A 100-fps, Time-Correlated Single-Photon-Counting-Based Fluorescence-Lifetime Imager in 130-nm CMOS

Ryan M. Field, *Member, IEEE*, Simeon Realov, *Member, IEEE*, and Kenneth L. Shepard, *Fellow, IEEE*

***Abstract*—A fully-integrated single-photon avalanche diode (SPAD) and time-to-digital converter (TDC) array for high-speed fluorescence lifetime imaging microscopy (FLIM) in standard 130-nm CMOS is presented. This imager is comprised of an array of 64-by-64 SPADs each with an independent TDC for performing time-correlated single-photon counting (TCSPC) at each pixel. The TDCs use a delay-locked-loop-based architecture and achieve a 62.5-ps resolution with up to a 64-ns range. A data-compression datapath is designed to transfer TDC data to off-chip buffers, which can support a data rate of up to 42 Gbps. These features, combined with a system implementation that leverages a x4 PCIe-cabled interface, allow for demonstrated FLIM imaging rates at up to 100 frames per second.**

***Index Terms*—Fluorescence lifetime imaging microscopy (FLIM), imaging, single-photon avalanche diodes (SPADs), time-correlated single-photon counting (TCSPC), time-to-digital converter (TDC).**

I. Introduction

FLUORESCENCE microscopy is a powerful imaging technique used in the biological sciences to identify labeled components of a sample with specificity. This is usually accomplished by labeling with fluorescent dyes and imaging these labels, isolating individual dyes by their spectral signatures with optical filters and determining signal from the intensity of the fluorescent response. Additional techniques, such as fluorescence energy resonance transfer (FRET), allow interactions between dyes to be monitored through measuring intensity ratios of the dyes' spectra[1]. Although these techniques are widely used, fluorescence intensity images can be negatively affected by intrinsic fluorescence of unlabelled molecules (autofluorescence), residual leakage of excitation illumination through the filters (bleedthrough), loss of fluorescence with continued illumination (photobleaching), and variations in fluorophore concentration [2].

Fluorophores have associated with them a characteristic lifetime, which defines the exponential fluorescent decay transient after the removal of the excitation source. These lifetimes, on the order of nanoseconds for organic dyes, are characteristic of the dye and its environment such as pH, local charge density, viscosity, and FRET interactions [3]–[5]. Consequently, the fluorophore lifetime can not only provide contrast in forming an image but can also serve as a sensing mechanism for the microenvironment of the fluorophore. FLIM has the property of being insensitive to fluorophore concentration and other factors that affect fluorescence intensity and has been applied to applications as diverse as bacteria detection, *in vivo* metabolic state identification, and FRET studies [6]–[9].

The two most common techniques for measuring the fluorescence lifetime are the modulated frequency-domain technique and time-correlated single-photon counting (TCSPC) [10]. Typically, wide-field frequency-domain techniques can record a few frames per second but are limited by the inability to detect small changes in lifetime or to resolve multi-exponential decays when more than one fluorophore is present at the same location. TCSPC allows for high accuracy in measuring lifetime and for the extraction of complex lifetime waveforms, as is necessary in chemical characterization studies [4]. However, TCSPC traditionally can require tens of seconds to acquire a single FLIM image in typical laser scanning systems.

In commercial TCSPC systems [11], one detector, typically an avalanche photodiode (APD) or photomultiplier tube (PMT), with one time-to-digital converter (TDC) measurement channel is raster scanned across a sample. At each point in the image, a laser is pulsed and the arrival time of the first fluorescent photon relative to the laser pulse is measured. With repeated laser pulses, a histogram of these individual photon arrival times is collected and the lifetime is extracted from the exponential fit to the resulting distribution. In order for the histogram distribution to match the true fluorescence lifetime, the fluorescence intensity should be sufficiently low such that a photon is only detected from around 1% of laser pulses [12]. For a typical 20 MHz laser repetition rate, a fluorescence intensity tuned for a 1% detection rate, an ideal scanning and detection system, and a minimum of 500 photon detections for lifetime extraction, it will usually take 250 μs to measure the lifetime at each pixel. Acquiring a 64-by-64 pixel image, therefore, requires at least 1 s.

Recent work has leveraged integrated arrays of CMOS single-photon avalanche diodes (SPADs) and TDCs to create parallelized TCSPC imaging systems [13]–[18]. Although improved imaging speeds have been demonstrated in some of these designs, the parallel acquisition channels generate

Manuscript received August 21, 2013; revised November 13, 2013; accepted November 16, 2013. This paper was approved by Guest Editor Jeffrey Gealow. This work was supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under contract number W911NF-12-1-0594 and by the National Science Foundation under grant 1063315.

R. M. Field is with Intel Corporation, Santa Clara, CA 95054 USA.

S. Realov is with Intel Corporation, Hillsboro, OR 97124 USA.

K. L. Shepard is with Columbia University, 1300 New York, NY 10027 USA.

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.2013.2293777

TABLE I
A TABLE SHOWING THE MAXIMUM THEORETICAL FRAME RATE FOR PREVIOUS SPAD ARRAYS WITH INTEGRATED TDCS. THIS ASSUMES AN EVENT-DRIVEN READOUT SCHEME, WHICH REQUIRES THAT EACH TDC DATA MUST ALSO BE TAGGED WITH THE PIXEL LOCATION FROM WHICH IT WAS GENERATED. THIS THEORETICAL MAXIMUM ASSUMES THAT 1000 PHOTON EVENTS ARE NEEDED TO EXTRACT THE LIFETIME AND THAT PHOTON EVENT DATA IS OUTPUT ON EVERY POSSIBLE I/O CLOCK CYCLE

Pixels	Bandwidth (Mbps)	Time Bits	Position Bits	Frame Rate (fps)	Ref.
1024	10240	10	10	500	[30]
1024	5120	10	10	250	[31]
4096	0.073	37	12	0.00036	[14]
16384	7680	10	14	19.5	[32]
20480	51200	10	15	100	[18]
4096	42000	10	10	466	This Work

off-chip data rates that limit the achievable frame rates. For a typical laser repetition rate of 20 MHz and a 64-by-64 array of pixels with 10-bit timing resolution and 12-bit position information, the required data rate reaches 1.8 Tbps. While event-driven readout approaches have reduced these data rates, previous SPAD array systems have still been limited in the number of parallel channels, frame rates, or number of acquired frames [19]. Table I lists published SPAD arrays with integrated TDCs and the theoretical data-bandwidth-limited maximum frame rate.

In this work, we present an FLIM imager containing a 64-by-64 array of SPADs in CMOS with per-pixel TDCs. An event-driven high-speed datapath supports a maximum imaging frame rate of 466 fps. The imager is designed using a standard 130-nm CMOS process, with an associated board-level data-handling system optimized for high-throughput operation. Section II describes an overview of the imager chip architecture, the detailed design of each on-chip component, and an overview of the system-level considerations for high-speed image acquisition. Section III presents measurement results that highlight the SPAD capabilities, characterize the TDCs, and demonstrate the high-speed FLIM performance of our system. Section IV concludes.

II. TCSPC FLIM IMAGING SYSTEM

A block diagram showing the entire imaging system is shown in Fig. 1. At the core of the system is the FLIM imager chip. The data output from the imager chip consists of raw arrival time data, which is arranged into histograms for each pixel by the four field-programmable gate arrays (FPGAs). Each FPGA bins the arrival time data from 1024 pixels, which is then transmitted to a computer where it is saved to disk before subsequent data processing to extract the lifetime. We now consider the design aspects of each of the major system components.

A. Integrated Circuit Architecture

In Fig. 2, a block diagram of the imager architecture is shown. The entire imager chip is synchronized to a 20-MHz laser signal using a phase-locked loop (PLL), which generates a 1-GHz clock signal that is distributed to the delay-locked loops (DLLs), as $\mathrm{clk_{dll}}$, and the datapath, as $\mathrm{clk_{datapath}}$. A trigger input signal allows the imager to synchronize the datapath controller and TDC start signals with any frequency that is an integer fraction of the 20-MHz laser repetition, allowing for laser pulse picking.

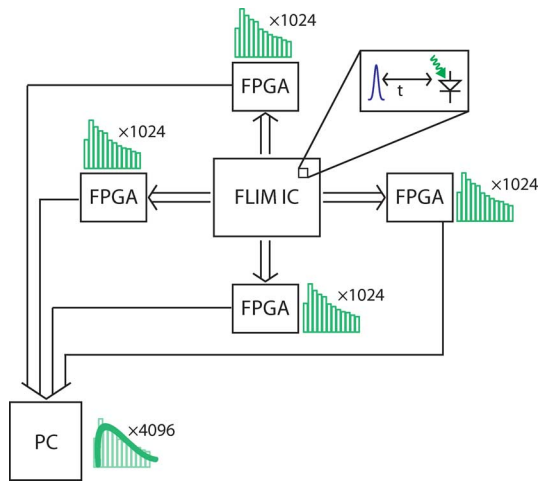

Fig. 1. System level block diagram showing a high-level overview of the connections between the IC, FPGAs, and PC.

Each pixel contains quench, reset, control, calibration, and output circuits. The output buffer for each pixel drives the stop signal for one of 4096 independent TDCs. The timing data recorded by the TDCs is shifted into a datapath for compression before passing to the chip periphery. Four banks of 22 LVDS buffers output a clock, a one-bit valid flag, and 20-bits, which consists of 10-bits of arrival time and 10-bits of position data, at up to 500 MHz.

B. SPAD Array

The SPADs used in this design are the same as those reported by the authors in [20]. Each pixel contains one octagonal SPAD with a 5-μm-diagonal active area and the pixel-level circuitry for control of the SPAD (see Fig. 3). With this circuitry, the SPADs are spaced on a 48-μm pitch, resulting in a fill-factor of 0.77%. The layout for a single pixel is presented in Fig. 4(a). When the SPAD is triggered by a photon, a voltage equal to the overvoltage potential, $\mathrm{V_{ov}}$, is applied across the gate of the PFET M4, causing it to turn on, triggering the output buffers. The threshold voltage of M4 is approximately 320 mV and the maximum gate voltage is 3.6 V, yielding a range of acceptable $\mathrm{V_{ov}}$ values from 0.32 V to 3.6 V as set by $\mathrm{V_{dd,diode}}$. The inverter U2 is connected to the core chip power supply of 1.5 V and level-shifts the output from M4, with a supply of $\mathrm{V_{dd,diode}}$, to this core logic voltage. Following a second inverter, U3, are two multiplexers for selecting among the SPAD output, an electrical calibration input (used for characterizing the TDCs), and

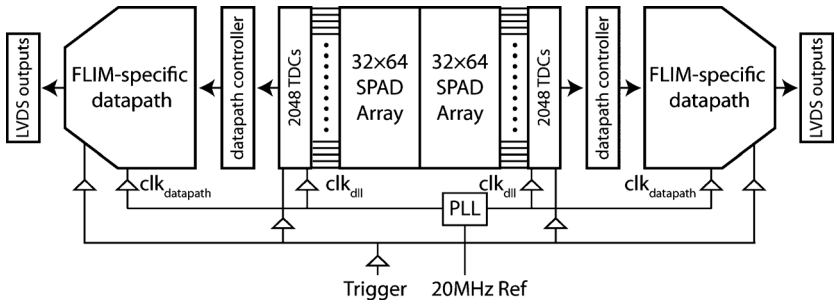

Fig. 2. Block diagram showing the major components of the imager chip.

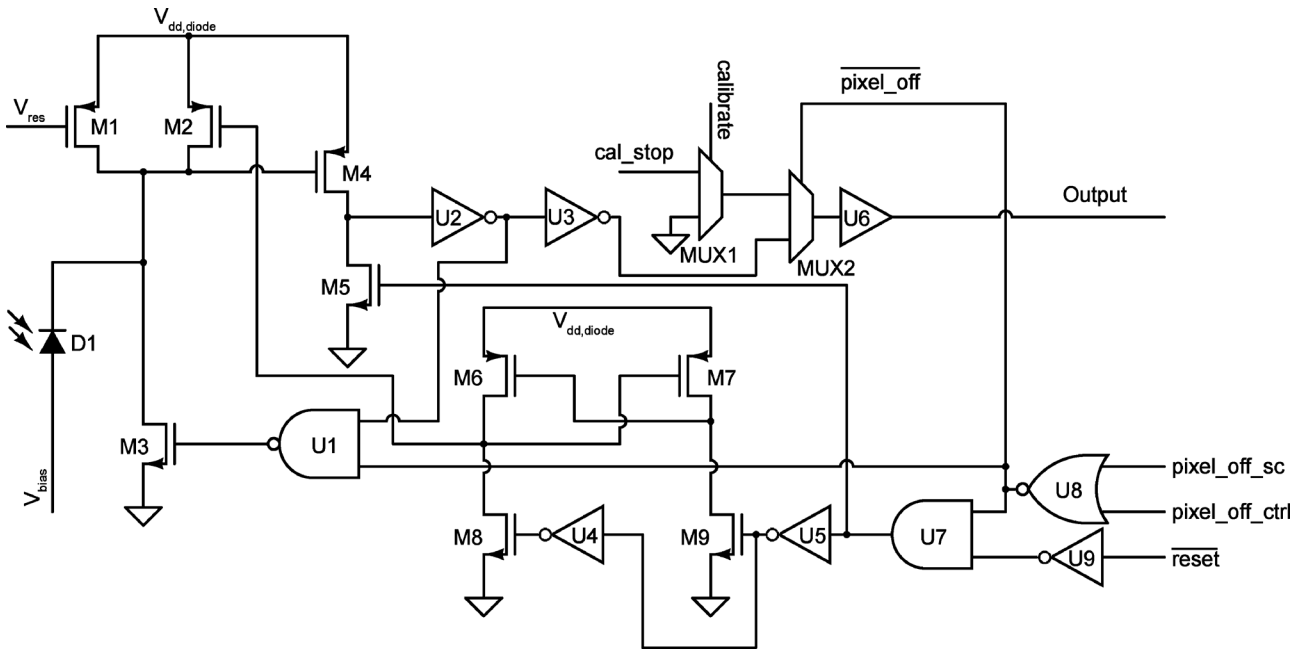

Fig. 3. Pixel circuit schematic that performs the quench, reset, TDC calibration, event output, and pixel control functions. Transistors M1-M9 and inverter U2 are designed using thick oxide devices.

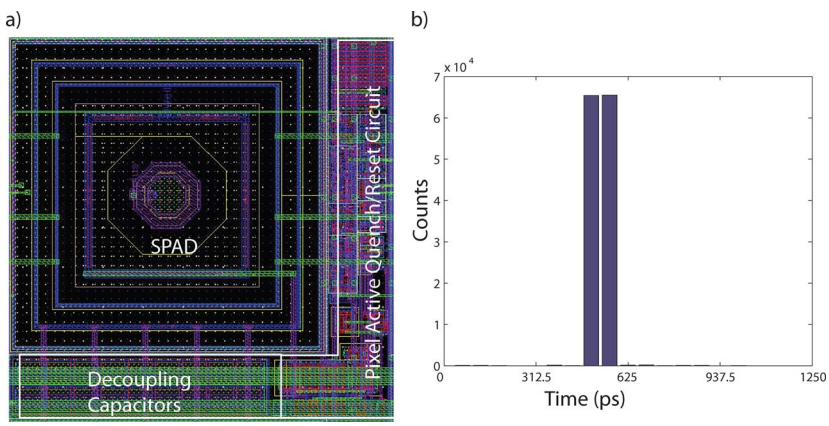

Fig. 4. (a) The layout of a single pixel in the array, including SPAD and the pixel control circuit of Fig. 3. The pixel and circuitry occupy an area that is 48 μm $\times$ 48 μm, of which a considerable amount is white space due to the conservative SPAD structure and guard rings used. (b) The impulse response of the SPAD as recorded by the on-chip TDCs is 125 ps. Each bar in the histogram represents a 62.5 ps wide timing bin.

a ground signal for turning off the output. The output from the multiplexers is then buffered to drive the stop signal of the TDC for this pixel.

When a photon event triggers a SPAD avalanche, the avalanche current must be stopped, or quenched, so that the SPAD can be reset and used in subsequent detection windows. In order to quench the device, the voltage across the SPAD must be reduced to below its breakdown voltage, $V_{\rm br}$. In this design, a PFET device, M1 in Fig. 3, is used as the quenching resistor. A tunable voltage, $V_{\rm res}$, is applied to the gate of M1,

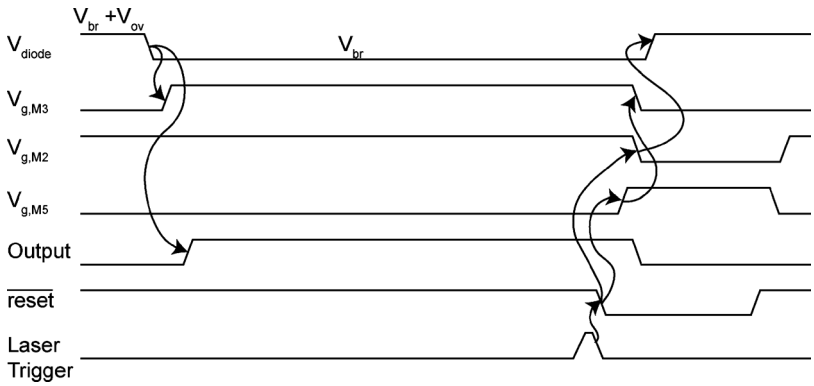

Fig. 5. Pixel circuit timing diagram showing a typical measurement and reset cycle. A pixel event occurs at the beginning of the cycle and triggers the output buffer. Immediately after the next laser pulse, $\overline{\text{reset}}$ is asserted and the SPAD recharges.

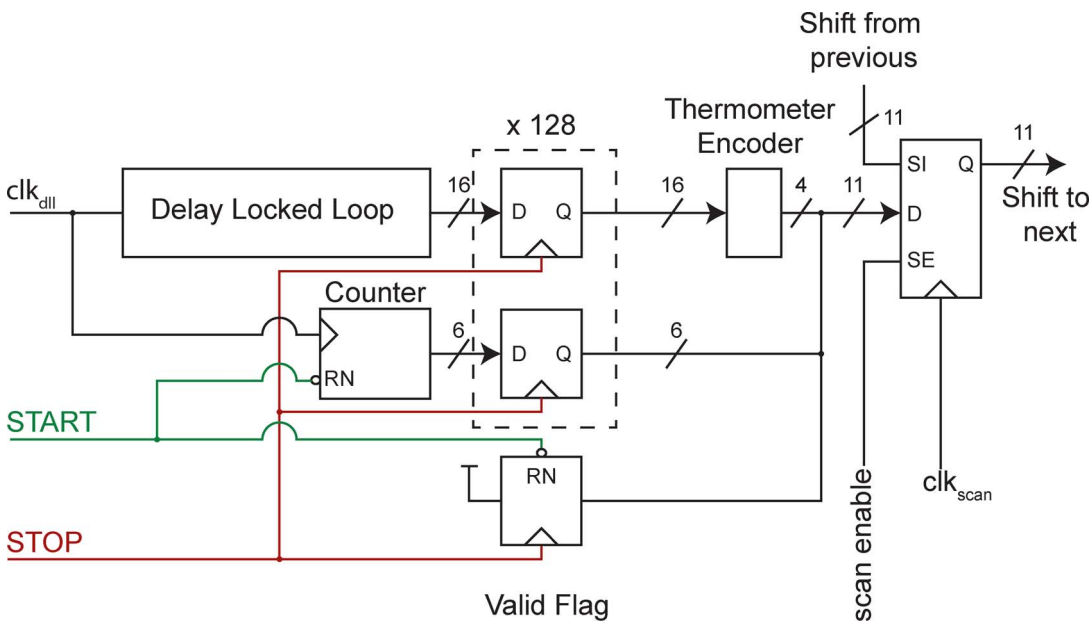

Fig. 6. A block level overview of the time-to-digital converter used in this work. A delay-locked loop subdivides the reference clock ($\mathrm{clk_{dll}}$) into 16 evenly spaced phases. $\mathrm{clk_{dll}}$ also increments a coarse counter. The phase and counter outputs are buffered to flip-flops to be used in TDCs for groups of 128 pixels. The thermometer encoder converts the 16-bit thermometer code into a 4 bit value, which, along with 6-bits from the counter and a valid data flag, is clocked into a chain of shift flip-flops at the end of each measurement window. $\mathrm{clk_{scan}}$ is a gated version of the datapath clock.

which allows the drain-to-source resistance, $\mathrm{R_{ds}}$, of the device to be adjusted from 10 kΩ to several MΩ as $\mathrm{V_{gs}}$ approaches the threshold voltage.

After the SPAD has been quenched, it must be reset before it can be used to detect another event. In an active quenching approach [21], [22], reset is performed by the wide-channel PFET device, M2, with a resistance between 1 kΩ and 400 Ω, depending on $\mathrm{V_{dd,diode}}$. In addition, the NFET, M3, is used to hold the bias across the SPAD below breakdown and can be used to prevent the SPAD from resetting or to disable the SPAD completely. M2 and M3 are independently controlled to minimize the probability for after-pulsing.

The $\overline{\text{reset}}$ signal, triggered by the laser pulse, passes through an AND gate, which allows for the option of disabling the pixel, and through a level-shifter that brings the $\overline{\text{reset}}$ signal to the $\mathrm{V_{dd,diode}}$ supply level. Device M2 is enabled and charges the cathode of the SPAD. When M5 turns on, it pulls the input to U2 low, and causes U1 to pull down, turning off M3. The timing diagram for event detection and reset is shown in Fig. 5.

Both pixel_off_sc and pixel_off_ctrl can be used to disable a pixel. The pixel_off_sc signal is a configuration bit that can be used to completely disable the pixel during all measurements. By using this control signal to disable abnormally noisy pixels, data bandwidth that would otherwise be used by these noise events is eliminated. The pixel_off_ctrl signal comes from a datapath controller and is used to disable the pixel at the end of a measurement window if no events have occurred. This feature reduces the impact of the dark count rate (DCR) on the lifetime measurement. A similar technique for defining the measurement window has recently been used to achieve extremely low noise levels in CMOS SPADs for measurements in which the photon arrival time is tightly constrained, like 3-D imaging [23].

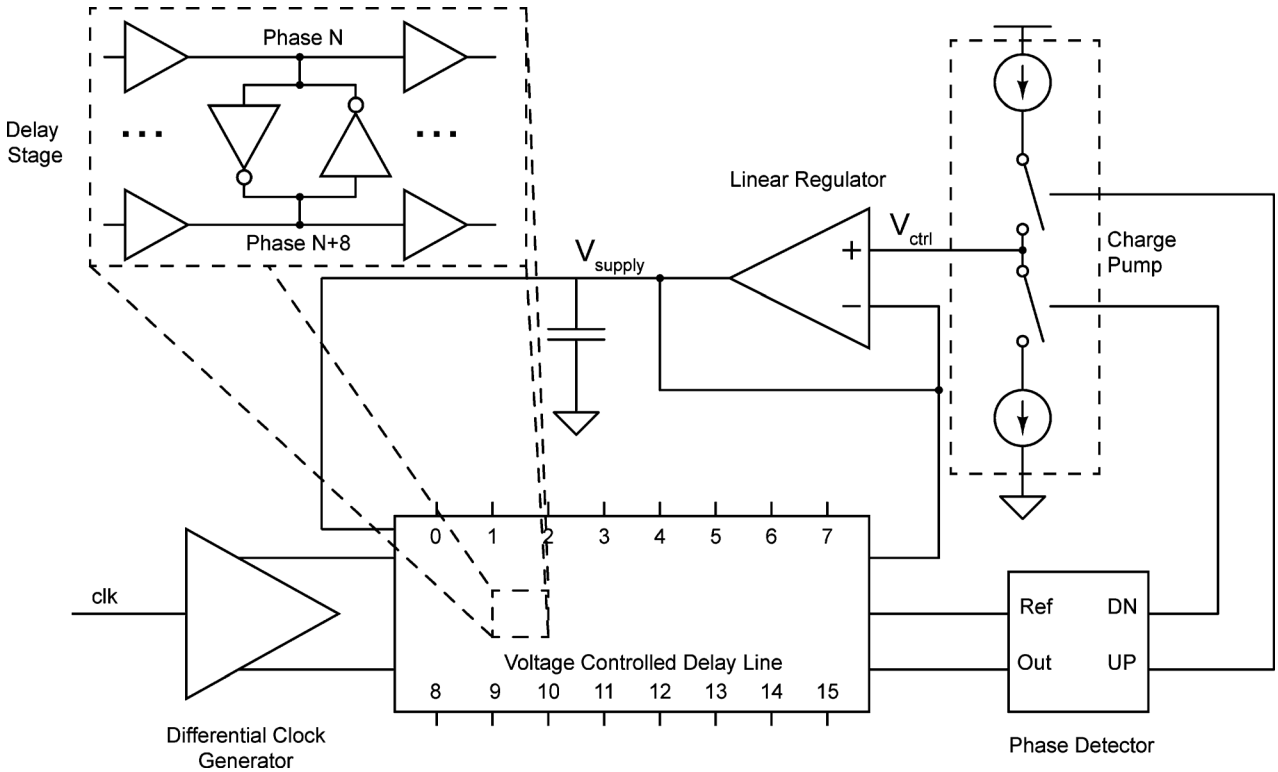

Fig. 7. Overview of delay-locked loop.

C. Time-to-Digital Converters

The TDC used in this work is based on a differential DLL architecture with a synchronous counter. This architecture provides a well-defined precision and dynamic range and fast conversion speed, and the DLL can be easily shared among groups of pixels in the array. In this design, the DLL and counter outputs are distributed to groups of 128 pixels, as shown in Fig. 6. The DLL [14] uses buffers as the delay element and includes a differential clock generator and charge pump (see Fig. 7).

A key concern when designing a differential voltage-controlled delay line (VCDL) with cross-coupled inverters between each stage is the matching of the two complementary clock edges. A slight difference in the $\mathrm{V_{dd}}/2$ crossing points of these complementary edges results in systematic timing errors. A complementary clock generator is designed using a pass-transistor circuit to align the crossing points of the complementary clocks (see Fig. 8(a)). The input buffer driving $\mathrm{clk_{in}}$ is designed such that its rise and fall times are greater than that of the inverter receiving $\mathrm{clk_{in}}$. Consequently, clk will be in the middle of its low-to-high transition when $\overline{\mathrm{clk_{in}}}$ reaches the pass-transistor gate. As a result, $\overline{\mathrm{clk_{in}}}$ is able to pass through its transmission gate with a constant low resistance because clk has already switched these transmission gates before $\overline{\mathrm{clk_{in}}}$ arrives, producing well-aligned transitions for both clk and $\overline{\mathrm{clk}}$. Simulation results including typical and skewed process corners are shown in Fig. 8(b)–(d).

The phase detector of the DLL generates equally sized UP and DN pulses when the input phases are aligned, such that any mismatch in the up and down currents of the charge pump will result in a static phase offset. Process-voltage-temperature (PVT) variations, in particular those that cause differences in the relative strength of NFETs and PFETs, can produce such offsets, necessitating calibration. The charge pump calibration control is designed such that the total combined differential width of the current mirror NFETs for the UP and DN currents can be adjusted in increments of 10 nm, corresponding to current steps of approximately 3 nA. A subset of the associated calibration coding scheme is shown in Table II.

A schematic of the charge pump is shown in Fig. 9, a switched, low-headroom, self-biased, cascoded current mirror. This architecture provides a high output resistance and closely matches an ideal current source. Both the UP and DN switches are implemented using only NFETs to minimize variability due to NFET-PFET process skew. Calibration control is implemented with six additional NFET devices in parallel with each of the switches (M24 – M35), allowing for fine current adjustments on the scale of 3 nA with minimum length devices. Replica biasing devices (M39 – M50) are also used to ensure that the UP and DN adjustment currents are well matched. 2.5-V devices are used, allowing the charge pump voltage to span the entire operating supply range of the VCDL from 0.75 V to 1.6 V. An off-chip reference current is used to bias the current mirrors.

D. Data-Compression Datapath

If a 10-bit time value is output for every pixel after each laser repetition with a laser pulse rate of 20 MHz, an off-chip data rate of 1.8 Tbps would be required for the array. This data rate, however, does not reflect the sparseness of the data. In particular, TCSPC experiments typically record a photon hit for only 1–2% of laser repetitions. Through the use of an event-driven

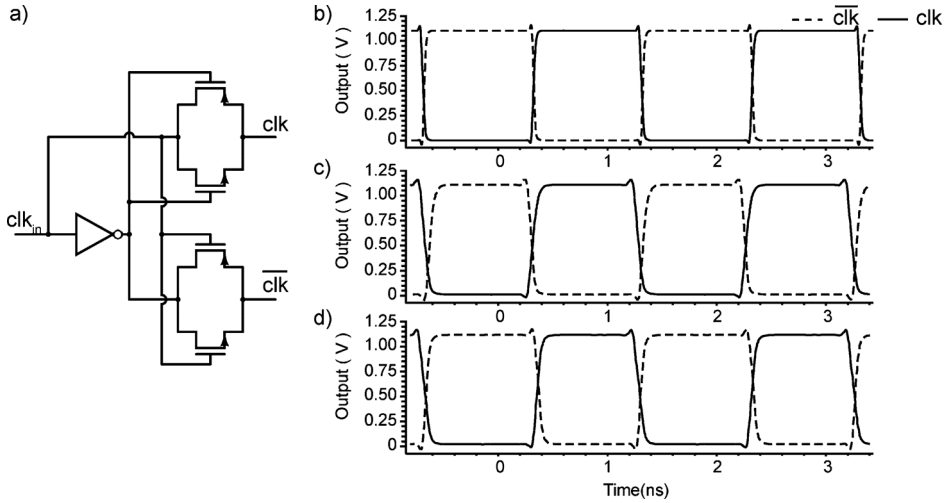

Fig. 8. (a) Schematic of the complementary clock generator. (b)–(d) Simulation results showing complementary clock edge alignment for nominal devices (b), skewed fast NFET, slow PFET devices (c), and skewed slow NFET, fast PFET devices (d).

TABLE II
SUBSET OF CALIBRATION CODES AND COMBINATIONS DEMONSTRATING THE WIDTH TUNING CAPABILITIES OF THE CALIBRATED CHARGE PUMP USED IN THIS DESIGN. THE DIFFERENTIAL VALUE OF THE CODES CONTINUOUSLY INCREASES IN INCREMENTS OF 10 NM FROM 0 TO 630 NM. ADDITIONAL DIFFERENTIAL WIDTHS BEYOND 630 NM CAN BE GENERATED FROM THESE CODES WITH THE MAXIMUM CALIBRATION DIFFERENCE AT 2470 NM. THE BITS IN THE CODE CORRESPOND TO THE FOLLOWING DEVICE WIDTHS IN ORDER FROM MSB TO LSB: 520 NM, 440 NM, 400 NM, 380 NM, 370 NM, 360 NM. THE MINIMUM DEVICE WIDTH FOR THESE TRANSISTORS IS 360 NM

Code 1	Code 2	$W_{total,1}$ (nm)	$W_{total,2}$ (nm)	Tuning Diff. (nm)	Current Diff. (nA)
000000	000000	0	0	0	0
000001	000010	360	370	10	3
000001	000100	360	380	20	6
000010	001000	370	400	30	9
000001	001000	360	400	40	12
000011	001100	730	780	50	15
000100	010000	380	440	60	19
...	...	...	...	...	...
000101	111000	740	1360	620	156
000011	111000	730	1360	630	159

readout approach, sparseness is exploited in our design to reduce the average data rate to approximately 18 Gbps. To achieve this, the time data for each pixel are appended with a valid bit that indicates whether a pixel event has occurred. This valid bit is used to control the flow of data out of the array such that only data associated with pixel events are allowed to pass.

The pipelined datapath shown in Fig. 2 is used to perform this data compression. At the end of a measurement window and before the next laser pulse occurs, all time data and associated valid bits are loaded into a set of registers, connected as shift registers on a half-row basis. The 10-bit time data are shifted out of each half-row into separate datapaths. A counter tracks the pixel position from which the data originated, resulting in a 5-bit position word appended to the time data, giving a combined 16-bit word (see Fig. 10). Within each datapath, the data for up to eight events per row are shifted into a bank of eight 16-bit registers as shown in Fig. 10. This happens within the 20 MHz laser pulse period that follows the one in which the data were captured.

On the rising edge of the next laser pulse trigger, the pixel event data is shifted in parallel into another set of registers and then shifted out as shown in Fig. 11. During this parallel shift operation an additional address bit is added to the data word in order to keep track of the row from which the data originated, increasing its length to 17 bits. A similar shifting process is repeated for each of the next four laser pulses (see Fig. 12). During each stage transition, an additional bit is added to the data word in order to indicate from which row it originated. Following Stage 5, the data are directly written into a first-in first-out (FIFO) buffer. At this point, the data words are 19-bits long with eight address bits. The diagram in Fig. 12 depicts the data for 8 half-rows of pixels. This identical datapath block is repeated 16 times on the chip.

The number of shift registers in each stage of the datapath is chosen such that an average pixel event rate of 1% of all laser pulses will result in an datapath overflow with a probability of less than 10^{-9}. The datapath is designed to operate at a clock frequency of 1 GHz and a laser pulse rate of 20 MHz, which provides 50 clock cycles within the datapath to complete all of the required shifting operations between stages.

E. Output Stage

After the valid pixel data has reached the FIFO at the end of each datapath, groups of four FIFOs are combined and their

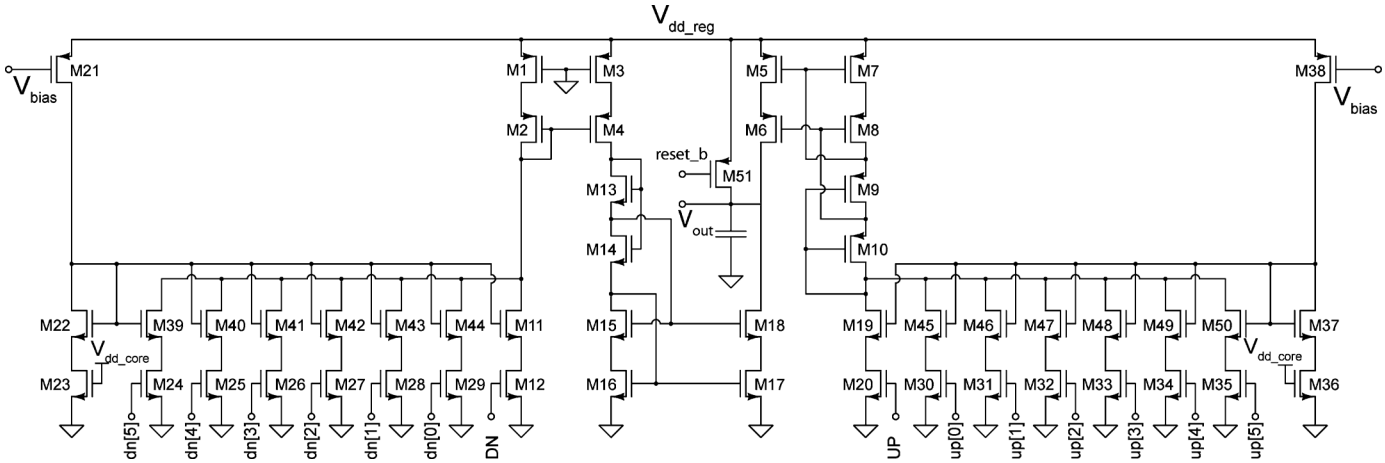

Fig. 9. Calibrated charge pump schematic.

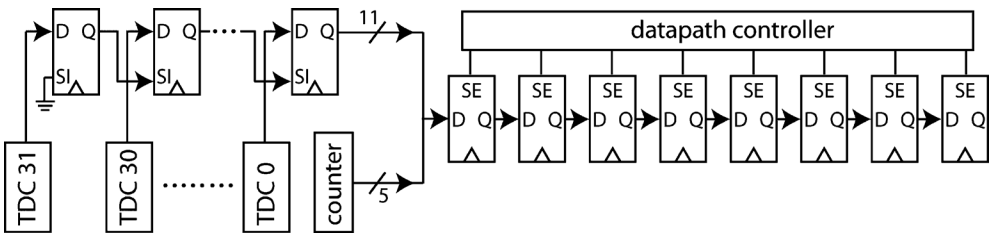

Fig. 10. The TDC data is shifted into a set of flip-flops at the end of each measurement window. This data is then shifted into the first stage of the datapath, which is sized to hold up to 8 pixel events per row.

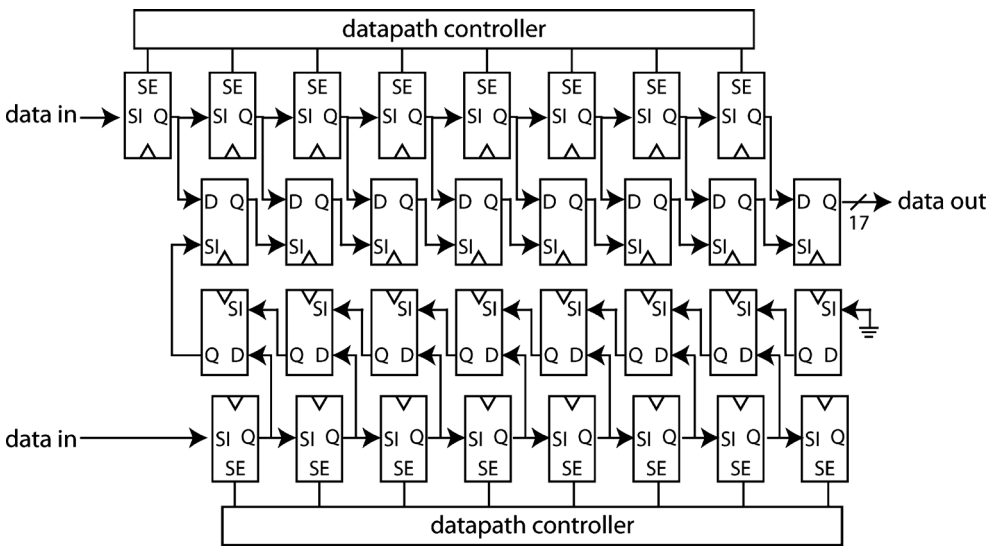

Fig. 11. The first stage of the datapath captures data shifted out of the TDCs and collects all valid pixel events. The datapath controller checks the valid bit of the incoming data and organizes the data into the rightmost flip-flops. After the data input shift is complete, the next laser pulse triggers a parallel shift operation into the central flip-flops that are connected in a U-shape. During the next measurement window, the data in these flip-flops will be scanned into the next stage of the datapath. The number of data bits increases by one to 17-bits, with the 17th bit representing the input row of the data.

data are transmitted over a bank of LVDS drivers, each designed to meet the TIA/EIA 644-A LVDS standard [24]. An output controller cycles between each FIFO in the group of four in a round-robin manner, adding two additional address bits that indicate the FIFO from which the data are retrieved. There are four banks of 22 LVDS buffers, a clock and 21 bits of data. These output drivers are capable of running at up to 500 MHz, providing a total output bandwidth of 42 Gbps.

F. System-Level Considerations

As shown in Fig. 1, each of the four LVDS banks communicates with a dedicated FPGA. In this design, four Virtex-6 XC6VLX130T-3 devices are used to capture the raw arrival time data and generate histograms of the arrival times for each pixel. The FPGA RAM is partitioned such that each 18-kb-block RAM is configured as a true dual-port memory and stores 128-

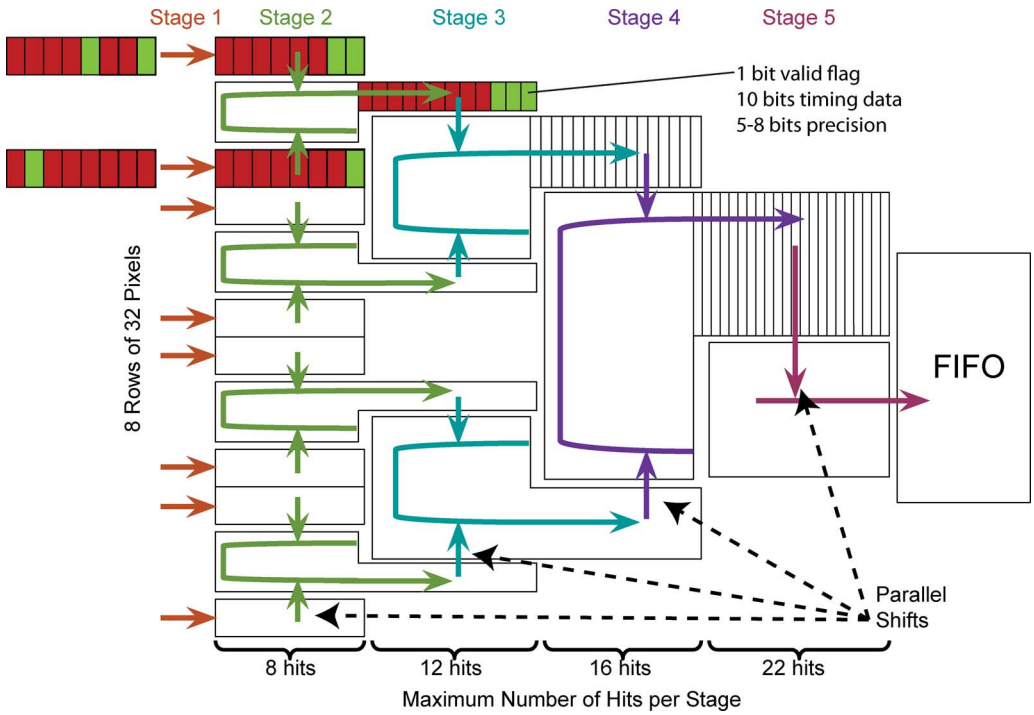

Fig. 12. Diagram showing the movement of data through the datapath. Each of these blocks is repeated sixteen times on the imager chip. An example showing how data is compressed in each stage is shown for the top two rows. Incoming valid data (green) initially has non-event data (red) between it. Each incoming data consists of 10-bits of timing information, 5-bits position information, and one-bit valid flag. As data enters the datapath, the non-event data is discarded, resulting in the two valid data packets in the top row finishing in the rightmost flip-flops. In the second stage, the two rows of data are shifted in parallel into the the U-shaped shift chain and then shifted clockwise with the non-event data being discarded once again. In this figure, all horizontal arrows represent serial data shifts while vertical arrows indicate a parallel data shift.

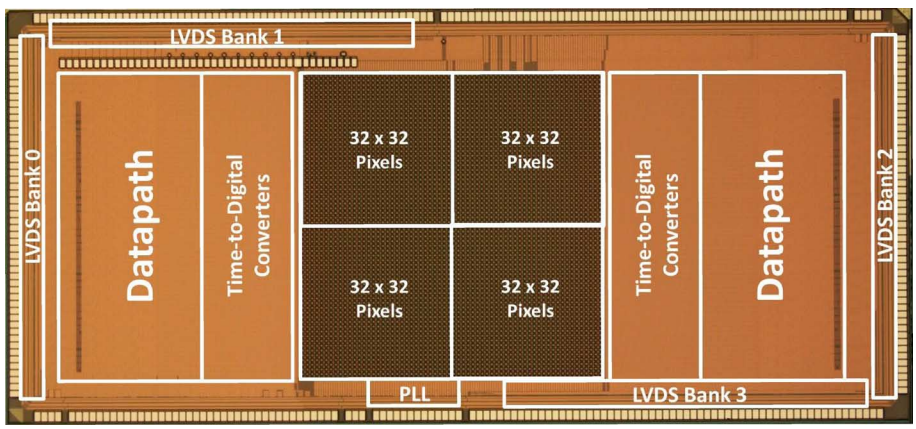

Fig. 13. Die photograph showing the major functional blocks of the FLIM imaging IC.

bin, 16-bit histogram information for two adjacent pixels in the array. Each pixel is allocated enough memory to record two histograms in the RAM such that one histogram can be read while the other is being captured, allowing for continuous recording of FLIM data. Each FPGA can process an incoming data stream at up to 10.5 Gbps.

Once the histograms have been formed, the data rate requirements for subsequent processing drop significantly. Each frame of the histogram dataset is 8 Mb, and the data rate requirement for transfer from the FPGAs to a computer scales with the desired frame rate. At 100 fps, the data rate to the computer is 800 Mbps. In order to reliably transfer data at this rate, we use PCIe interfaces on the Virtex-6 devices to perform direct memory access (DMA) writes from each FPGA directly to system memory on the computer. Each of the four FPGAs is configured with a x1 PCIe Gen 2 interface, which connects to a PCIe switch that combines four x1 links to a single x4 link. The switch used in this design is the PLX Technology PEX8608. We connect this x4 link using a cabled PCIe interface and x4 cabled PCIe adapter card. Using this PCIe interface, the system can support frame transfer rates of up to 754 fps.

III. Experimental Results

This design is fabricated in a standard 130-nm CMOS process, and a die photograph is presented in Fig. 13. Additionally, a printed circuit board (PCB) is designed with

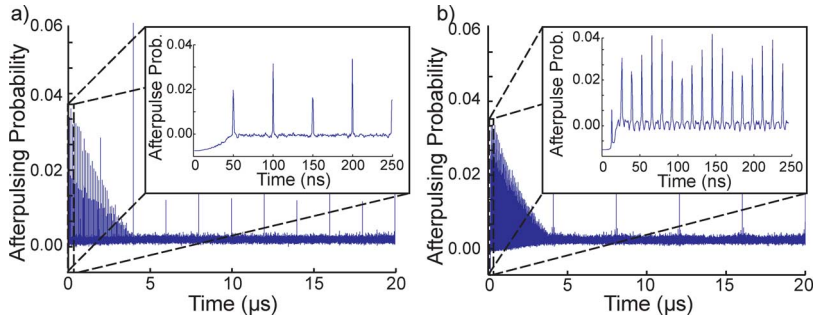

Fig. 14. Plots showing afterpulsing probabilities at (a) 20 MHz reset rate and (b) 100 MHz reset rate. Both measurements show periodic spikes in the afterpulsing probability plots at multiples of the reset frequency. For both measurements, V_{ov} was set at 3.5 V.

the appropriate high-speed PCIe interfaces and kernel driver modules to allow for DMA transfers between the FPGAs and system memory. Because of the complexity of the design, characterization is performed individually for each of the major system components.

A. Pixel Circuit Characterization

Located above the main SPAD array is an isolated pixel with the same circuitry as Fig. 3 but with its output connected directly to a pad for characterization. This pixel is used to evaluate the maximum count rate of our SPAD and the afterpulsing probability.

The maximum count rate for this device is evaluated using a bright, uncorrelated white light source with the standard reset rate of 20 MHz and a fast reset rate of 100 MHz. The test pixel is biased with an overvoltage, $V_{ov} = 3.5$ V. At a reset rate of 100 MHz, the pixel dead time, quench time, and reset time sum to 10 ns. The maximum count rate observed is 89.2 MHz.

The afterpulsing for the pixel using the active quench and reset circuitry is also evaluated at both 20-MHz and 100-MHz reset rates using uncorrelated white light. Afterpulsing probability can be measured by recording signal traces of pixel output pulses and computing the autocorrelation of the traces [25], [26]. The autocorrelation, $R(k)$, at the lag of k is given by

$$R(k) = \frac{1}{N-k} \sum_{n=1}^{N-k} \frac{x(n) \cdot x(n+k)}{\left(\frac{1}{N} \sum_{n=1}^{N} x(n) \right)^2} \qquad (1)$$

where N is the total number of samples used in the calculation and x is a discrete signal of pulse arrival times.

Fig. 14 shows the afterpulsing probabilities calculated from 3710 signal traces of 4000 ns with 800-ps precision for both 20 MHz and 100 MHz reset frequencies [26]. With either the 20-MHz or 100-MHz reset rate, afterpulsing probabilities are below 0.002 even with this SPAD biased with a relatively high V_{ov} of 3.5 V. The correlograms at both 20 MHz and 100 MHz reset frequencies do show periodic spikes, which are due to the synchronous SPAD reset at 50 ns and 10 ns intervals, respectively.

In addition to autocorrelation analysis, a histogram of the inter-spike interval (ISI) times can also be used to characterize afterpulsing. In a detector with afterpulsing, the histogram of the ISI times will show a bi-exponential decay, with a short decay time constant that is a consequence of afterpulsing and a long decay time constant that is related to the uncorrelated light source[27]. As seen in Fig. 15, the ISI histogram provides further evidence that afterpulsing is not significant for this detector.

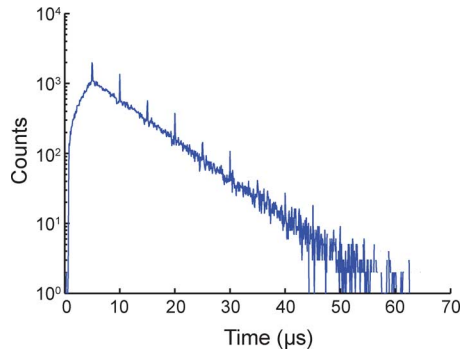

Fig. 15. Semi-log plot showing a histogram of the inter-spike intervals measured with a SPAD V_{ov} of 3.5 V and a reset rate of 20 MHz. The mono-exponential decay indicates that no afterpulsing is present. Spikes in the histogram are observed at multiples of the reset frequency.

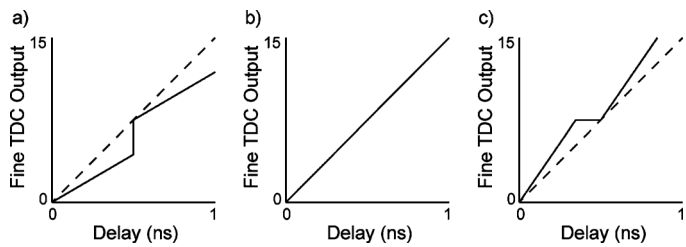

Fig. 16. Diagrams showing the different charge pump mismatch states. (a) The DN current is stronger than the UP current causing the VCDL to run slow and the TDC output to lag behind the input delay. This results in a vertical jump in the fine TDC transfer curve. (b) When the UP and DN currents are equal, the TDC output linearly tracks the delay input. (c) In the case when the UP current is stronger than the DN current, the VCDL runs fast and the TDC output leads the delay input. This causes a horizontal plateau in the fine TDC transfer curve. The charge pumps are calibrated by measuring this transfer curve and adjusting the UP and DN currents accordingly.

B. Time-to-Digital Converter Characterization

The TDCs are characterized and their charge pumps calibrated using the cal_stop signal in Fig. 3. To characterize the TDC, a 400-kHz reference signal is input into the trigger port of a Stanford Research Systems DG535 digital delay generator. Two outputs of the DG535 are used to generate the trigger signal and a tunable cal_stop signal with the delay between the trigger and cal_stop signal swept to characterize TDC performance.

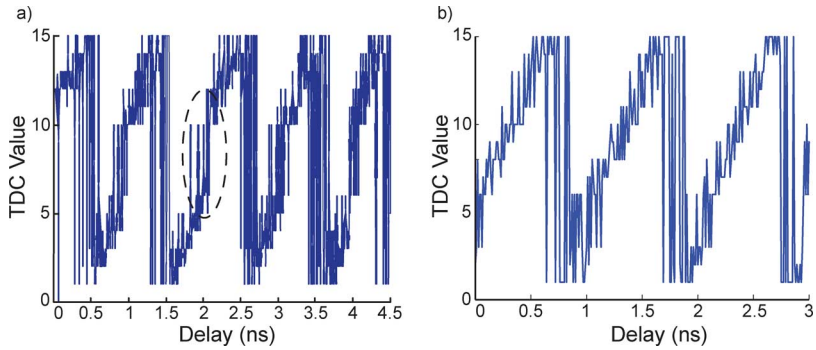

Fig. 17. (a) TDC output when the UP and DN charge pump calibration codes are 010000 and 010001, respectively. The dashed oval highlights the region of the transfer characteristic that indicates that the VCDL is running slowly (the DN is stronger than the UP current). (b) The TDC output after charge pump adjustments are made. The UP calibration code is 000101 and the DN calibration code is 101000. In this hardware, the UP devices are stronger than the DN devices.

The cal_stop signal is buffered on the PCB, which results in 200 ps of additional jitter in the measurement.

The charge pump in this design is manually calibrated by setting 12 control bits (6 bits each for UP and DN currents). Measurements of the fine TDC value while the start-stop delay of the cal_stop signal from the pixel into the TDC is adjusted are used to perform this calibration (see Fig. 16). Representative measurement results showing the adjustment of the fine TDC transfer characteristics through charge-pump calibration are shown in Fig. 17. In Fig. 17(a) the DN calibration bits have been set to 010001 and the UP calibration bits to 010000, creating a mismatch with the total DN current exceeding the UP. As a result, a vertical jump in the fine TDC transfer curve is observed (as described in Fig. 16(a)). In Fig. 17(b) the DN calibration bits are set to 101000 and the UP bits are set to 000101, which leads to the DN current matching the UP current. The TDC is also found to have a static delay offset of five fine-delay increments (312.5 ps) between the input to the DLL and the start of the coarse TDC counter increments, which is calibrated out digitally.

In order to measure the TDC linearity, the cal_stop delay is swept over the entire 64 ns range of the TDC. At each step of 10 ps, 10 samples are collected and averaged. The results of these measured delay sweeps are shown in Fig. 18. Fig. 18(a) shows the overall transfer curve for the TDC. The 200 ps measurement jitter is more than three times the LSB of the TDC, making accurate determination of the TDC linearity difficult. Subtracting the jitter, the measured DNL is better than ± 4 LSB (Fig. 18(d)). By using the code-density approach to determine the DNL with stop times derived from uncorrelated dark counts in the SPADs (eliminating the jitter from external measurement electronics), the maximum DNL is less than 2.27 LSB.

In the transfer curve in Fig. 18(a), periodic large spikes can also be observed. By separating the fine and coarse components of the TDC value in Fig. 18(d) and (e), it is clear that these spikes in the transfer curve are contributed by the coarse counter. This artifact is due to metastability brought on by the use of an asynchronous stop signal to latch counter values into flip-flops, as shown in Fig. 6. Because of these spikes, the INL is slightly more than 8 LSB. This error could be corrected by simply synchronizing the stop signal for the counter with the counter clock as in Fig. 19. Lifetime measurements presented in Section III.C show that this non-ideality does not significantly impact the imaging performance.

TABLE III
SUMMARY OF IC CHARACTERISTICS. THE DETAIL FOR THE AVERAGE POWER CONSUMPTION IS PROVIDED FOR EACH OF THE ON-CHIP SUPPLIES. $V_{dd,core}$ IS THE MAIN 1.5 V CORE SUPPLY AND POWERS ALL OF THE CONTROL LOGIC AND THE DATAPATH, $V_{dd,IO}$ IS THE 2.5 V I/O POWER SUPPLY FOR THE FOUR LVDS BANKS, $V_{dd,reg}$ IS THE 2.5 V POWER SUPPLY FOR THE VOLTAGE REGULATORS IN THE DLL, $V_{dd,diode}$ IS THE NEGATIVE BIAS FOR THE SPADS AND VARIES DEPENDING ON BIAS BUT IS TYPICALLY AROUND -12 V, $V_{dd,pll}$ IS A 1.6 V SUPPLY FOR THE PLL. THE POWER FOR BOTH $V_{dd,core}$ AND $V_{dd,diode}$ VARY WITH THE INCIDENT PHOTON FLUX AND THE VALUES OF THE TABLE ARE FOR A 'TYPICAL' FLUX THAT RESULTS IN A HIT RATE OF APPROXIMATELY 1%

Characteristic	Value	Unit
Peak PDP @ V_{ov}=1.5	30 [20]	%
Peak PDP Wavelength	425 [20]	nm
Average DCR @ V_{ov}=2.5 and Room Temp.	544	Hz
Afterpulsing Probability	<0.002	
TDC Resolution	62.5	ps
TDC Range	64	ns
TDC DNL	<4	LSB
TDC INL	<8	LSB
IC Output Bandwidth	42	Gbps
System Output Bandwidth	6.325	Gbps
Maximum Data Limited Frame Rate	466	fps
Maximum Attained FLIM Frame Rate	100	fps
Average Power IC	8.79	W
$V_{dd,core}$	6.10	W
$V_{dd,IO}$	2.15	W
$V_{dd,reg}$	0.125	W
$V_{dd,diode}$	0.35	W
$V_{dd,pll}$	0.06	W
Average Power System	26.4	W

The impulse response function (IRF) of the pixel and TDC combined was measured by repeatedly triggering the SPADs using a 500 nm band from a Fianium supercontinuum laser with pulse width of 10 ps. The outputs of the TDC were collected and formed into histograms, which resulted in a distribution with a peak of only two LSB, or 125 ps (Fig. 4(b)).

C. Imaging Array Performance

The imager draws a total of 8.79 W when running at full-speed and is water-cooled to avoid degraded performance of the SPADs due to heating [28]. To achieve this cooling, a custom BGA package with a copper core is directly soldered to the PCB, which also has a copper core. The dark count rate (DCR) for the

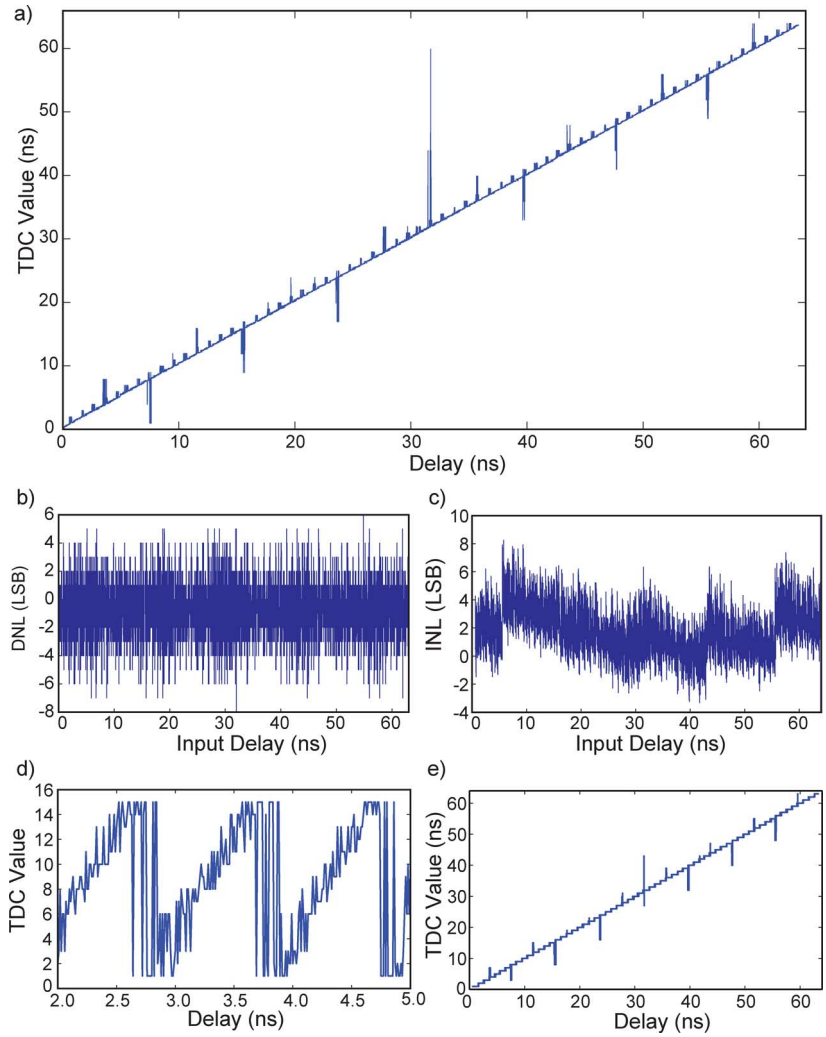

Fig. 18. Plots showing the (a) transfer curve of the TDC, (b) DNL of the TDC and (c) INL of the TDC. The transfer curves of fine and coarse components of the TDC value are presented in (d) fine and (e) coarse. In the DNL and INL plots, we have filtered the large spikes in the coarse TDC measurement by removing TDC value changes greater than 1 ns between any two of the 10 ps input delay steps.

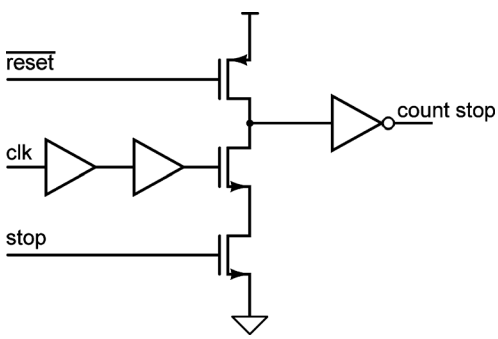

Fig. 19. TDC counter stop signal synchronizer.

system with cooling is 544 Hz, which is an improvement from 1036 Hz without cooling. Fig. 20 shows the distribution of DCR throughout the array. The overvoltage, V_{ov} for this measurement is 2.5 V and is consistent with the measurements taken in [20].

From this DCR data, a clear pattern in the number of hits recorded can be seen within groups of 8 rows of pixels, in which the seventh and eighth rows record lower counts. We attribute this to a voltage droop in the power distribution biasing the SPADs. A similar pattern is observed in the lifetime images and results in missing rows due to the pixels receiving an insufficient number of hits for lifetime extraction. This voltage droop problem does not affect the lifetime extraction for the other pixels within the array.

In preliminary testing of the imaging performance of the array, we use a ceramic cover to mask a portion of the SPAD array and then place a dish of fluorescein dye over the array and image using a 488 nm excitation wavelength filtered from a Fianium supercontinuum pulsed light source with a pulse repetition rate of 4 MHz. A 550-nm emission filter is also employed. Figs. 22 and 23 show the setup and the imaging results, respectively. The resulting image matches the expected fluorescence lifetime for fluorescein of 4–5 ns [29].

In order to test the fast acquisition capabilities of our system, we capture a total of 16 consecutive frames with each being acquired in 10 ms for a frame rate of 100 fps. The sixteen frames from this experiment are shown sequentially in Fig. 24. At this

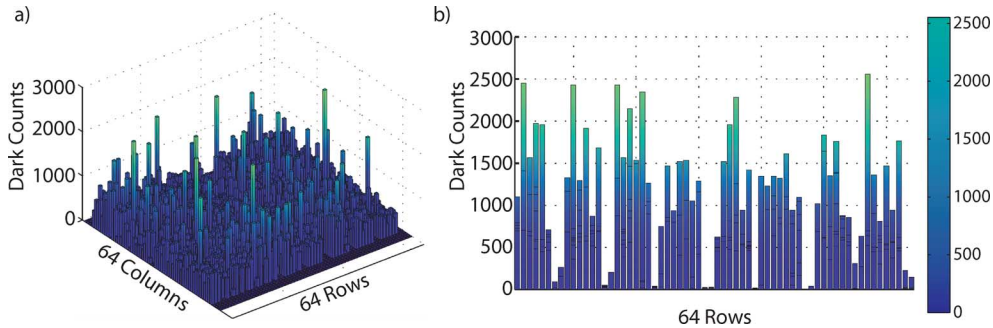

Fig. 20. (a) A plot showing the dark count rate distribution throughout the imaging array. (b) A cross-sectional view showing the periodic reduction in count rate across the array.

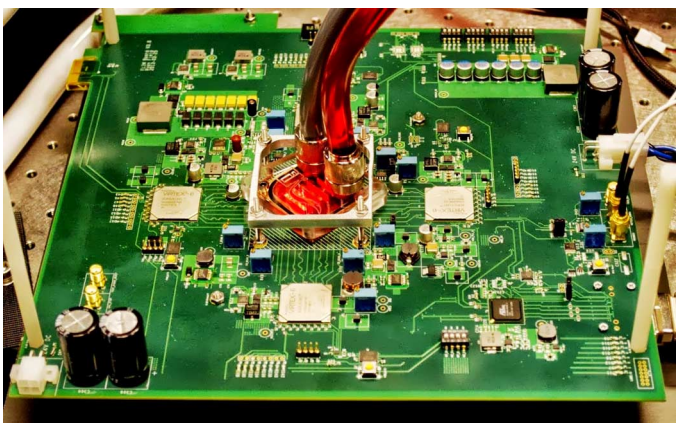

Fig. 21. A photograph of the FLIM PCB with FPGAs and the cabled PCIe interface. The liquid cooling system is in the center of the PCB with the imager mounted on the bottom side of the board.

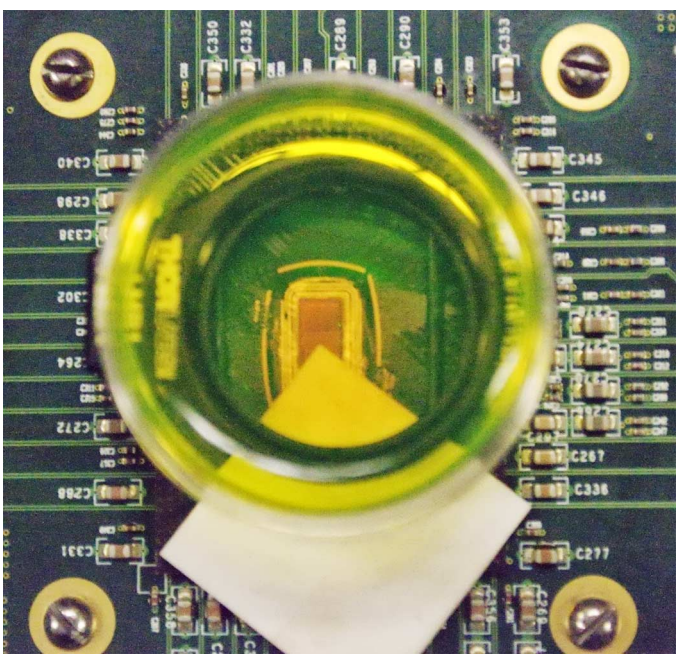

Fig. 22. A photograph showing the arrangement of the ceramic mask, band-pass filter, and fluorescein dye used in the imaging experiments. The excitation source is not shown but would be directed into the page over the sample.

high frame rate, there is more error in the lifetime estimation due to the limited amount of data that is collected for each frame. By increasing the image acquisition time, more accurate and uniform lifetime estimation is possible, as in Fig. 23 where an imaging rate of 50 fps is used. Currently the number of consecutive frames that can be captured is limited by the kernel driver we have written to handle the DMA transfer from the FPGAs to the computer. The techniuqe employed requires a coherent block of memory allocated on the computer for each of the FPGAs that can hold 16 frames of FLIM data. The CPU waits a predetermined amount of time that is based on the frame rate before reading this block of memory and performing additional processing. In order to leverage the full potential of the system, our driver must be modified to handle interrupt signals over the PCIe interface from the FPGAs such that the CPU can be instructed to process the data and free the memory so that additional frames can be written.

IV. Conclusions

In this work, we have built the fastest published TCSPC-based fluorescence lifetime imaging system, which is capable of acquiring FLIM images at 100 fps. The system consists of a FLIM-specific integrated circuit and a custom system architecture that are optimized for high-speed imaging. The integrated circuit contains 4096 SPADs with independent TDC channels allowing for fully parallel time-of-arrival recording. The timing resolution of the TDCs is 62.5 ps and they can record arrival times for up to 64 ns after a triggering laser pulse, supporting a range of possible fluorescence decay rates. A data-compression datapath provides a mechanism for efficiently transmitting data off-chip in an event-driven manner. The imager chip can support an output data bandwidth of up to 42 Gbps allowing for a maximum FLIM imaging rate of over 400 fps. The total power consumption of the IC is 8.79 W, or 2.15 mW/pixel, including the TDCs and all data handling circuitry.

We designed an FPGA-based system to capture the raw arrival time data from the imaging IC and generate a histogram for each pixel in the image. This histogram data can be transferred to a PC at over 750 fps. The entire system (including imager) consumes 26.4 W. A summary of our FLIM imaging system is presented in Table III. The unmatched performance of our design is a direct result of the circuit optimizations made to efficiently handle the FLIM data. Further improvements on this system could improve the lifetime variability between pixels and allow for a greater number of consecutive frames to be acquired at full speed. The image acquisition time achieved in this

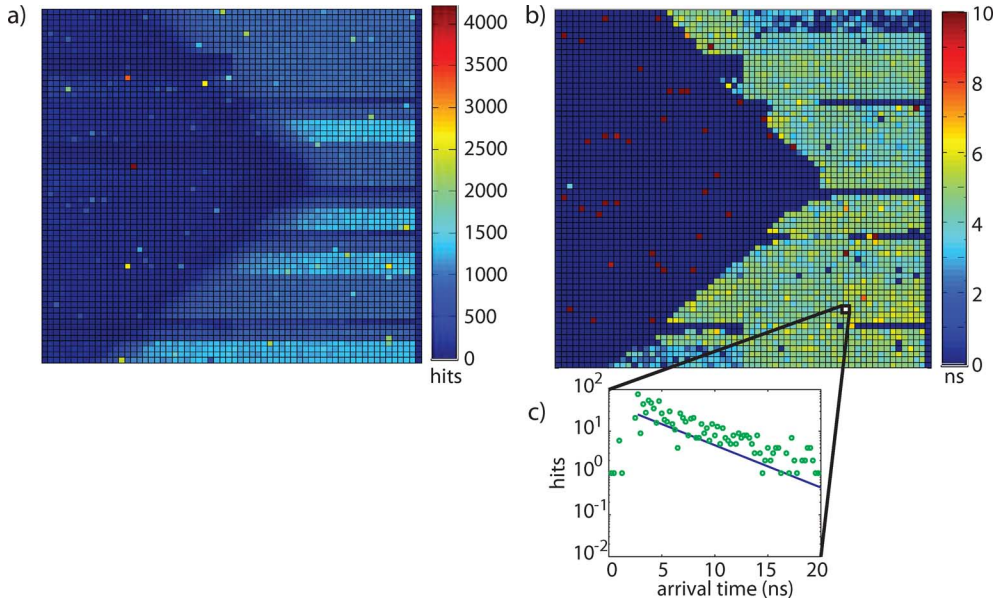

Fig. 23. (a) Intensity image from fluorescein dye measurement. (b) Lifetime image showing the well resolved mask. c.) A representative lifetime decay from pixel (49,10) in the image.

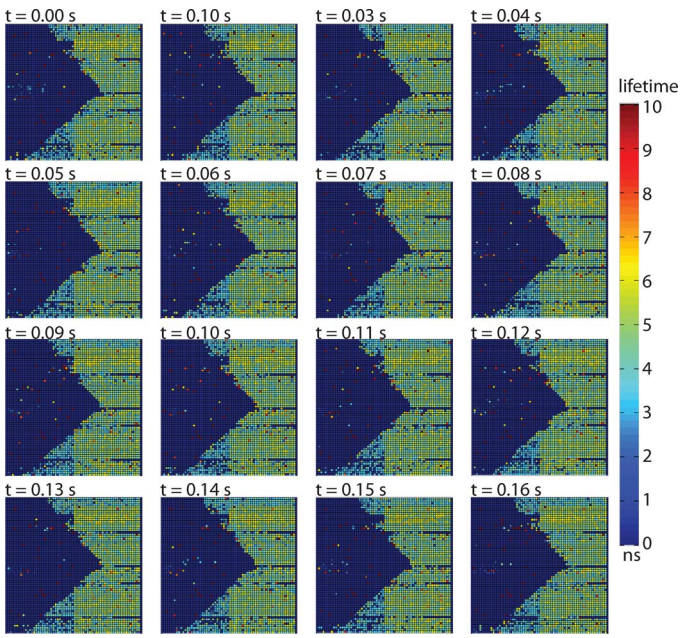

Fig. 24. Sixteen consecutive frames capture using our lifetime imaging system with an acquisition time of only 10 ms per frame. Color bar indicates the extracted lifetime value.

work has the potential to enable a wide range of FLIM applications involving dynamic samples.

REFERENCES

[1] A. Pietraszewska-Bogiel and T. W. J. Gadella, "FRET microscopy: From principle to routine technology in cell biology," *J. Microscopy*, vol. 241, pp. 111–8, Feb. 2011.

[2] J. C. Waters, "Accuracy and precision in quantitative fluorescence microscopy," *J. Cell Biol.*, vol. 185, pp. 1135–48, Jun. 2009.

[3] R. Sanders, A. Draaijer, H. Gerritsen, P. Houpt, and Y. Levine, "Quantitative pH imaging in cells using confocal fluorescence lifetime imaging microscopy," *Analyt. Biochem.*, vol. 227, no. 2, pp. 302–308, 1995.

[4] K. Suhling, P. M. W. French, and D. Phillips, "Time-resolved fluorescence microscopy," *Photochem. Photobiolog. Sci.*, vol. 4, pp. 13–22, Jan. 2005.

[5] D. Elson, J. Requejo-Isidro, I. Munro, F. Reavell, J. Siegel, K. Suhling, P. Tadrous, R. Benninger, P. Lanigan, J. McGinty, C. Talbot, B. Treanor, S. Webb, A. Sandison, A. Wallace, D. Davis, J. Lever, M. Neil, D. Phillips, G. Stamp, and P. French, "Time-domain fluorescence lifetime imaging applied to biological tissue," *Photochem. Photobiolog. Sci.*, vol. 3, pp. 795–801, Aug. 2004.

[6] M. S. Kim, B.-K. Cho, A. M. Lefcourt, Y.-R. Chen, and S. Kang, "Multispectral fluorescence lifetime imaging of feces-contaminated apples by time-resolved laser-induced fluorescence imaging system with tunable excitation wavelengths," *Appl. Opt.*, vol. 47, pp. 1608–1616, Mar. 2008.

[7] J. McGinty, N. P. Galletly, C. Dunsby, I. Munro, D. S. Elson, J. Requejo-Isidro, P. Cohen, R. Ahmad, A. Forsyth, A. V. Thillainayagam, M. a. a. Neil, P. M. W. French, and G. W. Stamp, "Wide-field fluorescence lifetime imaging of cancer," *Biomed. Opt. Expr.*, vol. 1, pp. 627–640, Jan. 2010.

[8] V. V. Ghukasyan and F.-J. Kao, "Monitoring cellular metabolism with fluorescence lifetime of reduced nicotinamide adenine dinucleotide," *J. Phys. Chemist. C*, vol. 113, pp. 11532–11540, Jul. 2009.

[9] J. W. Borst and A. J. W. G. Visser, "Fluorescence lifetime imaging microscopy in life sciences," *Measure. Sci. Technol.*, vol. 21, p. 102002, Oct. 2010.

[10] X. F. Wang, A. Periasamy, B. Herman, and D. Coleman, "Fluorescence lifetime imaging microscopy (FLIM): Instrumentation and applications," *Critical Rev. Analyt. Chemist.*, vol. 23, no. 5, pp. 369–395, 1992.

[11] C. Chang, D. Sud, and M. Mycek, "Fluorescence lifetime imaging microscopy," *Methods Cell Biol.*, vol. 81, no. 06, pp. 495–524, 2007.

[12] C. Harris and B. Selinger, "Single-photon decay spectroscopy. II, The Pile-up problem," *Australian J. Chemist.*, vol. 32, pp. 2111–2129, 1979.

[13] D. D.-U. Li, J. Arlt, D. Tyndall, R. Walker, J. Richardson, D. Stoppa, E. Charbon, and R. K. Henderson, "Video-rate fluorescence lifetime imaging camera with CMOS single-photon avalanche diode arrays and high-speed imaging algorithm," *J. Biomed. Opt.*, vol. 16, no. 9, p. 096012, 2011.

[14] D. E. Schwartz, E. Charbon, and K. L. Shepard, "A single-photon avalanche diode array for fluorescence lifetime imaging microscopy," *IEEE J. Solid-State Circuits*, vol. 43, pp. 2546–2557, Nov. 2008.

[15] L. Pancheri and D. Stoppa, "A SPAD-based pixel linear array for high-speed time-gated fluorescence lifetime imaging," in *Proc. ESSCIRC*, Sep. 2009, pp. 428–431.

[16] E. Charbon, "Highly sensitive arrays of nano-sized single-photon avalanche diodes for industrial and bio imaging," in *Proc. Nano-Net*, 2009, pp. 161–168.
[17] M. Gersbach, Y. Maruyama, E. Labonne, J. Richardson, R. Walker, L. Grant, R. Henderson, F. Borghetti, D. Stoppa, and E. Charbon, "A parallel 32x32 time-to-digital converter array fabricated in a 130 nm imaging CMOS Technology," *Proc. ESSCIRC*, pp. 196–199, Sep. 2009.
[18] C. Veerappan, J. Richardson, R. Walker, D.-U. Li, M. W. Fishburn, Y. Maruyama, D. Stoppa, F. Borghetti, M. Gersbach, R. K. Henderson, and E. Charbon, "A 160×128 single-photon image sensor with on-pixel 55 ps 10b time-to-digital converter," in *IEEE Int. Solid-State Circuits Conf. Dig., ISSCC*, 2011, pp. 312–314.
[19] C. Niclass, M. Sergio, and E. Charbon, W. Becker, Ed., "A single photon avalanche diode array fabricated in 0.35-μm CMOS and based on an event-driven readout for TCSPC experiments," in *Proc. SPIE: Advanced Photon Counting Techniq.*, Oct. 2006, vol. 6372, pp. 63720S–63720S-12.
[20] R. M. Field, J. Lary, J. Cohn, L. Paninski, and K. L. Shepard, "A low-noise, single-photon avalanche diode in standard 0.13 μm complementary metal-oxide-semiconductor process," *Appl. Phys. Lett.*, vol. 97, no. 21, p. 211111, 2010.
[21] N. S. Nightingale, "A new silicon avalanche photodiode photon counting detector module for astronomy," *Experimental Astronomy*, vol. 1, no. 6, pp. 407–422, 1990.
[22] F. Zappa, A. Lotito, A. Giudice, S. Cova, and M. Ghioni, "Monolithic active-quenching and active-reset circuit for single-photon avalanche detectors," *IEEE J. Solid-State Circuits*, vol. 38, pp. 1298–1301, Jul. 2003.
[23] E. Vilella and A. Diéguez, "A gated single-photon avalanche diode array fabricated in a conventional CMOS process for triggered systems," *Sensors Actuators A: Phys.*, vol. 186, pp. 163–168, Oct. 2012.
[24] "Electrical Characteristics of Low Voltage Differential Signaling (LVDS) Interface Circuits," TIA/EIA-644-A, 2001.
[25] J. C. Jackson, D. Phelan, A. P. Morrison, R. M. Redfern, and A. Mathewson, "Characterization of geiger mode avalanche photodiodes for fluorescence decay measurements," in *Photodetector Materials and Devices VII*, May 2002, vol. 4650, pp. 55–66.
[26] R. G. Brown, R. Jones, J. G. Rarity, and K. D. Ridley, "Characterization of silicon avalanche photodiodes for photon correlation measurements. 2: Active quenching," *Appl. Opt.*, vol. 26, pp. 2383–9, Jun. 1987.
[27] M. W. Fishburn, *Fundamentals of CMOS Single-Photon Avalanche Diodes*. Delft, The Netherlands: TU Delft, 2012.
[28] S. Cova, M. Ghioni, A. Lacaita, C. Samori, and F. Zappa, "Avalanche photodiodes and quenching circuits for single-photon detection," *Appl. Opt.*, 1996.
[29] D. Magde, G. E. Rojas, and P. G. Seybold, "Solvent dependence of the fluorescence lifetimes of xanthene dyes," *Photochemist. Photobiol.*, vol. 70, no. 5, p. 737, 1999.
[30] M. Gersbach, Y. Maruyama, R. Trimananda, M. W. Fishburn, D. Stoppa, J. A. Richardson, R. Walker, R. Henderson, and E. Charbon, "A time-resolved, low-noise single-photon image sensor fabricated in deep-submicron CMOS technology," *IEEE J. Solid-State Circuits*, vol. 47, pp. 1394–1407, Jun. 2012.
[31] J. Richardson, R. Walker, L. Grant, D. Stoppa, F. Borghetti, E. Charbon, M. Gersbach, and R. K. Henderson, "A 32x32 50 ps resolution 10 bit time to digital converter array in 130 nm CMOS for time correlated imaging," in *Proc. IEEE Custom Integrated Circuits Conf., CICC*, 2009, pp. 77–80, 029217.
[32] C. Niclass, C. Favi, T. Kluter, M. Gersbach, and E. Charbon, "A 128x128 single-photon image sensor with column-level 10-bit time-to-digital converter array," *IEEE J. Solid-State Circuits*, vol. 43, pp. 2977–2989, Dec. 2008.

Ryan M. Field received the B.S. degree in electrical engineering and the B.S. degree in physics from North Carolina State University, Raleigh, NC, in 2007 and the M.S. and Ph.D. degrees from Columbia University, New York, NY in 2008 and 2013, respectively. His Ph.D. work focused on integrated circuit and system-level design for improved biomedical imaging techniques.

After graduating in 2013, he joined Intel Corporation, Santa Clara, CA, USA, as a research scientist in the Integrated Biosystems Lab and is interested in electronic detection of biological signals. He was the recipient of the Astronaut Scholarship in 2006 and the National Science Foundation Graduate Research Fellowship, the National Defense Science and Engineering Graduate Fellowship, and the Columbia University Fu Foundation School of Engineering Presidential Distinguished Fellowship in 2007. He was also an associate in the Columbia University BioIGERT Program.

Simeon Realov received the B.S. degree in engineering from Swarthmore College, Swarthmore, PA, in 2006 and the M.S. and Ph.D. degrees in electrical engineering from Columbia University, New York, NY, in 2007 and 2012, respectively. During his time at Columbia University, he worked on various methods for detailed on-chip device characterization and statistical modeling of device variability in advanced CMOS processes.

He held internship positions at IBM T. J. Watson Research Center, Yorktown Heights, NY, where he developed techniques for on-chip capacitance measurement, combined C-V/I-V transistor characterization, as well as a fully-integrated SRAM fluctuations monitoring circuit. He also held a summer internship position with Rambus, Inc, Sunnyvale, CA, where he worked on the design of a high-bandwidth on-chip supply noise monitoring system. After his graduation from Columbia University in 2012, Dr. Realov joined Intel Corporation, Hillsboro, OR, where he is currently a member of the Advanced Design Library group working on digital standard cell design.

Kenneth L. Shepard (F'08) received the B.S.E. degree from Princeton University, Princeton, NJ, USA, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1988 and 1992, respectively.

From 1992 to 1997, he was a Research Staff Member and Manager with the VLSI Design Department, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, where he was responsible for the design methodology for IBM's G4 S/390 microprocessors. Since 1997, he has been with Columbia University, New York, NY, USA, where he is now Professor of Electrical Engineering and Biomedical Engineering. He also was Chief Technology Officer of CadMOS Design Technology, San Jose, CA, USA, until its acquisition by Cadence Design Systems in 2001. His current research interests include power electronics, carbon-based devices and circuits, and CMOS bioelectronics.

Dr. Shepard was Technical Program Chair and General Chair for the 2002 and 2003 International Conference on Computer Design, respectively. He has served on the Program Committees for International Electron Devices Meeting (IEDM), International Solid-State Circuits Conference (ISSCC), VLSI Symposium, International Conference on Computer-Aided Design (ICCAD), Design Automation Conference (DAC), International Symposium on Circuits and Systems (ISCAS), International Symposium on Quality Electronic Design (ISQED), Great Lakes Symposium on VLSI (GLS-VLSI), and International Conference on Computer Design (ICCD). He received the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992, a National Science Foundation CAREER Award in 1998, and the 1999 Distinguished Faculty Teaching Award from the Columbia Engineering School Alumni Association. He has been an Associate Editor of the IEEE TRANSACTIONS ON VERY LARGE-SCALE INTEGRATION (VLSI) SYSTEMS and is currently an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS and the IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS.