Institute of Technology, Cambridge, MA, USA, in 2005 and 2009 respectively. Since July 2009 he has been with the Department of Electrical Engineering at Princeton University, Princeton, NJ, USA, where he is currently an Associate Professor.

His research focuses on advanced sensing systems, including low-voltage digital logic and SRAMs, low-noise analog instrumentation and data-conversion, large-area sensing systems based on flexible electronics, and low-energy algorithms for embedded inference, especially for medical applications.

Prof. Verma serves on the technical program committees for ISSCC, VLSI Symposium on Circuits, and IEEE Signal-Processing Society (DISPS). Prof. Verma is recipient or co-recipient of the 2006 DAC/ISSCC Student Design Contest Award, 2008 ISSCC Jack Kilby Paper Award, 2012 Alfred Rheinstein Junior Faculty Award, 2013 NSF CAREER Award, 2013 Intel Early Career Award, 2013 Walter C. Johnson Prize for Teaching Excellence, 2013 VLSI Symposium Best Student Paper Award, 2014 AFOSR Young Investigator Award, 2015 Princeton Engineering Council Excellence in Teaching Award, and 2015 IEEE TRANSACTIONS ON COMPONENTS, PACKAGING AND MANUFACTURING TECHNOLOGY Best Paper Award.