(Keynote) Machine Learning and High-Speed Circuitry in Thin Film Transistors for Sensor Interfacing in Hybrid Large-Area Electronic Systems

ECS Transactions, 92 (4) 121-134 (2019)
10.1149/09204.0121ecst ©The Electrochemical Society

Machine Learning and High-Speed Circuitry in Thin Film Transistors for Sensor Interfacing in Hybrid Large-Area Electronic Systems

J.C. Sturm, Y. Mehlman, L.E. Aygun, C. Wu, Z. Zheng,
P. Kumar, S. Wagner, and N. Verma

Department of Electrical Engineering, Princeton University,
Princeton, New Jersey 08540, USA

The advent of flexible substrates with thin film transistors (TFTs) over large areas (meters) makes large-area electronics (LAE) an attractive platform for integrating very large numbers of sensors onto surfaces over large areas. While TFT's may directly interface to sensors and may be used for sensor addressing, to realistically communicate with the outside world, IC's will probably be bonded onto the "sensor sheets" to create a "hybrid" LAE/IC system. This paper examines novel architectures to minimize the number of physical interfaces to the IC, beyond the typical TFT-based active-matrix approach. Approaches demonstrated include (i) high-frequency TFT-based analog oscillators, and (ii) implementing elements of machine learning into TFT circuitry, so a higher-level information is sent to the IC's, thus requiring fewer physical connections.

Introduction

Large-Area (and flexible) Electronics (LAE) on substrates such as plastic is a technologically attractive platform for deploying large numbers (potentially millions or billions) of sensors over surfaces, from strain sensors on bridges and airplanes to flexible neural sensing caps worn on the head [1]. Eventually the sensor information must reach conventional IC's for external communication and system integration. For sensor sheets on the order of meters in size, sending all data by wires to the connectors on the edge of the sheet could be problematic. Rather, it is attractive to bond IC's onto the sheet, resulting in a "hybrid electronic" system (Fig. 1). With large numbers, clearly each sensor cannot be directly connected to an IC, for cost and practical reasons and also perhaps the reliability issues due to the substrate flexibility. The classic approach to reduce the number of physical interconnects from the large area domain to the IC domain is active-matrix addressing enabled by low-performance thin film transistors (TFTs) fabricated directly on the plastic, architecturally similar to many conventional imaging arrays.

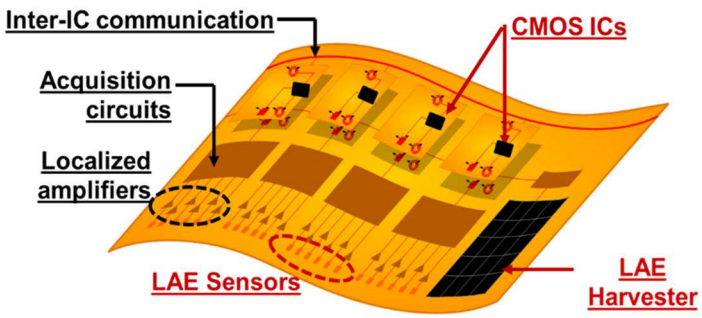

Figure 1. Concept of flexible hybrid LAE sensing sheet with embedded IC's and integrated TFT circuitry.

Our work at Princeton has focused on using novel circuit, algorithmic, and machine-learning approaches implemented in the TFT/LAE domain, to greatly further reduce the number of hard physical connections between the TFT/LAE domain and conventional IC's. (Fig. 2). The work has been demonstrated by prototype hybrid electronic sensor systems ranging from structural health monitoring [2-3], remote gesture sensing [4-5], brain wave (EEG) sensing [6], recognition of handwritten digits [7], pattern recognition of physical objects by weight distribution [8-9], to isolation of individual speakers (when many are speaking) via microphone arrays [10].

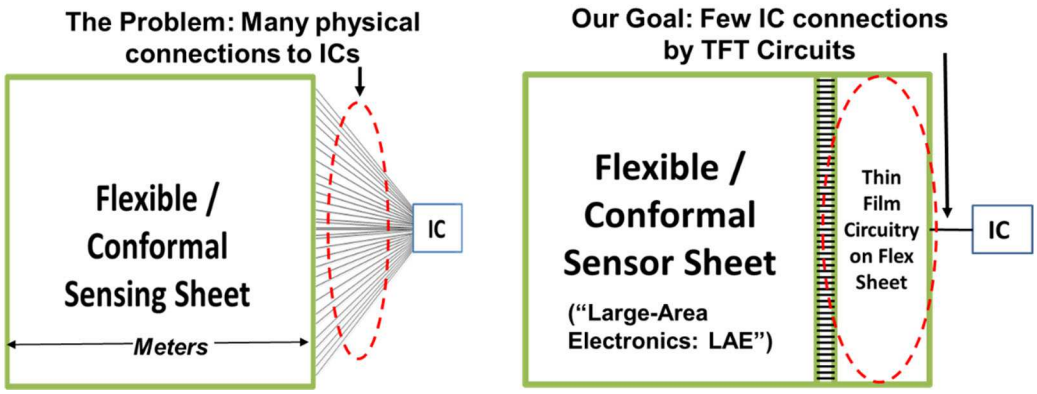

Figure 2. Illustration of the problem of many physical connections from many sensors to the IC (or ICs), and approach of this work to use TFT circuitry to reduce the number of connections.

A first set of work is based on the analog properties of TFT's, especially via high-frequency TFT oscillators. While large feature sizes and low process temperatures typically lead to "low performance" TFT's, with simple self-alignment methods we have extended ZnO TFT performance [11-12] to currently well over 2 GHz. Along with HF thin film diodes [13], this enables near field coupling of signals and power between adjacent large area sheets [14, 15] and sending sensor data "off-sheet" by wireless methods [5]. With sufficient bandwidth, frequency-hopping (spread spectrum) methods can be used to send far more data for a fixed number of connections than with a conventional active matrix approach [16].

Second, in many real-world problems, the goal of sensor arrays is often some kind of pattern recognition, either the in time domain (e.g. seizure detection via neural sensors) or in the spatial domain (e.g. initial detection of a crack on a wing). Machine learning (ML) is a powerful tool for this, conventionally done in software. This requires sending all of the raw data from the sensors to the CMOS ICs and perhaps then to the cloud. Instead, we have implemented various levels of machine learning on the sensor data directly in TFTs in the large-area domain, so that only the results are sent to the IC domain portion of the system. ML conceptually consists of two steps: (i) feature extraction and (ii) classification. Feature extraction can be implemented in several ways using digital TFT circuits. Successful classification requires "memory" to be integrated into TFTs to save the "learning." We have directly implemented both steps in TFT circuits for problems such as

handwriting recognition from images in amorphous silicon sensors arrays, seizure detection, and tactile sensing [6-9].

Thin Film Transistor Technology

Two TFT technologies were used in this work presented in this paper. Both are bottom-gate (top source/drain contact) approaches similar to those widely used in the display field. The first was an amorphous silicon (a-Si) technology with a maximum process temperature of 300 ºC. The gate insulator was typically 300 nm of SiN_x, on top of which the channel layer of ~100 nm a-Si was deposited. N-type source/drain layers were deposited and patterned on top of the a-Si and the metal source/drain contacts were on top of the n-type layers. The insulator and semiconductors were deposited in an a load-locked multi-chamber plasma-enhanced chemical vapor deposition system (PECVD). The gate layer typically was ~100 nm Cr, although sometimes a Cr/Al/Cr sandwich was used to lower the gate resistance. The upper Cr layer was used to reduce hillocks that typically form on Al surfaces, leading to reduced breakdown voltages on the gate insulator. The back (top) of the channel was passivated with SiN_x.

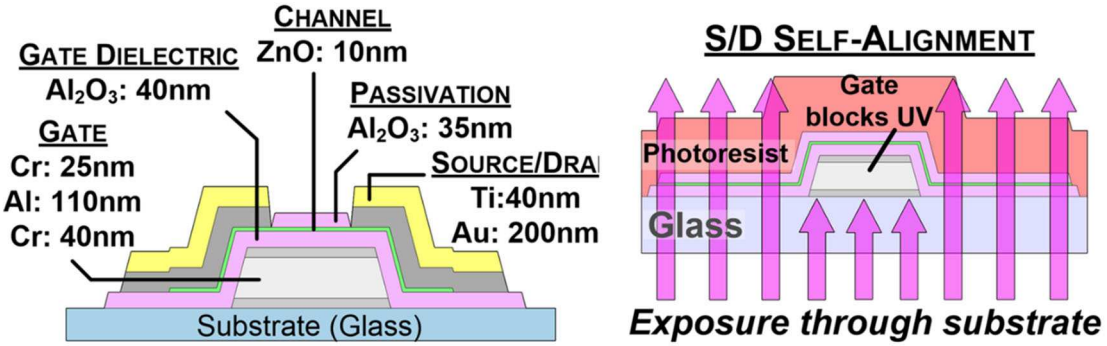

(a) (b)

Figure 3. (a) Cross section of ZnO TFT's with maximum process temperature of 200 ºC (compatible with fabrication on plastic), and (b) approach towards self-aligning the metal source/drain contacts to the gate metal to reduce overlap capacitances [11].

The second technology was similar in cross section (Fig. 3a) but used a thin (10 nm) polycrystalline ZnO layer as the semiconductor. A 40-nm Al_2O_3 layer served as the gate insulator, and the underlying gate was typically Cr or Cr/Al/Cr as with the a-Si TFT's. No intentionally-doped semiconductor was used, and the source drain contacts were made by the direct deposition of Ti/Au on the ZnO. Al_2O_3 was used for top passivation of the ZnO as well. Both the Al_2O_3 and the ZnO were deposited by plasma-enhanced atomic layer deposition (PEALD) at 200 ºC. While the work presented in this paper was performed on glass substrates, because of the low maximum process temperatures, the fabrication processes for both the a-Si and ZnO TFT's are compatible with fabrication on plastic.

The two types of TFT's had similar threshold voltages (1-2 V), but the ZnO TFT's had a channel mobility (evaluated by the drain saturation current at ~ 6V gate voltage) 10-20X larger than that of the a-Si TFT's (0.5-1 vs. 10-15 cm^2/Vs) (Fig. 4a). Typical channel lengths for the TFT's were 3-5 um, and the FET's had large gate-drain capacitances because the lateral location of the top drain contacts was set by a separate lithography step

than the one determining the location of the underlying gate, necessitating a gate much wider than the channel. For high performance, in some devices the source/drain contacts were "self-aligned" to the gate by exposing the photoresist which defined the source/drain from the bottom, through the substrate, so that the light was blocked by the opaque gate layer (Fig. 3b). By thus reducing the overlap capacitances, high frequency performance in the ~5 MHz range for the a-Si TFTs and in the 1 GHz range for the ZnO TFT's was achieved. The higher ZnO TFT performance comes from a factor of ~10 for the higher mobility, and another factor of ~10 from shrinking the channel length down to ~1 um.

Parameter	a - Si	ZnO, IGZO
Mobility (cm^2/Vs)	0.5 -1	10 - 50
V_T (V)	1 - 2	1 - 2
Channel	n-type	n-type
Cutoff freq. f_t	~5 MHz	~1 GHz

(a)

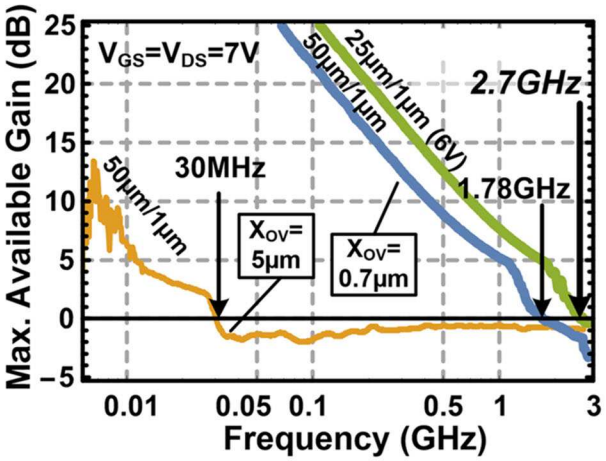

(b)

Figure 4. (a) Typical parameters of the TFTs used in this work, and (b) high frequency (f_{max}) performance of ZnO TFTs using self-alignment [11, 12].

Analog Oscillators for Wired and Wireless Data and Power Transmission

The usual approach to capture data from many sensors is to send the analog data directly by a wire from each sensor to an IC for amplification and digitization. Active matrices can reduce the number of connections to ICs (e.g. scanning one row at a time), but connections for each column are typically required. We now explore how TFT oscillators can be used to reduce the number of "large-area to IC" physical wire connections. The work is largely based on "cross-coupled" TFT LC oscillators, a simplified equivalent circuit model of which is shown in Fig. 5a. The capacitors are typically dominated by those of the TFT's, and the inductors are made using large-area patterned metal layers. The oscillator frequency f is simply $\sim \frac{1}{2\pi\sqrt{LC}}$. Oscillation requires power gain at this frequency, which is reflected in the f_{max} parameter of Fig. 4(b), in turn requiring low parasitics such as gate line resistance.

A first sensor system example is a large flexible sheet (40 cm x 40 cm) to designed to remotely sense gestures, a sort of "touch-pad at a distance" (Fig. 5b) [5]. An array of 5-cm electrodes senses the capacitance to the user's hand, which changes with the position and distance (10's of cm) of the hand. Each electrode adds to the capacitance of a TFT oscillator, so that the oscillator frequency depends on the hand distance. Rather than directly connecting each oscillator output to an IC for measuring frequency, note that the inductor radiates a time-varying magnetic field. A pick-up loop wire is placed around the

array of sensing oscillators, each of which is tuned to a different base frequency by choice of inductor. One row of oscillators is turned on a time, and a custom IC can measure the amplitude of each oscillator simultaneously and independently. Row control signals are sent from the IC to the sensing sheet, but all of the data is transmitted to the IC in a wireless fashion, without any physical wired connection between the capacitance-sensing oscillators and the IC.

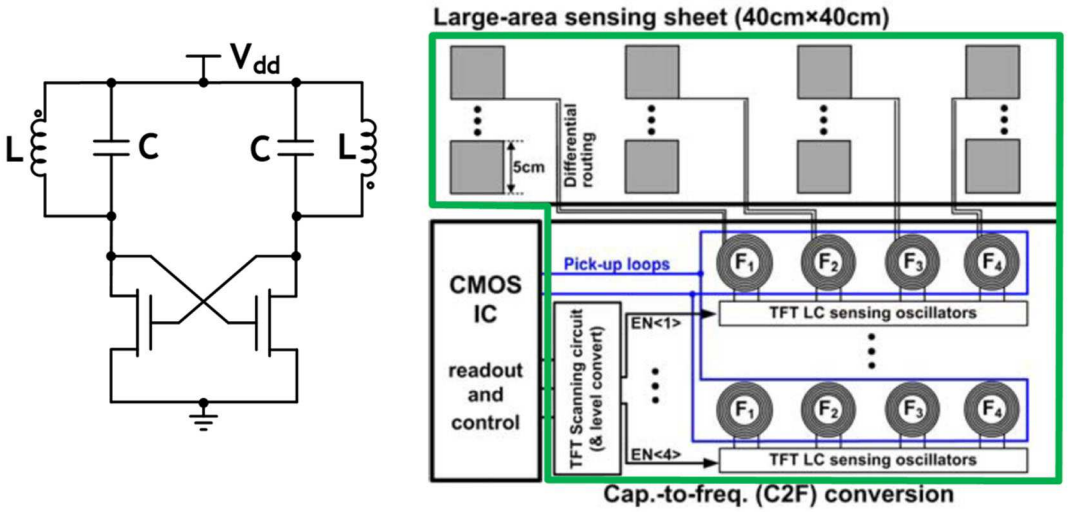

(a) (b)

Figure 5. (a) Schematic diagram of TFT-based LC oscillator, and (b) remote-gesture sensing sheet using LC oscillators to sense capacitance changes by frequency changes, and to inductively send sensed signal to and external pick-up loop [5].

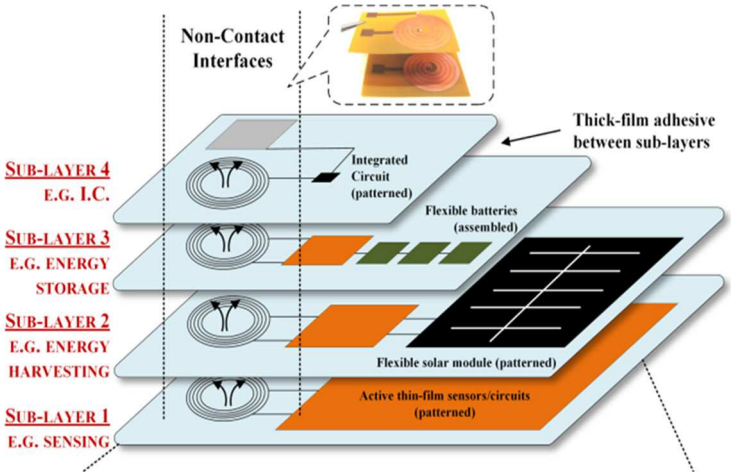

Figure 6. Use of oscillators and inductive coupling to send power and/or data from one sheet to an adjacent laminated sheet [14, 15].

The use of TFT oscillators can be used not just to send data "off-sheet" to an IC for data acquisition, but also from one sheet to another. Rather than integrating all system components (sensors, power sources, etc.) onto a single flexible sheet, one attractive approach for fabricating large-area systems is to fabricate each subsystem on a different sheet, and then laminate the sheets together to form the final system (Fig. 6) [14,15]. One then has to be able to send power and signals/data from one sheet to another. Physical wired connections, such as "bump bonds" between sheets to another could be problematic

from both cost and reliability points of view, especially in flexible systems. Thus we have used TFT oscillators to send both power and data from one sheet to another, using either inductive coupling (as shown in Fig. 6) or capacitive coupling. For example, one sheet of flexible solar cells could harvest energy, which can then be sent to other sheets. A critical point, however, is that a non-linear component (e.g. a thin-film diode) is needed to convert the AC power and data to DC signals (e.g. 5V power supplies or 0/5V digital data). Thus one must also fabricate thin film diodes [13]. With such an approach, DC to DC power transmission efficiencies in excess of 80% between sheets can be achieved.

As a last application of TFT oscillators, we examine a frequency-hopping system for acquisition of analog sensor data to potentially greatly reduce the number of data connections from a system of many sensors on a large-area sheet to the realm of ICs and digital computers. Consider again a cross-coupled LC oscillator, to which a "tail" TFT has been added, with its gate connected to an analog sensor signal (in this case the voltage from a resistive pressure sensor) (Fig. 7a). The sensor signal modulates (controls) the amplitude of the oscillation, so that the data is encoded in the oscillator amplitude, not in its frequency as in the earlier gesture-sensing example. (AM stands for Amplitude Modulation, as in AM radio.)

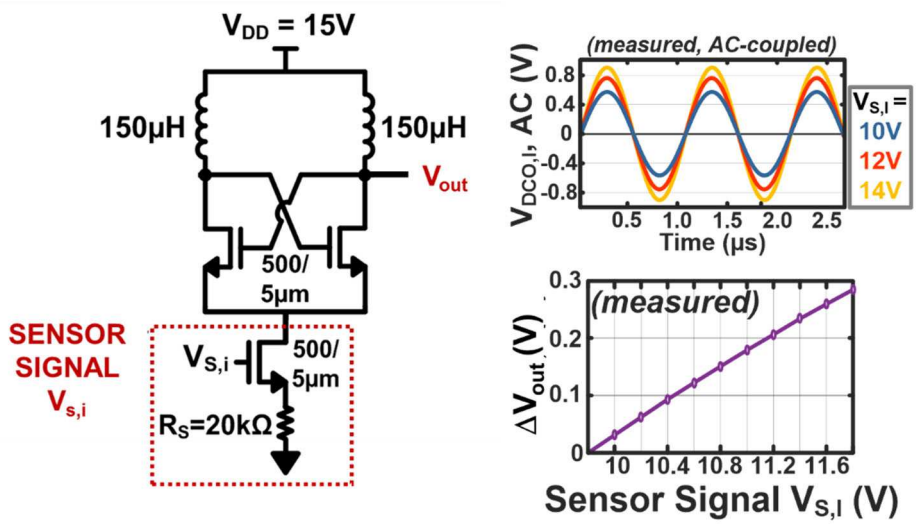

(a) (b)

Figure 7. (a) TFT circuit for AM (amplitude modulation) of resistive pressure sensor data ($V_{S.i}$) on top of 1 MHz carrier frequency, and (b) linearity of amplitude modulation with sensor data. In (a), the capacitors are omitted for simplicity [16].

To send data from the sensors to an IC, each oscillator can be given its own frequency (by design of C and/or L), and then a summing junction can be used to add the current output signals of all of the oscillators, so that only one wire goes to the IC. The IC can then independently recover the independent sensor signals since each one is at a different frequency (Fig. 8a). While attractive, this would limit the number of sensors per wired connection to the IC to the number of possible frequencies.

To overcome this limit, a single tunable oscillator design was developed (Fig. 8b). Digital control signals from the IC are used to switch in or out capacitors on each oscillator, thus tuning its frequency. For example, with 3 control signals, $2^3 = 8$ frequencies and 8

separate data channels are possible. To expand the number of independent sensor channels with this architecture, the digital frequency control codes were changed as a function of time. The control signals were wired differently (and uniquely) to the FET switches in each oscillator. Thus the frequency changes in a unique pattern as function of time in each oscillator. With advanced signal processing techniques, easily implemented in an IC, the oscillator magnitude and thus the data from each sensor can be recovered. Such a technique is known as "frequency hopping" or "spread spectrum." For example, with 5 digital control signals, 32 different frequencies can be selected from the circuit in Fig. 8b, leading to data from 32 sensors going to the IC on a single physical connection if each of the 32 oscillators had a fixed frequency. With the frequency hopping method, it can be shown that this number increases combinatorically, from 32 to ~350. With more digital controls signals, the factor of increase is substantially larger [16].

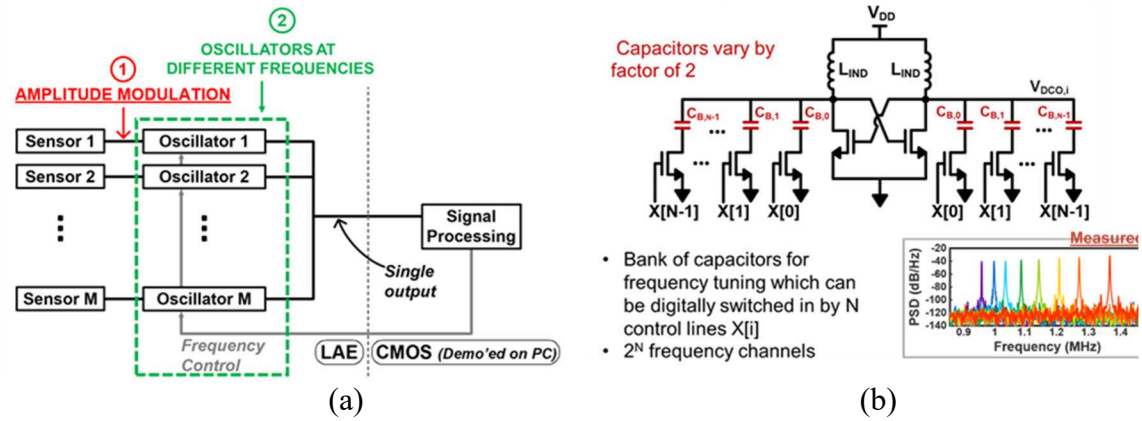

Figure 8. (a) Use of different TFT oscillator frequencies (one for each sensor) to send data from all sensors superimposed to CMOS IC's over a single data line, where they can be independently recovered, and (b) digital-control tuning of frequency of each oscillator by switching in different capacitors. (Sensor modulation TFT's are not shown for simplicity.) By changing the frequency of each sensor's oscillator vs. time (frequency-hopping), the number of addressable sensors for a fixed number of digital frequency control signals X[i] goes up dramatically [16].

Implementation of Machine Learning Algorithms into TFT Circuitry

Up to this point this paper has focused on novel approaches to send data from sensors to an IC with a minimal number of physical connections to the IC, through the use of novel TFT circuits, especially oscillators. The focus has been on sending all of the sensor data to the IC, which may be subject to bandwidth limitation, and in any case requires power. In many applications, however, we are not interested in all of the sensor data per se, but rather patterns in the data. Such patterns will tell us if a crack is beginning to form in an airplane wing, what sorts of hand gestures are being made, if EEG brain waves are indicating the onset of a seizure, and so forth.

In these applications, "machine-learning" algorithms executed in software have recently become popular for recognizing patterns. We now focus our attention on developing hardware circuits of TFT's, rather than software programs, for implementing portions of machine-learning algorithms on sensor data to recognize patterns. If successful, only high level information, such as if a crack on a bridge is forming [16], is then sent to

the ICs and computer side of the system. That will of course require far fewer physical connections and less interconnect bandwidth and power than sending every piece of sensor data to the IC's.

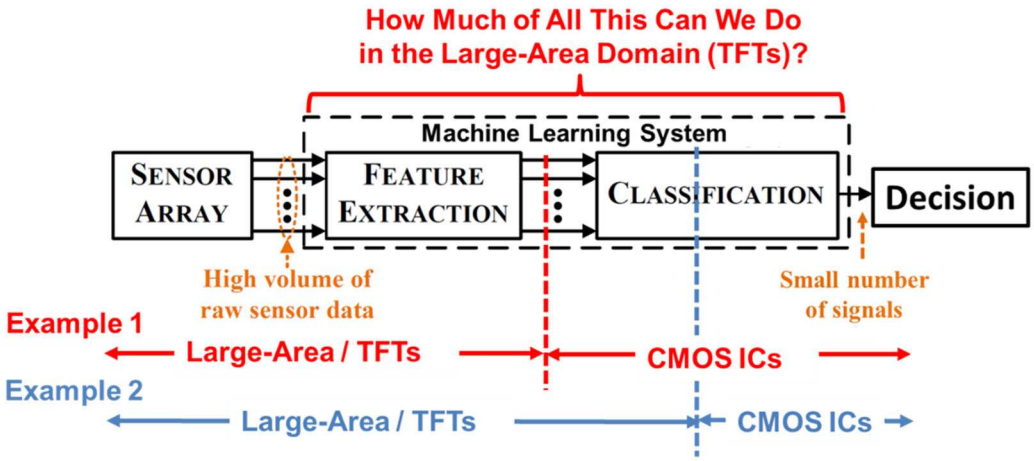

Figure 9. Conceptual stages of machine learning (feature extraction and classification) between sensor data and ultimate decision. We seek to implement these functions in hardware in TFT circuits to reduce the number of physical connections to the IC /computer domain and similarly reduce the amount of data being transferred to the IC/computer domain (adapted from [6]).

On a high level, the problem can be represented by the diagram of Fig. 9. Machine learning focused on pattern recognition can be represented two steps. The first step is to recognize certain simple features in the data, with an output of how strong the data matches each of these simple features. Certain types of data compression, such as compressed sensing, which reduces that data from N independent signals (where N is the number of sensors, potentially a very large number – e.g. millions) to a small number M of analog signals, one for each feature can serve this role (M is typically 5-10 in our work). After feature extraction, there is a second step of "classification," which involves looking at the different signals for each feature and deciding if a certain overall *general* pattern exists or not. We now examine two ML platform demonstrations, the first of which performs the feature extraction step in TFT's and the classification step in an IC, and the 2nd of which also performs part of the classification in TFT's as well.

The first system was designed to recognize handwritten digits (Fig. 10a) [7]. The digits were projected onto an array of a-Si photosensors. The data from each row of the data (N_C columns) was then compressed into M separate outputs by what effectively is a linear matrix multiplication step. Because the data is sparse (meaning mostly 0's, meaning no handwriting), it can be shown that a matrix coefficients of random +1's and -1's (random projection compression) (Fig. 10b) preserves key features in the data, and can thus be thought of as a type of feature extraction (the first step in Fig. 9). Because the key features are mathematically preserved, a high level of classification performance is possible directly from the compressed sensor data.

Figure 11a shows the TFT circuit used to implement the data compression step. On a row by row basis, the data of each sensor column is multiplied by 1 or -1, depending on the matrix coefficients. For each of the M compressed outputs, the TFT contributes current proportional to the sensor data to one of two horizontal lines, depending on the +/-

1 matrix coefficient. A custom circuit on the IC differences these two currents to get the compressed data.

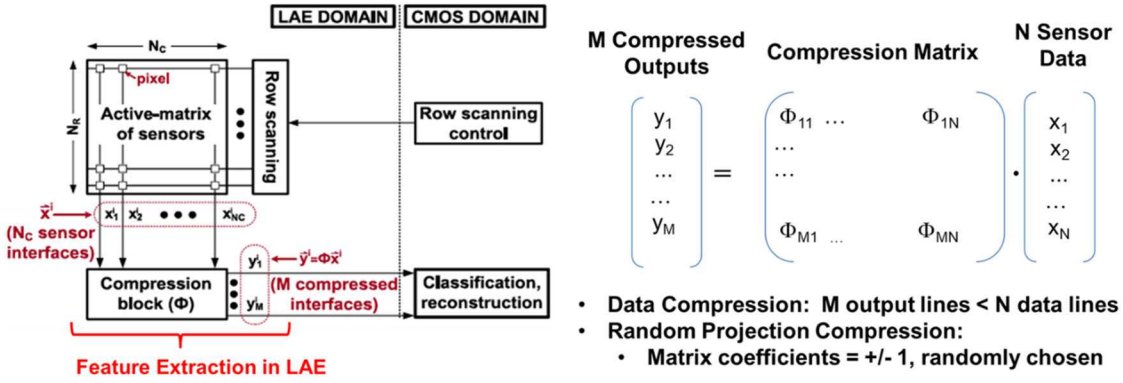

(a) (b)

Figure 10. (a) Architecture for performing data compression using TFT's in Large-Area electronics to reduce N_C columns of data to M compressed signals, and (b) mathematical operation performed in hardware implementing matrix multiplication by +/- 1 (random projection compression) [7].

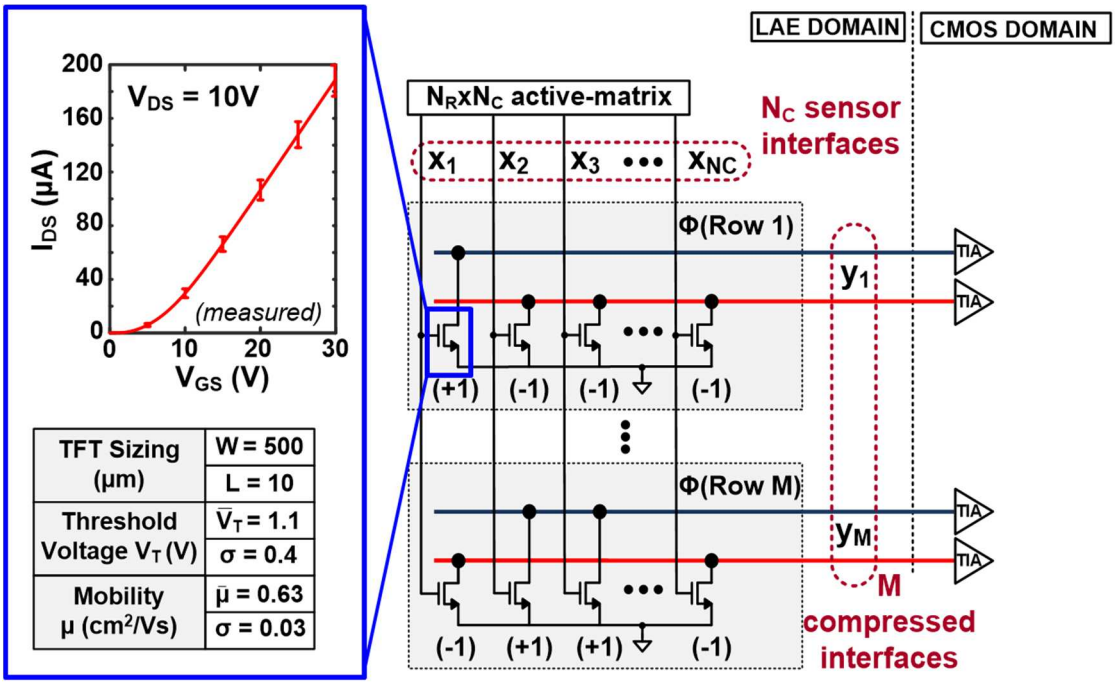

Figure 11. a-Si TFT characteristics and the TFT circuit used to implement matrix-multiplication compression algorithm of Fig. 10 in the large-area domain. Each compression block puts out two current signals, with one subtracted from the other after the signals go from the large-area domain to a CMOS and a transimpedance amplifier (TIA) [7].

For the step of the generalized machine learning system of Fig. 9, in this example, "one-versus-all" classification is performed for each digit using a SVM (support vector machine) after the large-area (TFT) signal processing – e.g. in the realm of IC's and/or

computers. For ease of testing, a MATLAB-implemented SVM classifier is used; however, such a classifier can be readily integrated in a CMOS IC on a hardware level [17]. Critically, successful classification relies on the ability to train the classifier, in this case directly from the compressed data.

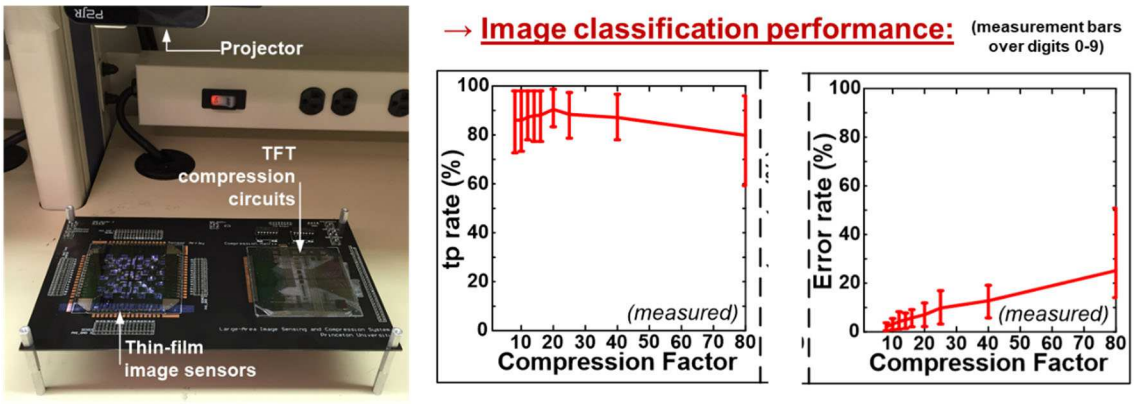

Figure 12. (a) View of image sensor array and TFT compression circuits, and (b) results of positive detection and error rate for recognition of handwritten digits 0-9, as a function of compression factor = N_C/M = 80/M [7].

Figure 12a shows the a-Si image sensor array and the a-Si TFT compression circuit shown schematically in Fig. 11. Fig. 12b shows the performance of the image classification step in recognizing handwritten digits 0 through 9, as a function of the compression factor in the data compression (feature extraction) step. The compression factor CF is defined as $CF = \frac{N_C}{M}$, where N_C = 80 (columns of the imager) and M is the number of extracted features (rows in the compression matrix of Fig. 10) – how many signals are actually sent to the IC for that row. Surprisingly, the true positive rate is consistently 80% of higher, even for a compression factor of 80 (meaning only one compressed output going to the IC for each row of the image). The error rate, on the other hand, increases with the compression factor, but only reaches ~20% for its maximum value of 80.

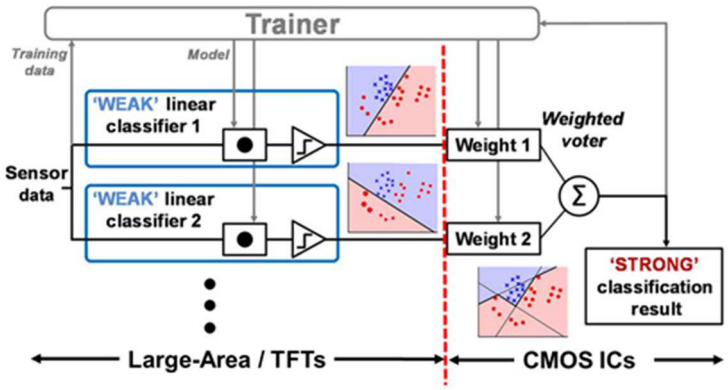

Figure 13. Architecture of "weak classification," including feature extraction, performed by TFTs. It consists of a matrix multiplication of the sensor data for each weak classifier output as in the data compression circuits of Fig's. 11 and 12 [8]. However, in this case the matrix coefficients are not simply +/- 1, and have values determined by learning

algorithms and feedback. The final decision is made by a weighted voter implemented in an IC.

We now examine a second pattern recognition system, employing machine learning implemented in TFTs, in this case a system to recognize and classify shapes, not handwritten digits, from their images, shown architecturally in Fig. 13. As in the previous case we start with images captured by an a-Si array, and use TFT hardware to implement matrix multiplication on each column of date. In this example the TFT hardware circuits will implement part of the classification step (part of the ultimate decision making), not just perform data compression/feature extraction. Architecturally, two things differentiate this sytem from the first example (Fig. 13). First, the matrix coefficients are not just +/-1, but are weighted according to a training algorithm, to achieve a so-called "weak classification" result. This is achieved by a series combination of TFT's, one with a gate voltage determined the the sensor output (V_{Si}) and the other with the gate connected to a stored "weight" value V_{Bi} (Fig. 14a).

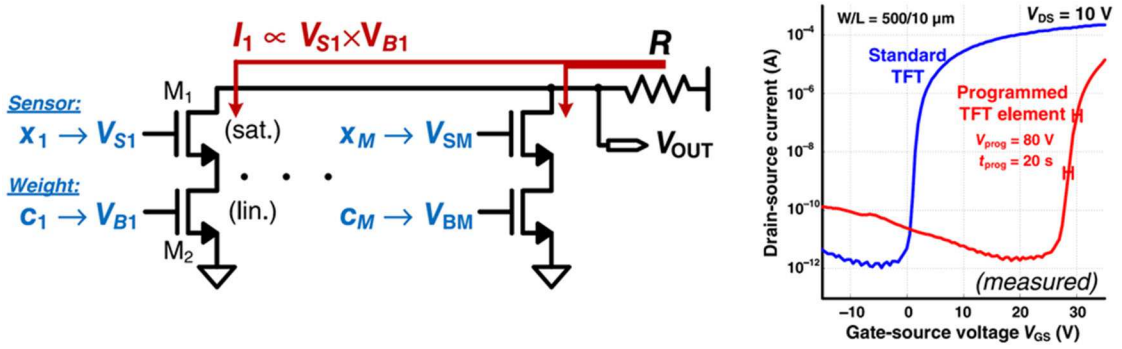

(a) (b)

Figure 14. (a) TFT circuit used to implement analog matrix-multiplication for one output of the weak classifier of Fig. 13. V_{S1} is the sensor signal and V_{B1} is the "weight" coefficient which is selected by training. (b) The weight coefficient is stored as the threshold of the TFT. The threshold is adjusted to a desired value by large gate voltage pulses to inject electrons into the gate dielectric. Amorphous silicon TFT's with a SiN_x gate dielectric were used for this example [8].

The value for each of the N_C x M weights is chosen by a training step. This means optimizing the weights based on presenting images of known shapes to the system. Because the learning is taking place on the level of TFT's, and not on an IC or software level, the "learning" must be somehow physically stored in the large/area TFT domain. This was effectively achieved by adjusting the threshold voltage of of the lower ("weight") TFT's in the circuit of Fig. 14a (and holding the actual gate voltage of all of the lower TFT's fixed when processing data. As an outcome of the learning step, a selected number of high gate voltage pulses where applied to each of the lower TFTs. This caused the injection of electrons into the gate dielectric, some of which get trapped to raise the TFT threshold voltage. Because the electrons are trapped, this adjusted threshold voltage is then "stored."

These TFT classifiers are referred to as "weak" because they do a poor job of recognizing complex patterns. After the weak classifier step, the M weak classifier outputs where individually sent to an IC. The final "strong classification" step was performed in

software using a "weighted voter" approach based on these M inputs, following the algorithm known as adaptive boosting (AdaBoost) [18]. Results for recognizing five different shapes are shown in Fig. 15. After training, the ability of the system to correctly recognize (true positive, tp) and reject (true negative, tn) each shape type is shown in Fig. 15 in blue. While the success rates depend somewhat on shape type, note that optimal results are achieved already using only 3 weak classifiers, i.e. three rows in Fig's. 13 and 14, representing only 3 signals going to the IC (M = 3). Also shown in red in Fig. 15 are results from a classical support vector machine (SVM) processing of the images in software. In most cases, 3 weak classifiers implemented in TFT's are sufficient for the quality of the pattern recognition to approach that of the conventional software based machine learning.

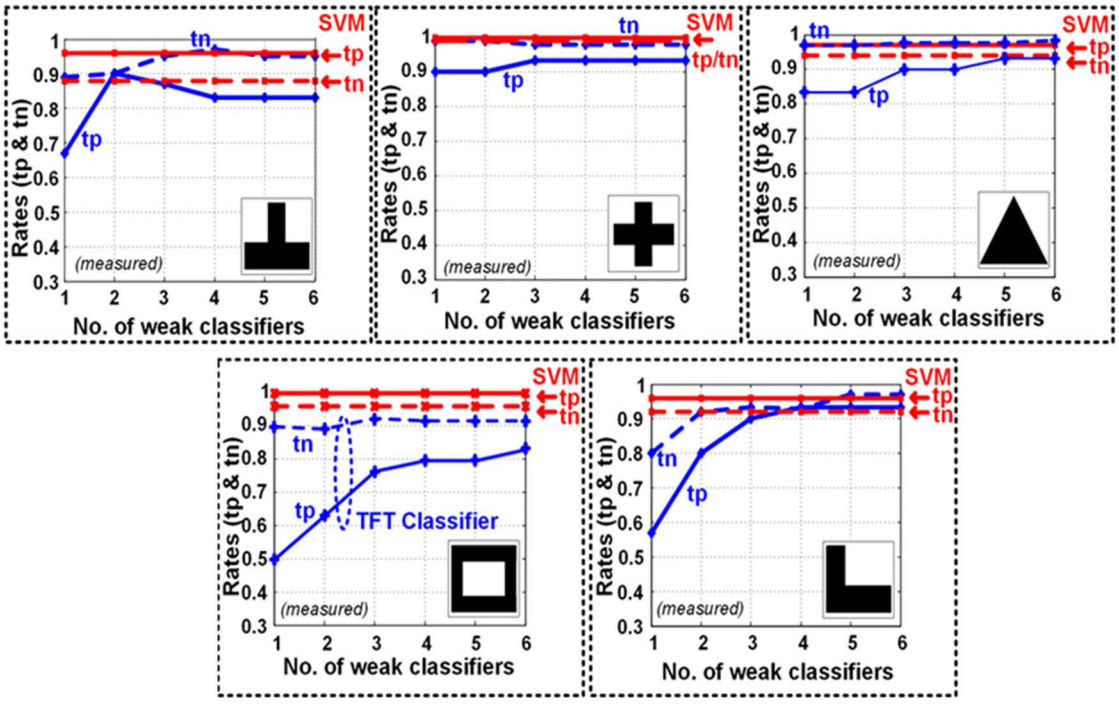

Figure 15. (a) Ability of the weak classifiers in the TFT domain (coupled with weighted voter in CMOS) to recognize shapes (insets), which were projected onto an array of a-Si detectors. The results (blue) improve with the number of weak classifiers, and are compare favorably to a conventional SVM machine learning algorithm performed on the images in software (red) [8].

Summary

Flexible hybrid TFT/IC systems are an attractive physical platform for the implementation of large numbers of sensors over areas on the size scale of people or larger. This paper has examined the use of novel TFT circuits to reduce the number of physical connections between the "large-area domain" (the distributed sensors and TFTs) and the "conventional computer domain," represented on a physical level by IC's. One enabling technology is oscillators implemented in TFT's, which can wirelessly transmit data and power over short distances, and also enable novel data encoding techniques such as amplitude modulation and/or frequency hopping. A second enabling approach is the use

of TFT circuits to execute machine-learning algorithms directly in TFT hardware, so that only higher-level information is sent to IC's and the external world.

Current and future directions include new capabilities enabled by the recent GHz frequency oscillators enabled by ZnO TFT's and self-alignment, compared to the MHz range of frequencies used in the systems examples of this paper. On the machine learning side, physical-related issues are thin film approaches for implementing memory (i.e. "learning"), and the effect of device to device TFT variations (or even faults) and device drift over time, on the ML algorithms implemented in hardware. An important issue in this case is expected to be the nature of the sensor signals themselves and their "sparsity." On a higher architectural level, emerging directions are systems which preserve certain contexts of the data, such as its physical location, and the reconstruction of the original signals from compressed data [8, 9].

Acknowledgments

This work was supported by the Semiconductor Research Corporation, DARPA, NSF, FlexTech, and the Princeton Program for Plasma Science and Technology (PPST).

References

1. N Verma, Y Hu; L Huang, W Rieutort-Louis; J. Sanz Robinson, T Moy, B. Glisic; S. Wagner; J. Sturm, Proc. IEEE 103, pp. 690-712 (2015).
2. Y. Hu, L. Huang, W. Rieutort-Louis, J. Sanz-Robinson, J.C. Sturm, S. Wagner, N. Verma, J. Solid-State Circuits 49, pp. 838-840 (2014).
3. Y. Hu, W. Rieutort-Louis, J. Sanz-Robinson, L. Huang, B. Glisic, J.C. Sturm, S. Wagner, and N. Verma, J. Solid-State Circuits 49, pp. 513-523 (2014).
4. Y. Hu, L. Huang, W. Rieutort-Louis, J. Sanz-Robinson, S. Wagner, J.C. Sturm, and N. Verma, Tech. Dig. IEEE Inter. Solid State Circ. Conf. (ISSCC), pp. 212-213 (2014).
5. Y. Hu, T. Moy, L. Huang, W. Rieutort-Louis, J. Sanz-Robinson, S. Wagner, J.C. Sturm, N. Verma, Dig. IEEE Custom Integ. Circ. Conf. (CICC) M-11 (2014).
6. T. Moy, L. Huang, W. Rieutort-Louis, C. Wu, P. Cuff, S. Wagner, J. C. Sturm, and N. Verma, IEEE Journal of Solid-State Circuits 52, pp. 309-320 (2017).
7. T. Moy, W. Rieutort-Louis, S. Wagner, J.C. Sturm, N. Verma, IEEE Trans. Circuits and Systems I 63, pp. 1833-1844 (2016).
8. W. Rieutort-Louis; T. Moy, Z. Wang; S. Wagner; J.C. Sturm; N. Verma, IEEE J. Solid State Circuits 51, pp. 281-290 (2016).
9. L.E. Aygun, P. Kumar, Z. Zheng, T-S. Chen, S. Wagner, J.C. Sturm, and N. Verma, Dig. IEEE Inter. Solid State Circ. Conf. (2019).
10. J. Sanz-Robinson, T. Moy, L. Huang, W. Rieutort-Louis, Y. Hu, S. Wagner, J.C. Sturm, and N. Verma, IEEE Journal on Emerging and Selected Topics in Circuits and Systems 7, pp. 38-49 (2017).
11. Y. Mehlman, N. Verma, and J.C. Sturm, Tech. Dig. Device Research Conf. (2017).
12. Y. Mehlman, C. Wu, S. Wagner, N. Verma, and J.C. Sturm, Tech. Dig. Device Research Conf. (2019).
13. J. Sanz-Robinson, W. Rieutort-Louis, Y. Hu, L. Huang, N. Verma, S. Wagner, J.C. Sturm, , IEEE Electron Dev. Lett. 35, pp. 425-427 (2014).

14. W. Rieutort-Louis, L. Huang, Y. Hu, J. Sanz-Robinson, S. Wagner, J.C. Sturm, and N. Verma, J. Photovoltaics 4, pp. 432-439 (2014).
15. W. S. A. Rieutort-Louis, J. Sanz-Robinson, T. Moy, L. Huang, Y. Hu, Y. Afsar, J.C. Sturm, N. Verma, and S. Wagner, IEEE Trans. Components, Packaging, and Manufact. Tech. 5, pp. 2019-2029 (2015).
16. Y. Afsar, T. Moy, N. Brady, S. Wagner; J. Sturm, N. Verma, IEEE J. Solid-State Circuits 53, pp. 297-308 (2018).
17. K. H. Lee and N. Verma, IEEE J. Solid-State Circuits 48, pp. 1625–1637 (2013).
18. Y. Freund and R. Schapire, J. Comput. Syst. Sci. 55, pp. 119–139 (1997).