An Architecture for Large-Area Sensor Acquisition Using Frequency-Hopping ZnO TFT DCOs

Yasmin Afsar, *Student Member, IEEE*, Tiffany Moy, *Student Member, IEEE*, Nicholas Brady, Sigurd Wagner, *Fellow, IEEE*, James C. Sturm, *Fellow, IEEE*, and Naveen Verma, *Member, IEEE*

***Abstract*— Hybrid systems combine large-area electronics (LAE) with silicon-CMOS ICs for sensing and computation, respectively. In such systems, interfacing a large number of distributed LAE sensors with the CMOS domain poses a key limitation. This paper presents an architecture that aims to greatly reduce both the number of physical connections and the time for accessing all of the sensors. Each sensor modulates the amplitude of a thin-film transistor (TFT) digitally controlled oscillator (DCO). All DCO outputs are combined, but each follows a unique frequency-hopping pattern (controlled by a code from CMOS), allowing recovery of the individual sensors. The architecture enables much greater scalability of sensors for a given number of connections than active-matrix and binary-addressing schemes. For demonstration, an 18-element large-area force-sensing system is demonstrated based on zinc-oxide (ZnO) TFT DCOs with a frequency-hopping rate of 4.2 kHz. Acquisition error ≤ 62 mV_{rms} is achieved over 30 weight patterns.**

***Index Terms*— Hybrid systems, large-area electronics (LAE), metal-oxide thin-film transistor (TFT), sensing systems, ZnO TFT.**

I. Introduction

LARGE-area electronics (LAE) is a technology that enables unique capabilities for sensing. This is a consequence of its characteristically low processing temperatures (≤200 °C), which promote compatibility with a wide range of materials and fabrication methods. This in turn enables a variety of transducers [1]–[5] fabricated on substrates, such as glass, plastic, and paper, which can be large (square meters) and conform to underlying surfaces. However, a key challenge to the creation of complete systems is the integration of other required functionality (instrumentation, computation, power management, and so on). Thin-film transistors (TFTs) could provide this functionality in LAE, as recent work demonstrating various TFT circuit blocks has shown [6]–[10]. However, low process temperatures cause the performance and energy efficiency of TFTs to be orders-of-magnitude lower than those of silicon-CMOS transistors.

Manuscript received April 15, 2017; revised July 12, 2017; accepted August 3, 2017. Date of publication October 2, 2017; date of current version December 26, 2017. This paper was approved by Guest Editor Pui-In Mak. This work was supported in part by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA. *(Corresponding author: Yasmin Afsar.)*

The authors are with Princeton University, Engineering Quadrangle, Princeton, NJ 08544 USA (e-mail: yafsar@princeton.edu; tmoy@princeton.edu; nbrady@princeton.edu; 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/JSSC.2017.2740383

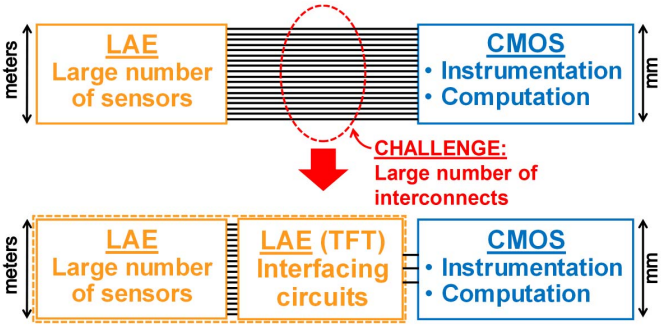

Fig. 1. Hybrid systems combine complementary strengths of LAE and CMOS, but are limited by the interconnections between technology domains, necessitating specialized TFT architectures.

For example, silicon-CMOS transistor f_T values fall between 200 and 300 GHz, while those of LAE TFTs are typically 1–10 MHz [11], [12].

To enable complete systems, hybrid systems have emerged [13]–[15] in which LAE is combined with silicon-CMOS ICs in architectures that selectively delegate functionality between both technology domains. However, research in hybrid systems has shown that a primary bottleneck is the interconnections required between the two technologies [15], particularly when sampling large numbers of distributed transducers. As shown in Fig. 1, specialized TFT architectures can overcome this challenge; TFT active matrices have accomplished this in large-area, flat-panel displays and imagers. However, active matrices yield only a square-root reduction in interconnections and are most conducive to tight, highly regular arrangements of sensors (due to the impact of gate-/data-line capacitances) [16]. Envisioning future systems capable of scaling to very large numbers of distributed sensors, the objective of this paper is a TFT architecture that enables both a much greater reduction in interconnections and a high rate of access over all sensors, despite low TFT performance.

Section II describes the overall architecture, based on an array of amplitude-modulated, frequency-hopping TFT oscillators [17]. Section III focuses on the circuit-level implementation, which exploits high-Q, digitally controlled LC oscillators to substantially increase the signal-processing bandwidth compared to typical TFT circuits. Section IV presents the system prototype, wherein the architecture is applied to a large-area force-sensing system. Finally, Section V provides conclusions.

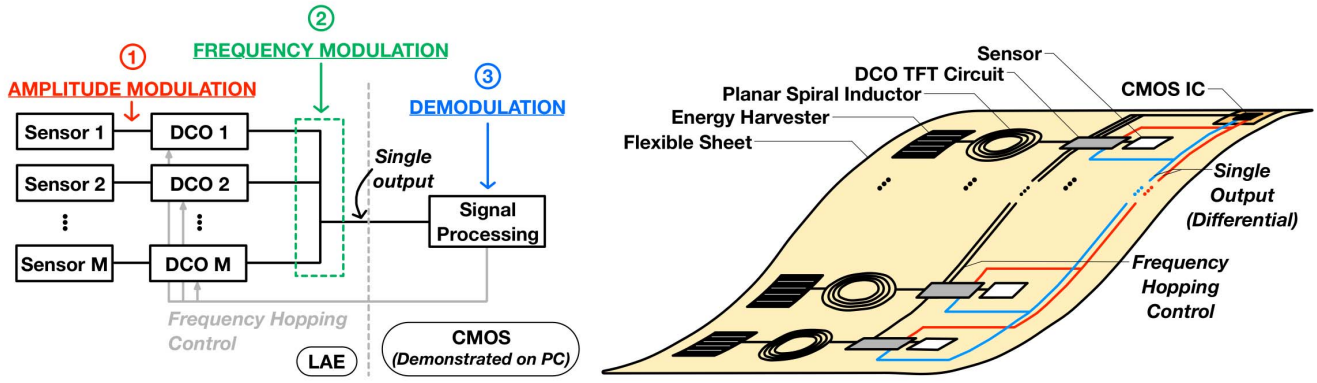

Fig. 2. Left: system architecture. Right: conceptual implementation. System presented consists of an array of sensor-DCO pairs, where DCO amplitude is modulated by sensor signal and frequency is modulated by CMOS frequency-hopping control signal, yielding a single differential interface to CMOS.

II. System Architecture

A. Overview

The system architecture is shown in Fig. 2 as a block diagram (left) and as a conceptual sketch of a fully realized system (right). The LAE domain consists of an array of M sensors, each connected to a TFT digitally controlled oscillator (DCO), giving a single differential interconnection to the CMOS domain. The CMOS domain consists of instrumentation and signal processing for demodulation. The signal from each sensor modulates the *amplitude* of its associated DCO. Meanwhile, a frequency-control signal from CMOS modulates the *frequency* of all DCOs, in a way that enables all sensor signals on the single interconnection to be separated. Each DCO, having N frequency-control bits, outputs to one of 2^N frequency channels. While standard frequency multiplexing would yield a linear increase in the number of sensor signals with the number of frequency channels, a frequency-hopping scheme can yield a much greater increase, and accordingly a much greater reduction of interconnections. While the sensor density is limited by the area occupied by the LAE circuitry, dominated by the DCO inductors as seen in Fig. 2, access to a very large number of distributed sensors is now possible.

Fig. 3 (top) shows the frequency-hopping scheme employed, for the case of $N = 3$ (i.e., eight frequency channels). The hopping-control code $H[N-1:0]$ from CMOS provides N bits. Critically, each of these is hardwired *differently* to the frequency-control inputs $X[N-1:0]$ of the M DCOs. For readout, the DCO outputs are coupled through capacitors C_C to the single differential interconnection to CMOS.

Fig. 3 (bottom) shows how scanning through the values of the hopping-control code causes each DCO to follow a unique frequency-hopping pattern through the frequency channels. The patterns are shown for two DCOs, where the sensor signals are assumed to be fixed for the duration of one hopping-code scan cycle (i.e., fast frequency hopping). As shown, the DCO outputs are summed by transimpedance amplifiers (TIAs) in the CMOS domain for readout. Starting with $H[2:0] = 001$, we see that the frequency-control connections cause DCO 1 to output in frequency channel 6, and DCO M to output in frequency channel 0 (note that the output levels of the DCOs are shown to be different, as set by their sensor signals). Next, with $H[2:0] = 010$, both DCOs output in frequency channel 1, resulting in a temporary superposition of the two sensor signals. But subsequently, with $H[2:0] = 011$, the DCOs again output in different frequency channels.

The unique and predetermined frequency-hopping patterns taken by the sensors over the course of the hopping-code scan enable all of the sensor signals on the single interconnection to CMOS to be separated, despite the sensor superpositions observed above. In fact, the mapping of sensor signals to frequency channels can be represented by a "hopping matrix" $\mathbf{T}$, and separation can be achieved via the matrix equation shown in Fig. 4. In $\mathbf{T}$, each group of eight rows forms a submatrix, corresponding to a particular hopping-code value. Each row within a submatrix corresponds to a particular frequency channel f. Thus, each column of $\mathbf{T}$ corresponds to a particular sensor-DCO pair m, having a nonzero entry $t_{f,\mathrm{DCOm}}$ only if the DCO m outputs in the corresponding frequency channel f, during the corresponding hopping-code value. $\mathbf{T}$ then multiplies with a vector $\vec{s}$, whose elements represent the values of the M sensor outputs. The resultant vector $\vec{y}$ consists of elements $y_{f,H[2:0]}$ that correspond to voltages acquired from the CMOS TIA in each frequency channel f, for each hopping-code value $H[2:0]$. Using the example from above, during $H[2:0] = 001$, DCO 1 outputs in frequency channel 6 and DCO M outputs in frequency channel 0, corresponding to the nonzero entries shown in $\mathbf{T}$.

Given the matrix relationship shown, the sensor data are easily recovered: $\vec{s} = \mathbf{T}^{-1} \times \vec{y}$. However, the inversion of $\mathbf{T}$ is only possible if it is a full-rank matrix. Thus, the rank of $\mathbf{T}$ sets the maximum number of sensors M that can be accessed via the architecture. This is in turn a function of the number of frequency channels, which is set by the number of DCO frequency-control bits N, equal to the number of hopping-code bits from CMOS. In this way, the number of interconnections from CMOS (N) ultimately sets the number of sensors M. Fig. 5 shows this relationship for the architecture (computed based on the matrix rank of $\mathbf{T}$). The many different ways in which the hopping-code bits can be connected to DCO frequency-control bits lead to rapid

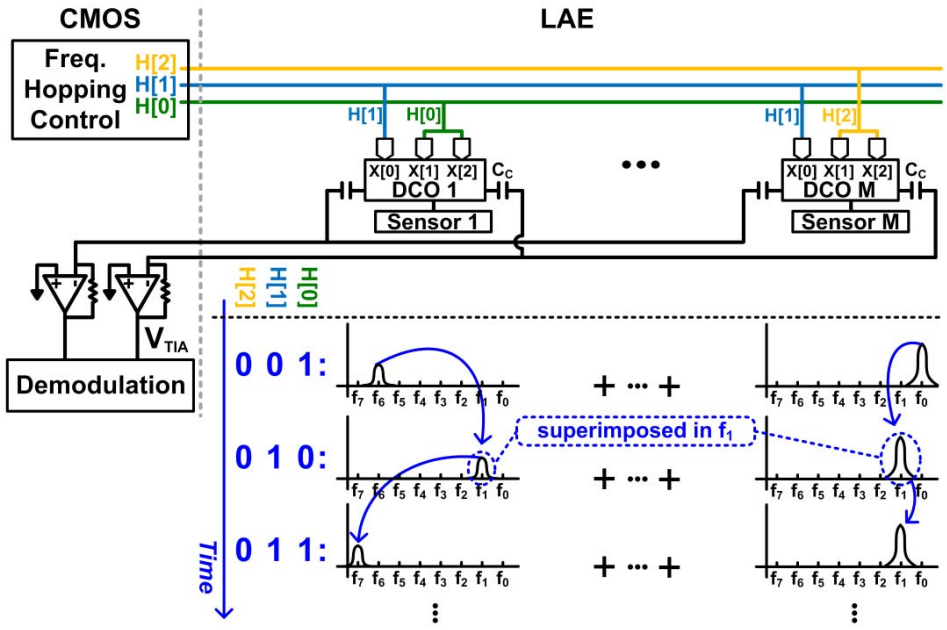

Fig. 3. Unique hardwired connections between CMOS hopping control-code and DCO frequency-control bits give each sensor a unique, predetermined hopping pattern used to recover sensor data.

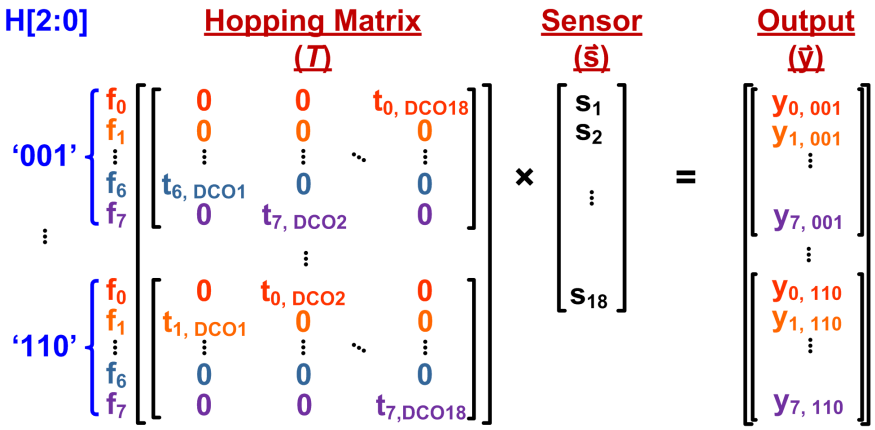

Fig. 4. DCO hopping patterns are represented in matrix **T**, giving a mathematical relation representing sensor signals $\vec{s}$ and CMOS-TIA outputs $\vec{y}$.

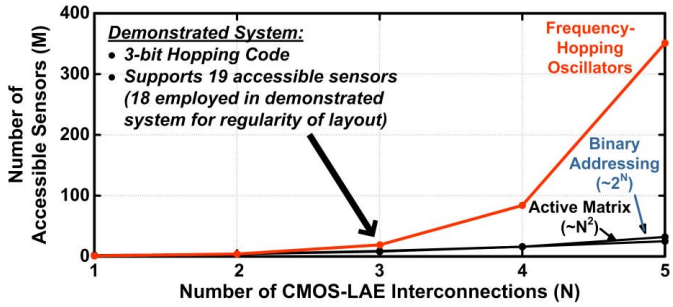

Fig. 5. Comparison of a number of sensors M versus a number of CMOS-LAE interconnections N for the presented and existing architectures.

combinatorial scaling, outpacing active-matrix and even binary-addressing schemes (shown in Fig. 5 for comparison). For example, even for just $N = 5$, frequency hopping can access 351 sensors, a greater than ten times improvement over active matrices (25 sensors) and binary addressing (32 sensors). The prototyped system employs a 3-bit hopping code ($N = 3$), resulting in a hopping-matrix rank of 19, enabling up to 19 accessible sensors. This corresponds to the first point for which the frequency-hopping scheme shows an advantage over the other approaches. In the demonstrated system, 18, rather than 19, sensors are employed ($M = 18$) for greater regularity in the sensor arrangement.

B. Dynamic-Range, Acquisition-Rate, and Circuit-Complexity Tradeoffs

The presented architecture has the potential to greatly enhance the scalability of hybrid systems, but also raises a number of other design tradeoffs related to dynamic-range requirements, sensor-acquisition rate, and circuit complexity.

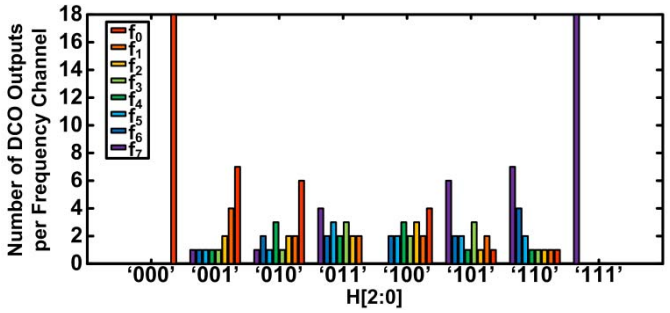

Fig. 6. Distribution of DCO outputs across the frequency channels for different hopping codes.

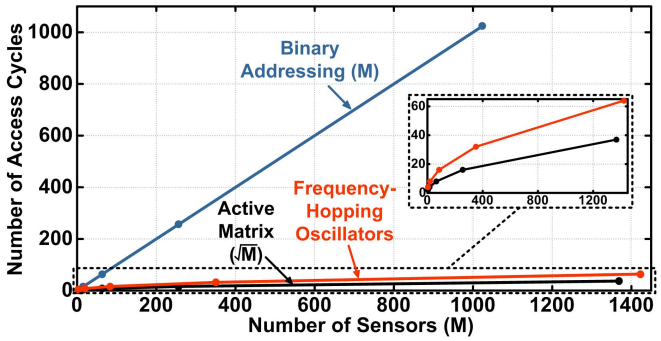

Fig. 7. Comparison of access cycles required in different sensor-accessing schemes.

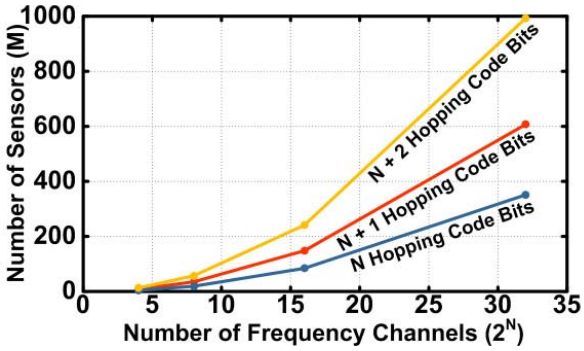

Fig. 8. Scaling in the number of sensors is supported by adding hopping-code bits beyond those needed for a particular number of frequency channels.

Regarding dynamic range, for every hopping-code value, all DCO outputs corresponding to the M sensors are summed by a TIA in CMOS. This requires that the CMOS electronics support a larger dynamic range. However, a key attribute of the architecture is that the increased dynamic range appears only in the CMOS domain (following current summation at the TIA outputs). Given that CMOS circuits can have much higher energy efficiency and performance, and can readily employ architectural enhancements for signal processing (e.g., baseline cancellation to only preserve incremental changes in DCO amplitude from sensor-signal modulation), this is a promising tradeoff.

Additionally, with summation of DCO outputs performed across the 2^N frequency channels, CMOS demodulation can take advantage of architectures that process frequency channels separately. This can substantially mitigate dynamic-range requirements, particularly by selecting hopping codes that result in relatively uniform distributions of DCO outputs across the frequency channels. Fig. 6 shows a histogram of how DCO outputs distribute across frequency channels, for all values of the hopping code in the prototyped $N = 3$ system. Since using all hopping-code values results in more rows in the $\mathbf{T}$ matrix than its rank, many hopping codes can be omitted (e.g., in the demonstrated system, using all eight hopping-code values results in $2^N \times 8 = 64$ rows, but only 18 are needed). As seen, after dropping the values $H[2:0] = 000/111$ (as done in the demonstrated system), most hopping-code values yield much more uniform distributions, reducing the maximum dynamic range. Similar distribution characteristics are seen for systems with larger N.

Regarding acquisition rate, the presented system, as well as binary-addressing and active-matrix schemes, acquires all sensors over a number of access cycles, yielding a frame rate. This is important to consider in addition to the number of interconnections, as it sets the sensor bandwidth that can be supported. Fig. 7 compares the number of access cycles between the different schemes. In the presented scheme, the number of cycles required corresponds to only the number of hopping-code values needed. As seen, this is close to the number of cycles required in an active matrix, yet with far fewer interconnections. As seen, one-by-one sensor accessing with binary addressing yields a very large number of cycles. Furthermore, the presented system could require less time per access cycle, since the DCO output capacitance is absorbed in a resonant tank and no explicit capacitance of a data line needs to be driven (rather, the current is sensed through a TIA virtual ground node).

Regarding circuit complexity, the primary TFT block is the DCO, whose complexity depends on the number of frequency channels supported. While the circuit implementation has significant potential to ultimately scale to many channels, the presented architecture supports scaling in the number of sensors without scaling the number of frequency channels. Specifically, while keeping DCOs with 2^N frequency channels (i.e., N frequency-control bits), the number of hopping-code bits can be increased beyond N. This yields new ways of connecting hopping-code bits to the DCO frequency-control inputs, enabling more DCO-sensor pairs while ensuring they follow unique frequency-hopping patterns. Fig. 8 plots the achievable rank of corresponding $\mathbf{T}$ matrices, showing significant potential to support more sensors without increasing TFT-circuit complexity.

III. System Implementation

A. TFT-Based LC DCO

The heart of the architecture is a TFT-based LC DCO. The DCO's maximum frequency must be as high as possible: maximum frequency and minimum channel separation dictate the total number of frequency channels, which sets the scaling in Fig. 5. The frequency of many circuits is limited by the transistor f_T, which is low in LAE TFTs. However, LC oscillators offer the benefit that they can operate beyond f_T by resonating

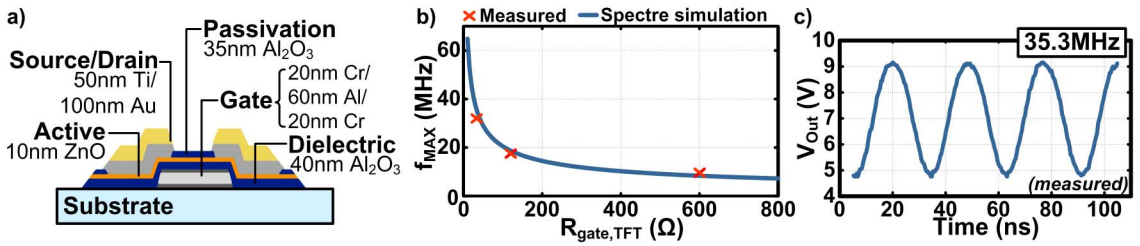

Fig. 9. ZnO TFT. (a) Cross section. (b) f_{MAX} versus $R_{\mathrm{gate,TFT}}$. (c) Measured oscillator output.

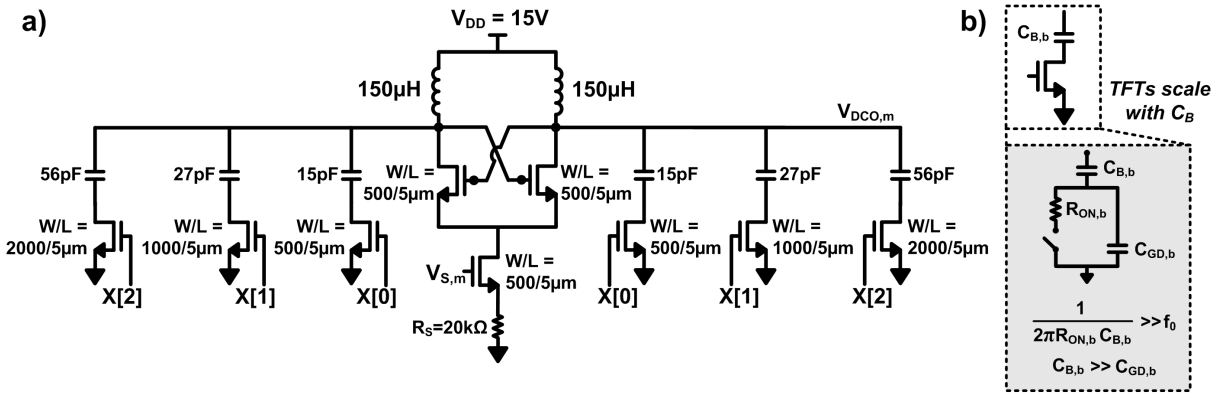

Fig. 10. DCO design based on (a) bank of switched, binary-weighted capacitors and (b) TFT switches designed for appropriate ON resistance $R_{\mathrm{ON},b}$ and drain capacitance $C_{\mathrm{GD},b}$.

the transistor capacitances with a tank inductor. In particular, to achieve oscillations, the following is a necessary condition:

$$\frac{g_m}{C_{\mathrm{Par}}} \times \frac{L_{\mathrm{ind}}}{R_{\mathrm{gate,TFT}} + R_{\mathrm{ind}}} > 1 \tag{1}$$

where g_m is the TFT transconductance, C_{Par} is the total parasitic capacitance (roughly $5 \times C_{\mathrm{GS}/D}$, where $C_{\mathrm{GS}/D}$ is the TFT overlap capacitance to which Miller multiplication is applied in the case of a cross-coupled oscillator topology), L_{ind} is the tank inductance, $R_{\mathrm{gate,TFT}}$ is the TFT gate resistance, and R_{ind} is the inductor resistance.

LC oscillators provide a key benefit—although g_m/C_{Par}, a quantity related to f_T, is limited in TFTs, LAE enables the formation of physically large, and therefore high-Q, inductors (i.e., inductors with many turns and thick traces), yielding sufficiently high $L_{\mathrm{ind}}/R_{\mathrm{ind}}$ to meet the oscillation condition [18]. With this as a starting point, the following sections describe: 1) TFT optimizations performed to maximize achievable oscillator frequencies and 2) circuit design for amplitude-modulated DCOs with many well-separated frequency channels.

1) TFT f_{MAX} Optimization: Following from the oscillation condition, the maximum frequency is limited by TFT f_{MAX}, rather than f_T. The expression for f_{MAX} is as follows:

$$f_{\mathrm{MAX}} = \frac{f_T}{\sqrt{2\pi f_T C_{\mathrm{GD}} R_{\mathrm{gate,TFT}} + \frac{R_{\mathrm{gate,TFT}}}{r_{o,\mathrm{TFT}}}}} \tag{2}$$

where the TFT output resistance $r_{o,\mathrm{TFT}}$ is much larger than $R_{\mathrm{gate,TFT}}$, making the last denominator term negligible. In addition to the dependence on g_m and $C_{\mathrm{GS}/D}$ (both explicitly and through f_T), we see a dependence on $R_{\mathrm{gate,TFT}}$, a parameter generally overlooked in TFT optimization. Our previous work has focused on $R_{\mathrm{gate,TFT}}$, through TFT gate-electrode materials and layer-thickness optimization [12]. Specifically, as shown in Fig. 9(a), our zinc-oxide (ZnO) TFTs employ a gate electrode composed of thin chrome (20 nm) for superior adhesion below an aluminum layer, whose thickness (60 nm) is set to minimize resistance but also ensure integrity for step coverage by the thin aluminum-oxide dielectric (40 nm). An additional Cr layer prevents formation of hillocks in the Al that otherwise form during the oxide deposition process, compromising the gate oxide. As shown in Fig. 9(b), the Al layer in the gate electrode reduces the gate resistance by 20 times for the TFTs employed (from 600 to 30 Ω), significantly enhancing their f_{MAX}. Fig. 9(c) shows a measurement of a corresponding oscillator's output, with frequency near the measured f_{MAX} [12]. Additional TFT optimizations (reducing channel length and lowering overlap capacitances through self-alignment) show further promise for pushing f_{MAX} [19], potentially benefitting the presented architecture.

2) Frequency-/Amplitude-Modulated TFT DCO: The high TFT oscillation frequency described above enables the creation of several frequency channels in the DCO. Fig. 10(a) shows the basic DCO topology. The DCO incorporates a bank of N binary-weighted capacitors, configured via TFT switches driven by a digital frequency-control signal $X[N-1:0]$. So long as the oscillation condition is still satisfied, this configuration alters the tank capacitance, thereby modulating the free-running oscillation frequency to one of 2^N values.

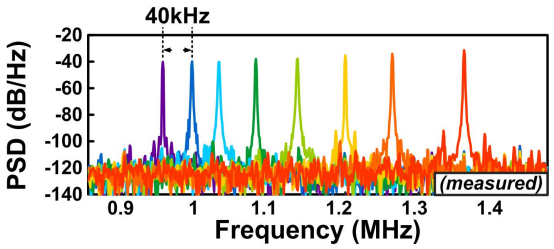

Fig. 11. Measured PSD for one oscillator showing minimum inter-channel spacing of 40 kHz.

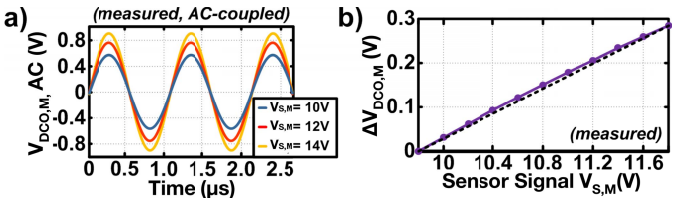

Fig. 12. Measured amplitude modulation of a DCO output $V_{\mathrm{DCO},M}$ over a range of input sensor signals $V_{S,M}$, and (a) time-domain waveform and (b) $V_{S,M}$-to-$V_{\mathrm{DCO},M}$ transfer function in purple.

Fig. 10(b) shows the considerations for designing the TFT switches. For the "ON" condition, the deep-triode TFT resistance $R_{\mathrm{ON},b}$ should be small, such that the high-pass cutoff frequency (set with the bank capacitor $C_{B,b}$, from the DCO output to the bank capacitor's bottom plate) is well above the highest oscillation frequency: $f_{\mathrm{HP}} = (1/2\pi R_{\mathrm{ON},b} C_{B,b}) >> f_{\mathrm{DCO}}, \mathrm{MAX}$. This requires the width of the TFT to be large. However, large width increases the parasitic gate-to-drain overlap capacitance $C_{\mathrm{GD},b}$, which must be kept significantly smaller than the corresponding bank capacitor for the "OFF" condition ($C_{\mathrm{GD},b} << C_{B,b}$). Thus, the DCO channel frequencies must be considered in this optimization. For the TFTs employed, $R_{\mathrm{ON},b} \approx 500$ kΩ/W and $C_{\mathrm{GD},b} \approx 0.01$ pF×W (where W is the TFT width in μm), yielding $(1/2\pi R_{\mathrm{ON},b} C_{\mathrm{GD},b}) \approx 30$ MHz. Based on these values, the eight nominal channel frequencies are designed to be 1.36, 1.27, 1.20, 1.13, 1.07, 1.03, 0.99, and 0.95 MHz, and the bank capacitor values and TFT sizes used are as shown in Fig. 10(a). This design point results in a value of 2–6 in the oscillation condition (1), ensuring robust oscillations.

Fig. 11 shows the measured power-spectral density (PSD) plot for one oscillator for the different values of $X[2:0]$. Because the capacitance-to-frequency relationship is nonlinear ($f = 1/2\pi(LC)^{1/2}$), the separation between frequency channels is expected to be non-uniform. At the design point, the minimal nominal channel spacing is seen to be 40 kHz.

For amplitude modulation of the DCO output $V_{\mathrm{DCO},m}$, the sensor signal $V_{S,m}$ is fed to the gate of a tail TFT, as shown in Fig. 10(a). This sets the current through the DCO, and thus the amplitude of the output oscillation. We note that $V_{S,m}$ reconstruction via the matrix inversion discussed in Section II requires a linear $V_{S,m}$-to-$V_{\mathrm{DCO},m}$ transfer function; source degeneration enhances linearity and is employed in the tail TFT. Fig. 12(a) shows the change in measured DCO output in one of the frequency channels over a 4-V input range (actual sensor range used is 1 V in the demonstrated system).

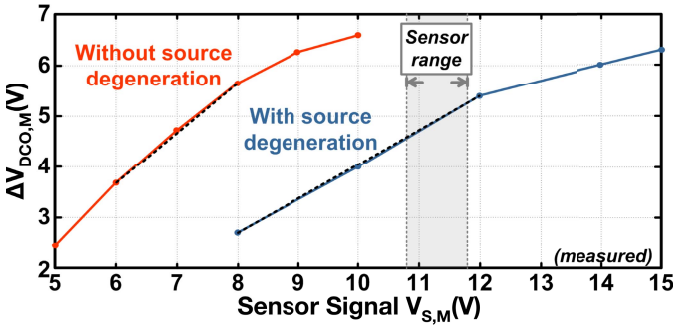

Fig. 13. $V_{S,M}$-to-$V_{\mathrm{DCO},M}$ transfer function with and without source degeneration of a tail transistor, showing increased linear range of sensor input. Linear fit within range for each in black.

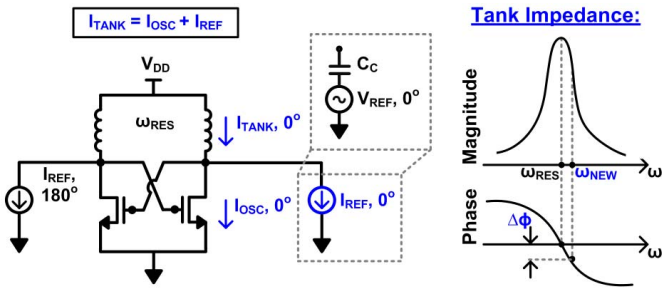

Fig. 14. Injection-locking principle.

Fig. 12(b) plots the source-degenerated transfer function, showing the linearity achieved. Fig. 13 shows the transfer function with and without source degeneration, and indicates that the maximum linear range for the sensor input doubles after source degeneration.

B. Injection-Locked DCO Array

Properly representing the summation of sensor values by the summation of DCO outputs requires all DCOs in the array to be phase/frequency synchronized. Phase variation of DCOs arises because oscillators turn on at different times, and frequency variation is a result of tank inductance and capacitance variations. In particular, the tank capacitance is set in part by TFT overlap capacitances (as described in Section III-A), and is subject to process variation.

To achieve phase/frequency synchronization, injection locking is employed [20]. As shown in Fig. 14, this is done by introducing currents $I_{\mathrm{REF},0°}/I_{\mathrm{REF},180°}$ into the LC tank at a desired reference frequency, which add with the TFT currents $I_{\mathrm{OSC},0°}/I_{\mathrm{OSC},180°}$, to give the tank currents $I_{\mathrm{TANK},0°}/I_{\mathrm{TANK},180°}$. For reference frequencies near the free-running resonant frequency of the oscillator, phase/frequency locking to the reference signal is achieved. We note that in being pulled off its resonant point, the tank's voltage exhibits a phase offset with its current, set by its impedance, as shown schematically. Furthermore, because $I_{\mathrm{OSC},0°}/I_{\mathrm{OSC},180°}$ change with $V_{S,m}$ (for DCO amplitude modulation), $I_{\mathrm{TANK},0°}/I_{\mathrm{TANK},180°}$ change in both their amplitude and phase, with respect to the reference signal. Nonetheless, as seen in Fig. 15 for a typical DCO, measurements of $V_{S,m}$-to-$V_{\mathrm{DCO},m}$ transfer functions (left) and linearity

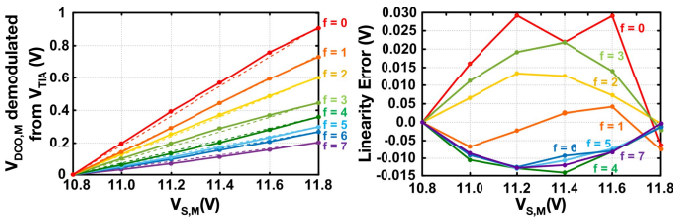

Fig. 15. Measured relationship between sensor input voltage and output amplitude of locked DCO in each frequency channel with linear fits in dashed lines (left) and linearity error (right).

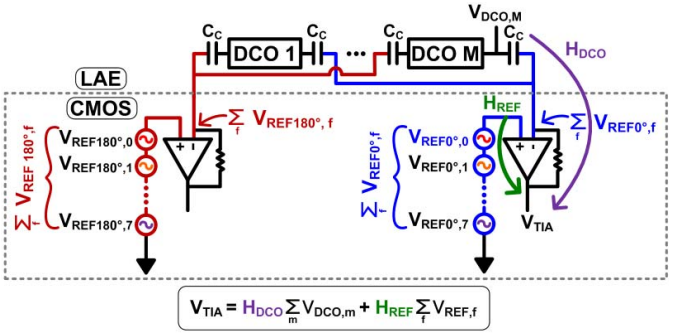

Fig. 16. Injection-locking implementation, with sum of reference signals applied at TIAs, and coupled to oscillators from virtual ground nodes.

error (right) in each channel for locked DCOs show that good linearity is maintained over the sensor range, which is designed to be 10.8–11.8 V.

In the system, the reference frequencies are generated in the CMOS domain as voltages and distributed to the DCOs via the coupling capacitors C_C. Locking of all DCOs is achieved simultaneously in all frequency channels. This is possible since the impact of $I_{\mathrm{REF},0°}/I_{\mathrm{REF},180°}$ is effectively filtered by the tank-impedance magnitude response (as shown in Fig. 14). Thus, reference signals sufficiently far from the free-running frequency will have negligible effect. The precise effect can be controlled by setting the magnitude of $I_{\mathrm{REF},0°}/I_{\mathrm{REF},180°}$ [20]. In this way, the lock range can be selectively limited. Accordingly, Fig. 16 shows the approach for achieving simultaneous locking in all channels within the $N = 3$ system demonstrated. The sum of $2^N = 8$ reference sinusoids (at the desired channel frequencies) is applied to the reference node of the CMOS TIA. The TIA's feedback condition causes the sum to then appear at the virtual ground node, which couples back to all DCOs through the C_Cs. This is done differentially through two TIAs, as shown. The amplitude of each reference sinusoid is chosen to be large enough to ensure locking over the DCO frequency spread expected within each channel, but small enough to ensure that adjacent channels experience negligible frequency pulling. Fig. 17 shows how tuning the $V_{\mathrm{REF},f}$ amplitude changes the lock range in each frequency channel ($V_{S,m}$ is set to a value in the middle of the sensor range).

Fig. 18 shows the measured PSD for 11 DCOs, each configured to output in each of the frequency channels. The free-running DCOs (top) show substantial frequency spread, which corresponds with the measured variation of

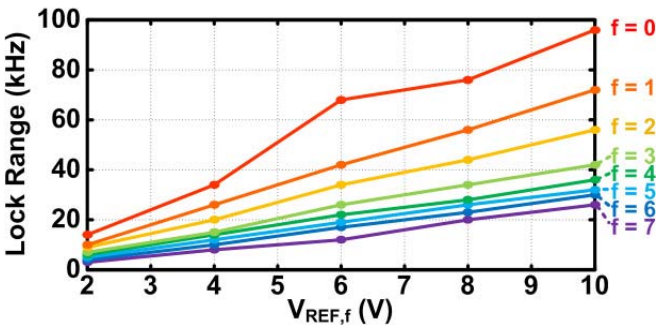

Fig. 17. Simulated lock range versus reference-signal voltage amplitude.

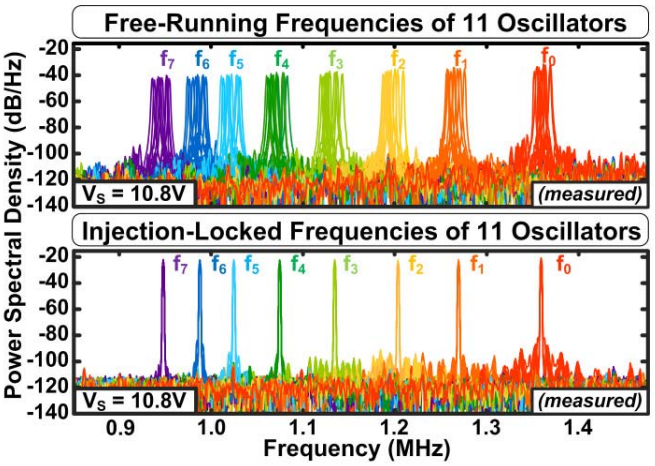

Fig. 18. Measured PSDs for 11 free-running (top) and injection-locked (bottom) DCOs.

the TFT capacitances. The measured coefficient of variation (σ/μ) is 0.035 over the oscillator TFTs. The injection-locked DCOs (bottom) show that the frequency spread is eliminated. Furthermore, injection locking is seen to enhance the spectral purity in all of the frequency channels, raising the possibility of introducing even more intermediate channels to enhance the scaling in number of sensors in Fig. 5.

In addition to process-induced TFT-capacitance variation, implementation of the system on flexible foils would introduce inductance variation upon bending. Based on our prior measurements of such variations [12], the system is estimated to tolerate bending radii of ~25 mm (corresponding to <1% frequency shift from nominal values).

C. Sensor-Signal Demodulation

Recall that reconstruction of the sensor signals $V_{S,m}$, represented by the vector $\vec{s}$, is enabled by the matrix relationship $\vec{y} = \mathbf{T} \times \vec{s}$, where $\mathbf{T}$ is the "hopping matrix" (described in Section II), and $\vec{y}$ is a vector whose elements represent the sum of sensor signals in each frequency channel for each value of the hopping-control code. Ideally, the elements of $\vec{y}$ could be retrieved by demodulating the output of the CMOS TIA V_{TIA} in each of the frequency channels. In the system presented, demodulation is complicated due to injection locking. As shown in Fig. 16, V_{TIA} consists of two components:

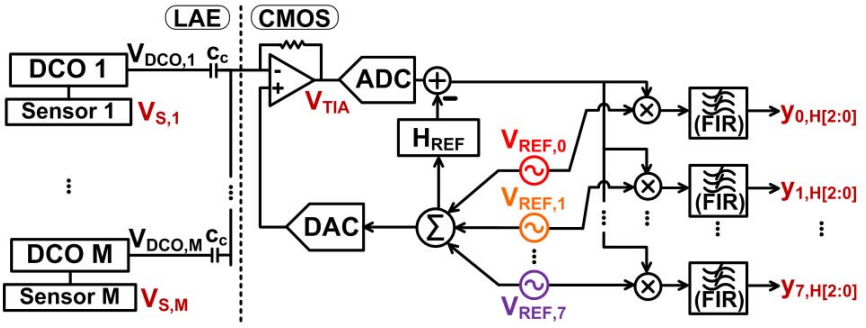

Fig. 19. Demodulation in CMOS domain involves the removal of reference-signal component, followed by mixing with corresponding reference signal and low-pass filtering.

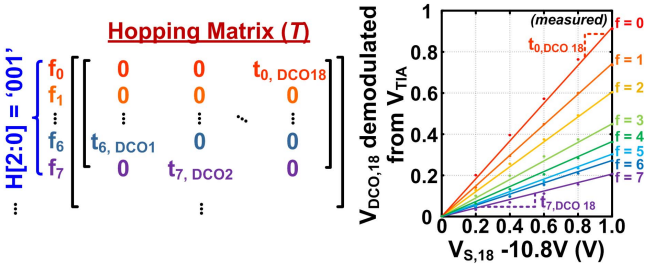

Fig. 20. **T** matrix entries derived from $V_{S,m}$-to-V_{TIA} transfer function in each frequency channel.

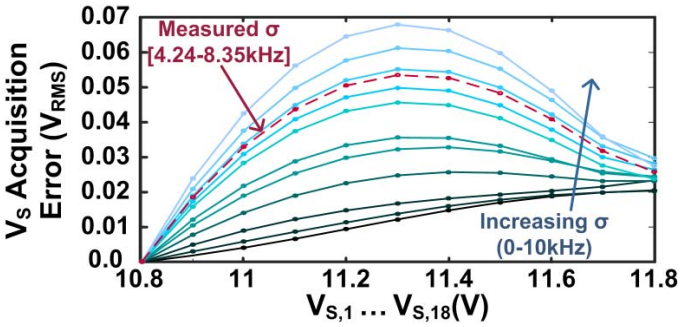

Fig. 21. Simulated impact of TFT capacitance variation on sensor-acquisition error.

1) the sum of the DCO outputs, through the transfer function H_{DCO} and 2) the sum of the injection-locking reference signals, through the transfer function H_{REF}. Thus, the component due to the reference signals must be removed before the elements of $\vec{y}$ can be demodulated.

The demodulation process is shown in Fig. 19. First, a sum of the injection-locking reference signals is generated in CMOS. This is applied to the TIA reference node via a DAC, but is also applied to the transfer function H_{REF} and then subtracted from the digitized V_{TIA} output. The resulting signal is then demodulated, using each of the injection-locking reference signals. Finally, low-pass finite impulse response (FIR) filtering is applied to each demodulated output to filter signals from other frequency channels, recovering only the sum of sensor signals in the corresponding frequency channel. The demodulated and filtered outputs from each frequency channel then form the elements of $\vec{y}$.

For $\vec{s}$ reconstruction from $\vec{y}$ via the inverse of **T**, a final requirement is to determine the nonzero elements of **T**. These are derived from the linear transfer function from $V_{S,m}$-to-$V_{\mathrm{DCO},m}$ and then from $V_{\mathrm{DCO},m}$-to-V_{TIA} (H_{DCO}). Nominally, the nonzero elements would all be equal. In the architecture, however, the transfer functions are slightly different for each frequency channel, since $V_{S,m}$-to-$V_{\mathrm{DCO},m}$ depends on the tank Q and $V_{\mathrm{DCO},m}$-to-V_{TIA} depends on the impedance of C_C. To determine the transfer functions, we employ calibration, as shown in Fig. 20. For each frequency channel, the sensor-voltage input for a DCO is swept, and the value of V_{TIA} is demodulated. A linear fit to the demodulated output is then performed to establish the corresponding entry of **T**.

Aside from dependence on frequency channel, the transfer function associated with each DCO is also impacted by process variations. First, variation in the mobility and threshold voltage of the DCO tail TFT directly affects the conversion of $V_{S,m}$ to $V_{\mathrm{DCO},m}$, as is typical in a modulator; this variation is found to be modest (measurements provided in Section IV). Second, injection locking introduces a new dependence on TFT capacitance variation, which is observed to cause sensor-acquisition error in the demonstrated system. Specifically, as seen in Fig. 14, the tank impedance is a function of the offset between the reference frequency and the DCO free-running resonant frequency, set in part by the capacitances of the cross-coupled TFTs. We analyze the impact of this on the presented system. To simulate the variation, we assume Gaussian distributions for DCO free-running frequencies in each of the frequency channels, and randomly sample from these for 18 DCOs. We then determine the sensor-voltage acquisition error for a system where all 18 sensor voltages $V_{S,1} \ldots V_{S,18}$ are swept together over the input range. This process is repeated for distributions with different standard deviations. Fig. 21 shows the result, also plotting a case corresponding to the measured free-running-frequency standard deviations in each channel of the prototype. We note that the measured error curve falls in the expected range predicted by simulation, indicating that the DCO free-running frequency variation, not tail-TFT transconductance variation, is the dominant source of system acquisition error. Finally, we observe that acquisition error can be notably improved beyond that of our prototype by controlling variation, as done in industrial fabrication environments.

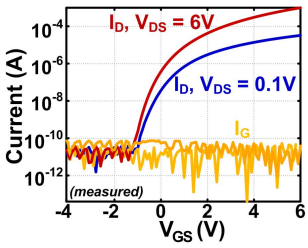

Mobility μ (cm^2/Vs)	13
Threshold Voltage (V)	2
$C_{GS/D}$ (fF/μm)	5
μ Coefficient of Variation*	0.13
V_T Coefficient of Variation*	0.069
$C_{GS/D}$ Coefficient of Variation*	0.035
Max Process T (°C)	180
Channel Width W (μm)	500
Channel Length L (μm)	5

*Over 15 TFTs

Fig. 22. Typical I–V characteristics and device parameters of LC tank ZnO TFTs in this paper.

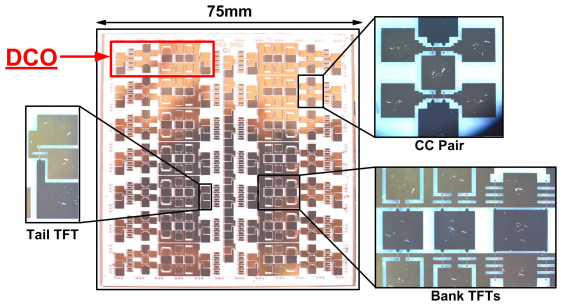

Fig. 23. Photograph of DCO sample, highlighting DCO components.

IV. System Prototype and Demonstration

The presented system was prototyped and demonstrated in a large-area, force-sensing application for characterization. The DCO array was fabricated in-house using ZnO TFTs, employing plasma-enhanced atomic-layer deposition (PEALD) at temperatures <200 °C (i.e., plastic-compatible temperatures). PEALD provides uniform, reproducible thin films at these temperatures, enabling an active ZnO layer and dielectric layer just 10 and 40 nm thick, respectively [21]. Devices are *in situ* back-channel passivated for improved reliability. For illustration, typical characteristics (the same size as cross-coupled TFTs in DCOs) are provided in Fig. 22, showing the drain and gate-leakage currents. The extracted mobility and threshold voltage are 13 cm^2/Vs and 2 V, and the measured coefficients of variation are 0.13 and 0.069 (from 15 devices), respectively. Furthermore, the measured TFT gate-to-source/drain capacitance is 5 pF, with a coefficient of variation 0.035. While mobility and threshold voltage variations are comparatively small, capacitance variations cause errors as modeled in Section III-C and measured in Fig. 26. A fabricated DCO-array sample is shown in Fig. 23 (the system demonstration employs DCOs from multiple such samples). Though the circuits are verified on flexible polyimide foil, a glass substrate is employed for this system to simplify testing with a standard probe card.

The demonstrated system is shown in Fig. 24. The system uses a 3-bit hopping code (eight frequency channels). Although this allows for 19 sensors, 18 sensors are employed to give a regular arrangement. The sensor array consists of commercial, polymer-based piezoresistive force sensors [22], distributed across a 30-cm rigid PCB. Each sensor forms one leg of a voltage divider, with the other leg being a fixed resistor; the intermediate voltage provides the sensor voltage $V_{S,m}$, which is designed to be in the range of

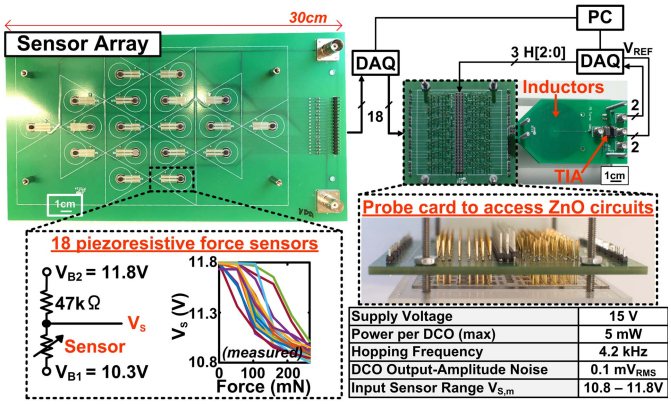

Supply Voltage	15 V
Power per DCO (max)	5 mW
Hopping Frequency	4.2 kHz
DCO Output-Amplitude Noise	0.1 mV_{RMS}
Input Sensor Range $V_{S,m}$	10.8 – 11.8V

Fig. 24. Demonstrated system, showing sensor array and probe card for interfacing with a TFT sample, as well as a DAQ system for post processing and configuration.

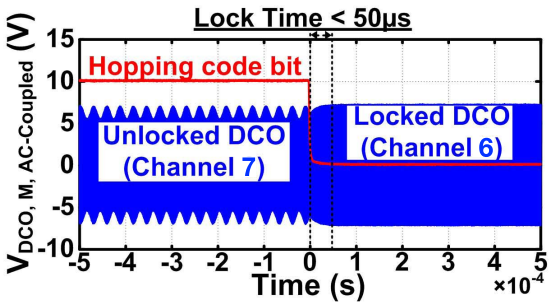

Fig. 25. Time required for DCO to lock is <50 μs.

10.8–11.8 V. This is then fed to the tail transistor of an associated DCO. The inset in Fig. 24 shows $V_{S,m}$ for all sensors as weight is added. Significant variation in the sensors is observed.

A data-acquisition (DAQ) system reads the sensor data, which is then provided to the DCO array for measurement. The DCO outputs are captured with a second DAQ through a single differential interface. The probe card PCB used to interface the TFT sample on glass to the system is shown at the right. The probe card contains an array of gold pins that contact the sample, visible in the side view, and also contains the bank capacitors, implemented as surface-mount components for ease of testing. With one additional lithographic mask, the capacitors could be integrated onto the TFT sample; prior measurements of metal–metal capacitors with a 40-nm-thick aluminum-oxide dielectric give a capacitance density of 2 $fF/\mu m^2$. The DCO inductors are also fabricated on a PCB, and connect to the probe card via rigid header connections. We note that in an actual system realization (Fig. 2), the DCOs would be distributed alongside the sensors, rather than separated, as done here for testability and characterization. Hence, the DCO area, dominated by the large-area inductor (50 cm^2), would likely set the sensor density.

The hopping frequency used was 4.2 kHz, limited by the length of FIR filtering required for demodulation. As this rate compares favorably with active matrix scan rates, but the number of sensors that can be accessed is higher, cycling through the hopping codes at this rate enables a high rate

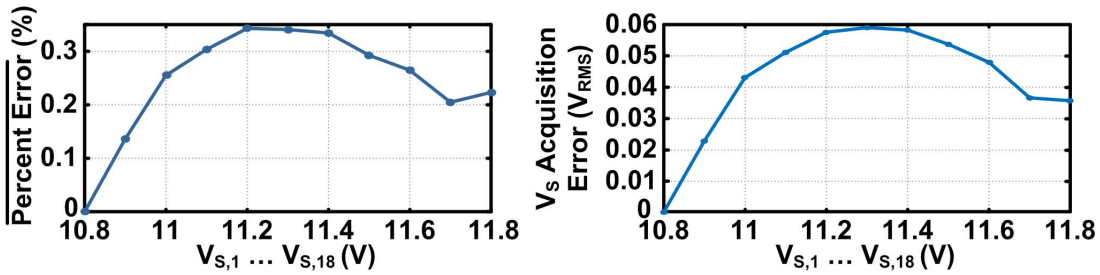

Fig. 26. Average error as a percentage of ground-referenced signal $V_{S,m}$ (left) and as absolute voltage (right) of a DCO-based acquisition system as a function of sensor-voltage level.

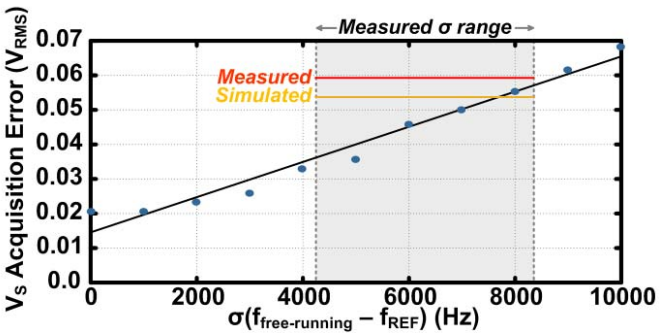

Fig. 27. Comparison of simulated maximum error (versus standard deviation of free-running DCO frequencies) and measured maximum percentage error.

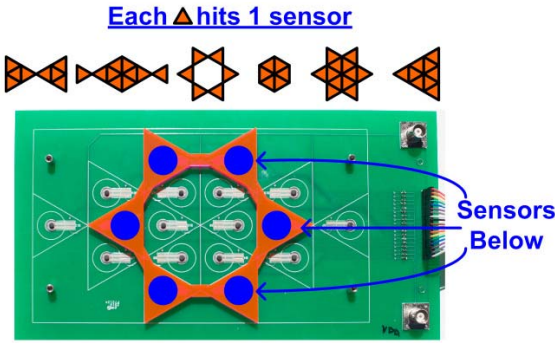

Fig. 28. Force-sensor data obtained by placing shapes of calibrated weight on a sensing plane.

for accessing sensors. The time required for injection locking is less than this FIR window, and does not ultimately limit scan speed. Fig. 25 shows a DCO output (in blue) changing from frequency channel 7 to 6. The reference signal is applied only for channel 6, with the reference signal for channel 7 explicitly turned off so that beating in channel 7 due to slight frequency pulling makes the locking transition time visually apparent. The left-hand side of the plot shows that before the least significant bit (LSB) of the DCO frequency-control signal (red) is switched, the DCO does not lock to the reference for channel 6. Once the LSB switches, the DCO locks in <50 μs.

The DCOs operate from a 15-V supply, with each consuming a maximum of 5 mW (depending on $V_{S,m}$). This raises an additional consideration for an eventual system (Fig. 2). While previous work has shown TFT oscillators operating from thin-film photovoltaic (PV) harvesters [18], such PVs would occupy an additional area of 1–100 cm^2 per DCO (depending on lighting conditions), further reducing the sensor density. However, TFT improvements (such as the work in [19] pushing f_{MAX} of PEALD ZnO TFTs from 40 MHz to >1 GHz) offer substantial promise to reduce power and PV area requirements.

Before reconstructing specific patterns by measuring the weight applied to the sensor array, we first characterize the acquisition error of the system itself over the entire sensor voltage range. To do this, all sensor voltages $V_{S,1} \ldots V_{S,18}$ were swept for all DCOs in the array simultaneously. Fig. 26 shows the maximum acquisition error measured for each sensor voltage level as percent error of the ground-referenced input signal $V_{S,m}$ (left), and as absolute error of the input signal $V_{S,m}$ (right). The maximum percentage error is 0.34% over the full sensor range. This is slightly different than the percent error measured in [17] (0.27%), as a different set of fabricated DCOs have been used, from which the standard deviation of the DCO free-running frequencies in each of the channels was also measured to allow simulation of the error due to variations. This percent error corresponds to a voltage error of at most 60 $\mathrm{mV_{rms}}$. Comparing the measured error with the simulation results in Fig. 21 derived from the measured standard deviations indicates good agreement. Specifically, Fig. 27 plots the simulated maximum error versus the standard deviation of the free-running frequencies (in each channel). Also shown is the range of measured standard deviations (across all frequency channels), as well as both the simulated maximum error based on these standard deviations and the actual measured maximum error.

The sample-to-sample noise of the DCO output sinusoids is measured to be $\leq$0.1 $\mathrm{mV_{rms}}$ (Fig. 24), which is small compared to the $>$200-mV DCO amplitude change over a 1-V sensor range. Thus, particularly after FIR filtering in the demodulation process, noise is not expected to be a prominent source of error compared to other sources (e.g., non-linearity, capacitance variations).

To generate force-sensor data, weights are cut from acrylic sheets into six different shapes, as shown in Fig. 28. Each weight pattern contacts a different set of sensors, and the total weight of the pattern is increased in discrete steps by stacking multiple identical weights on the array. Each shape applies a weight of 52 mN to each sensor it contacts. Generating calibrated weights in this way enables detailed characterization of acquisition data and error over the input range.

For each of the six weight patterns at each of the five weight levels, Fig. 29 shows the measured results both for direct

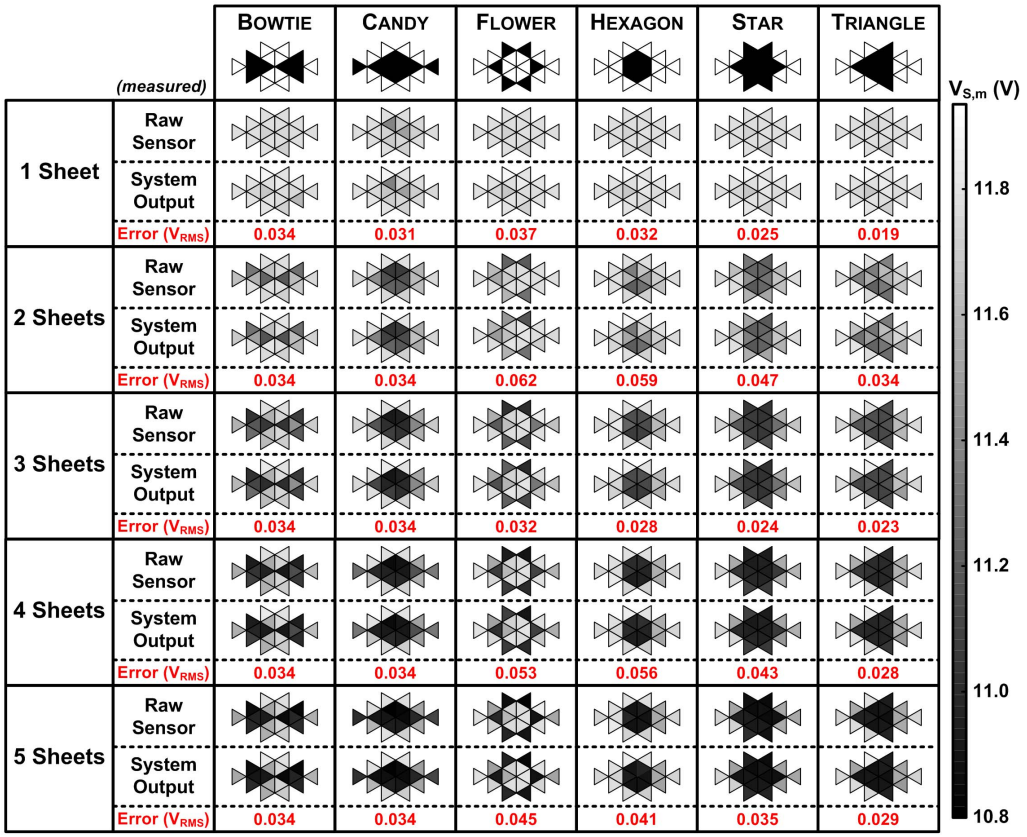

Fig. 29. Results comparing sensor voltages acquired through the demonstrated system and those acquired directly from a DAQ system.

acquisition of the 18 sensor voltages through 18 interfaces, and for reconstruction of the sensor voltages via the system through a single differential interface. As seen, the sensor voltages themselves (directly acquired) exhibit variation, due to variation in the force sensors. This causes some deviation of the acquired patterns from the actual shapes (particularly at low weight levels). However, across all shapes and all weight levels, the error for the sensor voltages acquired via the system is $\leq$62 mV_{rms}.

V. Conclusion

In hybrid systems that combine LAE, for large-scale, distributed sensing, with silicon CMOS, for sensor-data processing, the interfaces required between the technologies present a dominant limitation to system scalability. This paper demonstrates a system based on amplitude-modulated, frequency-hopping TFT DCOs that reduces this number of interconnections, while maintaining a high rate of sensor accessing, much beyond what can be achieved with existing active-matrix and binary-addressing schemes. The primary tradeoff is increased dynamic-range required for signal processing. However, this affects only the CMOS domain, where circuits can have much higher energy efficiency and performance thanks to the availability of a high-performance transistor technology. For demonstration and characterization, an 18-element force-sensing system is prototyped, employing ZnO TFT DCOs. The system performs frequency hopping over the 18-DCO array at a rate of 4.2 kHz and achieves a maximum acquisition error $\leq$0.31% over 30 weight patterns.

References

[1] W. Gao *et al.*, "Fully integrated wearable sensor arrays for multiplexed *in situ* perspiration analysis," *Nature*, vol. 529, no. 7587, pp. 509–514, Jan. 2016.

[2] S. Cheng and Z. Wu, "Microfluidic electronics," *Lab Chip*, vol. 12, no. 16, pp. 2782–2791, Jun. 2012.

[3] T. Yokota *et al.*, "Ultraflexible, large-area, physiological temperature sensors for multipoint measurements," *Proc. Nat. Acad. Sci. USA*, vol. 112, no. 47, pp. 14533–14538, Sep. 2015.

[4] C. M. Boutry, A. Nguyen, Q. O. Lawal, A. Chortos, S. Rondeau-Gagné, and Z. Bao, "A sensitive and biodegradable pressure sensor array for cardiovascular monitoring," *Adv. Mater.*, vol. 27, no. 43, pp. 6954–6961, Sep. 2015.

[5] K. Ishida *et al.*, "Insole pedometer with piezoelectric energy harvester and 2 V organic circuits," *IEEE J. Solid-State Circuits*, vol. 48, no. 1, pp. 255–264, Jan. 2013.

[6] K. Myny, E. van Veenendaal, G. H. Gelinck, J. Genoe, W. Dehaene, and P. Heremans, "An 8-bit, 40-instructions-per-second organic microprocessor on plastic foil," *IEEE J. Solid-State Circuits*, vol. 47, no. 1, pp. 284–291, Jan. 2012.

[7] K. Myny *et al.*, "A thin-film microprocessor with inkjet print-programmable memory," *Sci. Rep.*, vol. 4, Dec. 2014, Art. no. 7398.

[8] E. Cantatore *et al.*, "A 13.56-MHz RFID system based on organic transponders," *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 84–92, Jan. 2007.

[9] P. Harpe, E. Cantatore, and A. van Roermund, "A 2.2/2.7fJ/conversion-step 10/12b 40kS/s SAR ADC with Data-Driven Noise Reduction," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 270–271.

[10] V. Fiore *et al.*, "A 13.56MHz RFID tag with active envelope detection in an organic complementary TFT technology," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2014, pp. 492–493.

[11] W. Rieutort-Louis *et al.*, "Current gain of amorphous silicon thin-film transistors above the cutoff frequency," in *Device Res. Conf. Dig. Tech. Papers*, Jun. 2014, pp. 273–274.

[12] Y. Afsar *et al.*, "Impact of bending on flexible metal oxide TFTs and oscillator circuits," *J. Soc. Inf. Display*, vol. 24, no. 6, pp. 371–380, Jun. 2016.
[13] T. Moy *et al.*, "An EEG acquisition and biomarker-extraction system using low-noise-amplifier and compressive-sensing circuits based on flexible, thin-film electronics," *IEEE J. Solid-State Circuits*, vol. 52, no. 1, pp. 309–321, Jan. 2017.
[14] Y. Hu *et al.*, "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.
[15] N. Verma, Y. Hu, L. Huang, W. S. Rieutort-Louis, J. S. 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.
[16] 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," *Proc. SPIE*, vol. 3336, p. 444, Feb. 1998.
[17] Y. Afsar, T. Moy, N. Brady, S. Wagner, J. C. Sturm, and N. Verma, "Large-scale acquisition of large-area sensors using an array of frequency-hopping ZnO thin-film-transistor oscillators," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 256–257.
[18] Y. Hu *et al.*, "Flexible solar-energy harvesting system on plastic with thin-film LC oscillators operating above f_t for inductively-coupled power delivery," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2012, pp. 1–4.
[19] Y. Mehlman, Y. Afsar, N. Verma, S. Wagner, and J. Sturm, "Self-aligned ZnO thin-film transistors with 860 MHz f_T and 2 GHz f_{MAX} for large-area applications," in *Device Res. Conf. Dig. Tech. Papers*, Jun. 2017, pp. 1–2.
[20] B. Razavi, "A study of injection locking and pulling in oscillators," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1415–1424, Sep. 2004.
[21] Y. V. Li, K. G. Sun, J. I. Ramirez, and T. N. Jackson, "Trilayer ZnO thin-film transistors with *in situ* Al_2O_3 passivation," *IEEE Electron Device Lett.*, vol. 34, no. 11, pp. 1400–1402, Nov. 2013.
[22] Interlinkelectronics. *Force-Sensing Resistor*. [Online]. Available: http://www.interlinkelectronics.com/

Yasmin Afsar (S'17) received the B.A. degree (Hons.) in physics and the M.S. degree in electrical engineering from Columbia University, New York, NY, USA, in 2010 and 2011, respectively, and the M.A. degree in electrical engineering from Princeton University, Princeton, NJ, USA, in 2013, where she is currently pursuing the Ph.D. degree.

Her current research interests include expanding large-area system applications through metal oxide devices and circuits.

Ms. Afsar was a recipient of the National Science Foundation Graduate Research Fellowship.

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, where she is currently pursuing the Ph.D. degree in electrical engineering.

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

Ms. Moy was a recipient of the Princeton University School of Engineering and Applied Science Award for Excellence.

Nicholas Brady is currently pursuing the B.S.E. degree in electrical engineering with Princeton University, Princeton, NJ, USA.

He was with InPhase Research, Fairfax, VA, as a Software Integrator for software defined radio applications.

Mr. Brady was a recipient of the Mid-InfraRed Technologies for Health and Environment Research Experience for Undergraduates Fellowship and conducted research under Dr. N. Verma.

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

Following a Post-Doctoral Fellowship at The Ohio State University, Columbus, OH, he worked at Bell Telephone Laboratories, Murray Hill, NJ, USA, and Bell Telephone Laboratories, Holmdel, NJ, USA, on semiconductor memories and heterojunction solar cells from 1970 to 1978. He then joined the Solar Energy Research Institute (now NREL), Golden, CO, USA, as the founding Chief of the Photovoltaic Research Branch. Since 1980, he has been a Professor of electrical engineering with Princeton University, Princeton, NJ, USA; in 2015, he became a Professor Emeritus and a 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 from Stanford University, Stanford, CA, USA, in 1981 and 1985, respectively.

In 1979, he joined Intel Corporation, Santa Clara, CA, USA, as a Microprocessor Design Engineer, and in 1981, he was a Visiting Engineer with 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, Princeton, and its successor, the Princeton Institute for the Science and Technology of Materials, Princeton. From 1994 to 1995, he was a von Humboldt Fellow at the Institut fuer Halbleitertechnik, University of Stuttgart, Stuttgart, Germany. He has involved in the fields of silicon-based heterojunctions, thin-film and flexible electronics, photovoltaics, the nano–bio interface, 3-D integration, and silicon-on-insulator.

Dr. Sturm received 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 the International Electron Devices Meeting from 1988 to 1992 and from 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, Woburn, MA, and SpaceTouch, Princeton.

Naveen Verma (S'04–M'09) received the B.A.Sc. degree in electrical and computer engineering from The University of British Columbia, Vancouver, BC, Canada, in 2003, and the M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2005 and 2009, respectively.

Since 2009, he has been with the Department of Electrical Engineering, 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.

Dr. Verma serves on the technical program committees for the International Solid-State Circuits Conference, the VLSI Symposium on Circuits, and the IEEE Signal-Processing Society (DISPS). He was a recipient or co-recipient of the 2006 DAC/ISSCC Student Design Contest Award, the 2008 ISSCC Jack Kilby Paper Award, the 2012 Alfred Rheinstein Junior Faculty Award, the 2013 NSF CAREER Award, the 2013 Intel Early Career Award, the 2013 Walter C. Johnson Prize for Teaching Excellence, the 2013 VLSI Symposium Best Student Paper Award, the 2014 AFOSR Young Investigator Award, the 2015 Princeton Engineering Council Excellence in Teaching Award, and the 2015 IEEE *Transactions on Components, Packaging and Manufacturing Technology* Best Paper Award.