The CMOS APS Digital Camera Sensor

Douglas A. Kerr, P.E.

Issue 2.1
June 8, 2006

ABSTRACT

Many digital still cameras today use a "CMOS" sensor array. The name comes from the fact that this type of sensor uses the same construction, and can be fabricated with much the same technique, as the familiar CMOS (complementary metal-oxide-semiconductor) integrated circuit chip. However, the designation also implies an architecture and readout technique dramatically different from that of the other popular sensor type, the CCD (charge-coupled device). In this article we discuss the principles and operation of an important form of the CMOS sensor, the active pixel sensor (APS) form. We also describe a specific application of the design in the sensor used by Canon, Inc. in various of its digital SLR cameras.

INTRODUCTION

Many digital still cameras today use a "CMOS" sensor array. The name comes from the fact that this type of sensor can be built in a process that largely derives from, and can be built on the same toolset as, CMOS (complementary metal-oxide-semiconductor) logic devices, but with some sensor-specific modifications.

But these sensors are unique far beyond the matter of their construction and fabrication. They utilize a photodetection, pixel readout, and digitization architecture dramatically different from that used in the other common digital camera sensor family, the CCD (charge-coupled device) type.

We will examine here the architecture and operation of an important subtype of the CMOS sensor class, the active pixel sensor (APS)[1] type. This is so-called because the individual pixel cells incorporate active elements (transistors).

The form we will examine uses three such transistors per pixel cell (and accordingly is sometimes called the "3T" form of the APS form

---

[1] APS is also the acronym for the Advanced Photo System, a modern film and cartridge system. Although this actually has nothing to do with digital cameras, the designation has been hijacked into that field to denote a range of digital camera sensor sizes whose dimensions are something like the dimensions of one or another of the three frame sizes used in the Advanced Photo System. Thus there is ample opportunity for misunderstanding when "APS" is mentioned in connection with sensors.

of the CMOS sensor). Note at this point that there may be from perhaps 3 to 25 million such pixel cells in the entire sensor!

AN ILLUSTRATIVE IMPLEMENTATION

The figure

The figure shows the circuit schematic of an illustrative "basic" implementation of single pixel cell of this type of sensor, along with a portion of the common supporting circuitry. (Please note the caveat at the end of the article.)

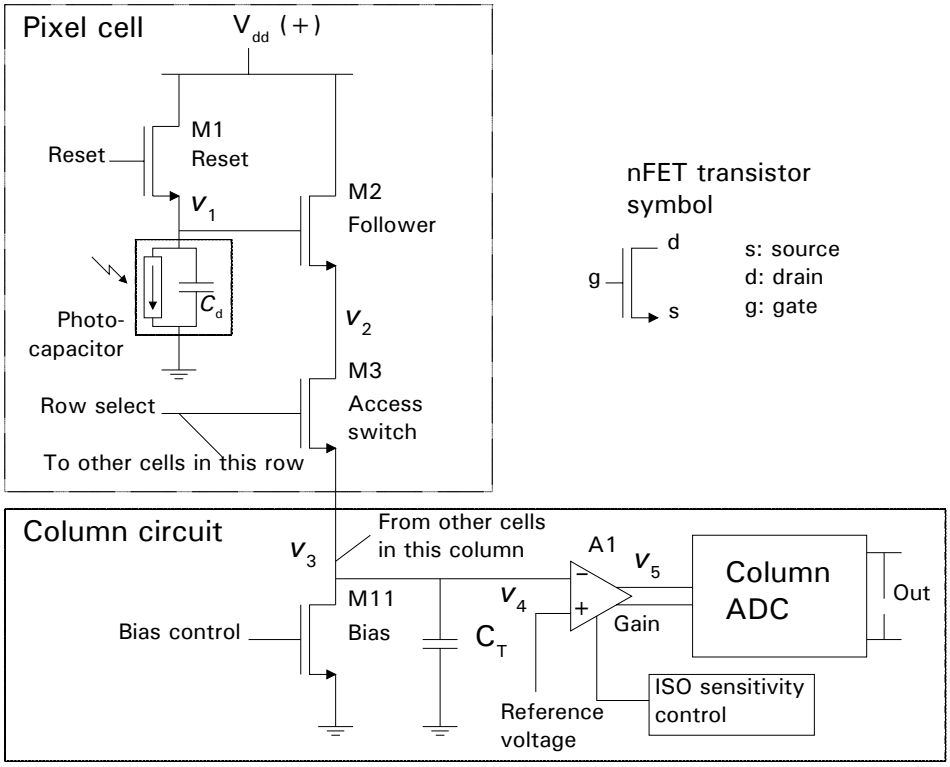

The full name of the transistor type used in a CMOS circuit is "metal-oxide-semiconductor field effect transistor" (MOSFET). MOSFET transistors come in two "genders", differing in the "doping" of their channels, and thus differing in the polarity of current flow between their two main electrodes, the *source* and the *drain*. The two flavors are the p-channel type (today usually called pFET) and the n-channel type(nFET). The APS cell we show here uses only the nFET type. We see the symbol for an nFET on the figure, with its three electrodes labeled.

Simplistically, the flow of current through the channel of the transistor (between source and drain) is controlled by the voltage on the gate.

Principle of photodetection

The actual photodetection is done with an MOS photocapacitor. It is common to call this a "photodiode", and indeed any two-element electronic device can be called a "diode", but this is in fact quite different from the "rectifying" diodes we usually mean when we refer to a semiconductor diode.[2] Thus my use of a distinctive name.

The actual internal working of this device is quite complex, and is quite different from the working of a conventional, non-semiconductor capacitor (perhaps the reason that is not customarily called a "capacitor"). We will, however, treat it with a simplified model that emphasizes its circuit behavior as seen from "outside" the device.

In the figure, we show the photocapacitor with a capacitor symbol accompanied by a "current regulator" symbol, which is meant to be evocative of the discharge of the capacitance by the impact of arriving photons, as will be described shortly.

Before the photographic exposure is commenced, the capacitance of this device is charged[3] to a carefully controlled "reset voltage". The photocapacitor is then "turned loose".

During the exposure interval, as each photon[4] of the arriving light strikes the photocapacitor, this results in the dissipation of a quantum of the charge on the photocapacitor. Accordingly, the voltage across the photocapacitor (generally proportional to the remaining charge) declines.[5] We use that decline in voltage to indicate the total number of photons that have become incident up to the time the exposure is ended (normally by an actual optical shutter).

This is in turn proportional to the illuminance-time product, *H*, on the sensor array at that pixel over the course of the exposure. This is the photometric quantity in which we are interested from a photographic

---

[2] Thanks to David Greenberg for suggesting the term "photocapacitor".

[3] We will subsequently refer to this as "resetting", not "charging", the device, out of respect for the sensibilities of semiconductor designers, who are alert to the fact that the actual internal mechanism of this process doesn't quote fit the simplistic notion of "charging".

[4] More precisely, "many of the photons". Not every one is fruitful, a matter that is reflected by stating the "quantum efficiency" of the detection process.

[5] We say "generally" because the mechanism of discharge results in the capacitance of the device changing as it discharges, disrupting the normal linear relationship between charge and voltage.

standpoint. (It is the same photometric quantity to which photographic negative film essentially responds.)

DETAILS OF OPERATION

Reset charging of the photocapacitor

The reset of the photocapacitor is done, just before exposure begins, by nFET transistor M1, which operates in the source follower mode.[6] In that mode, the source electrode is forced to a voltage close to that placed on the gate (the control electrode). The current needed to do so passes through the channel of the transistor from its drain to its source, drawn from the drain supply voltage, V_{dd}. The voltage applied to the gate to enable charging is carefully controlled, since it is vital that the initial voltage of the photocapacitor be consistent from cycle to cycle. Just before the exposure proper begins, transistor M1 is "turned off", thus "turning loose" the photocapacitor.

At the end of the exposure, the photon stream is blocked by the closure of the optical shutter and readout of the collection of photocapacitors begins. Readout consists of picking up the final voltage of each photocapacitor, comparing it to the assumed initial voltage, amplifying that difference, and passing it to an analog-to-digital converter (ADC) where a digital representation of the voltage is developed, to be passed on to the camera's image processing chain.

The ADCs

In the common implementation of the APS CMOS sensor system, there is not an ADC for each "column" of pixel cells (as the figure shows), but perhaps just one altogether, or perhaps one for each of a number of "sections" of the entire sensor. The ADC and its supporting amplifier is then "multiplexed" over all the columns it serves. This multiplexing is not shown on the figure, in the interest of clarity in showing the end-to-end operation. In effect, here we assume the presence of one ADC for reach column. Each ADC sequentially deals with each of the pixel cells in its column. We will see shortly how that happens.

In Appendix A, where we show a specific implementation of this sensor design principle, we will actually see the multiplexed ADC organization.

---

[6] This mode is analogous to the *emitter follower* mode of bipolar transistors or, for the old timers among us, the *cathode follower* mode of a vacuum tube.

The voltage follower

First note that the voltage of our photodetector, v_1, will, during readout, be "repeated" by transistor M2, operating in the source follower mode. That means that, as soon as there is any path from the transistor's source to "ground" (which will happen during readout, when the access switch "closes"), its source voltage, v_2, becomes close to its gate voltage, v_1.[7] The purpose of including this follower circuit in the chain is that the input impedance of this element (the gate impedance of M2) is almost infinite. Thus the connection to the gate of M2 does not constitute any drain on the remaining charge in the photodetector, which would of course corrupt the determination we are trying to make. The output voltage, v_2, can then be used as a proxy for v_1 to feed later parts of the circuit that may not have infinite input impedances.

Readout of the photocapacitor voltage

During the overall readout cycle for the entire sensor, when the pixel cell's row "comes up" for readout, access switch transistor M3 (in every cell of the row, actually) is "turned on hard" by raising its gate voltage substantially. Its source-to-drain path becomes a low impedance—a "closed switch".

This completes the path from the source of M3 to ground through "bias" transistor M4, which is kept turned on by a bias control voltage, causing it to operate in a "current regulating" mode as seen from its drain. (A certain voltage on the gate of the transistor causes a certain drain current, assuming that there is some voltage at the drain to propel this current.)

Thus M4 becomes a "constant-current load resistor"[8] for the source electrode of M2. (The term "bias", I'm afraid, doesn't describe that very well, but I have used it for consistency with the literature!) M2 now goes into the source follower mode earlier described.

With M3 "turned on hard", its source voltage, v_3, becomes close to its drain voltage, v_2. Thus, storage capacitor C_T[9] in the cell's assigned column circuit becomes charged to a voltage that is a proxy for the v_1

---

[7] There is in fact a significant "offset" voltage between the gate and source voltage, which in fact may vary with the amount of current flowing thorough the transistor.

[8] Often this would be called a "pull up" resistor, but in the source follower mode, this name doesn't seem quite apt to many people, so I will avoid it here.

[9] Actually, this is not necessarily a discrete capacitor, but rather is the intrinsic drain capacitance of M4 plus the intrinsic input capacitance of amplifier A1.

of the photodetector. This capacitor will hold that voltage while it is digitized in the steps to be next described.

Evaluation of the photodiode voltage

Amplifier A1, a differential type, essentially subtracts the cell voltage from a reference voltage, which is the voltage we would expect on C_T in the situation that the photocapacitor had not been discharged any by photon traffic. Thus the input to the amplifier (v_4) is the negative of the apparent decline in photodetector voltage, which is the manifestation of the discharge of the photodetectors capacitance by the incidence of the photons. The output of the amplifier, v_5, is then fed to the ADC to be digitized (perhaps in 12 bits or so).

In the sensors used in most digital cameras, these amplifiers can be set to any of several gains. This changes the scaling of the eventual digital representation of the photocapacitor voltage, in effect changing the "ISO sensitivity" of the sensor.

Shortly after, it will be the turn of the next pixel cell in this column to have its proxy photocapacitor voltage (v_3) transferred to C_T. That new value of v_3 may be higher or lower than the present v_3. If higher, current though the M2-M3 path of the new cell will charge up C_T to that new voltage. If lower, current through M4 will discharge C_T until its voltage reaches that which the M2-M3 chain would like to see (based on v_1), at which time current begins to flow through the M2-M3 chain to achieve equilibrium with the M_4 ("bias") current, at which time the new voltage on C_T has been established.

ABOUT "PIXELS"

The photocapacitor cell we discuss here is usually identified as a "pixel cell", and the moniker APS ("active pixel cell") reflects that. To avoid misunderstanding, note that, in most digital cameras today, there is indeed one sensor cell for each pixel of the image, but the recorded "color" of that pixel does not come directly from the associated cell. This is because the cell cannot actually determine the color of the arriving light.

Rather, the cells are divided into three groups, with each photodetector having in front of it a "color filter" whose spectral response is different for the three groups (the so-called *color filter array*, or CFA, structure). (We think of these as being "red", "green", and "blue" filters.) After the outputs of all the cells are passed to the camera, as we have described above, a complex algorithm "guesses" the color of the light at each pixel location by considering the outputs of a cluster of cells at and surrounding that location, taking into

account their different spectral responses. The details of this "demosaicing" process are beyond the scope of this article.

Because the individual photodetector cells do not actually report "pixel" values, we sometimes refer to them as "photosites" rather than "pixels". Thus, strictly speaking, the "active pixel sensors" are really "active photosite sensors" (still APS!).

CAVEATS

I have shown here an illustrative arrangement of the 3T APS form of the CMOS digital still camera sensor. The internal details, especially beyond M4, are illustrative and to some extent conjectural, and do not necessarily represent those of any specific implementation (not to mention the fact that I ignore the multiplexing of the ADC). The details through M4 are based on consistent information found in the literature. In any case, the interpretations are those of the author.

There have been developed a number of improvements in the CMOS APS sensor beyond the basic implementation we have examined, intended for such purposes as allowing the whole cycle to proceed more rapidly, reducing uncertainty in the output values, and so forth.

For example, there are reset charging schemes more sophisticated than the one we show with just transistor M1. These are intended to reduce variation on the actual initial charge voltage. Such variations affect the cell outputs for any given photon impact, and thus contribute a source of noise (called *reset noise*) to the overall image result.

CANON'S IMPLEMENTATION

Appendix A describes and discusses in detail the architecture and operation of the CMOS APS sensor design used by Canon, Inc. in the sensors for such cameras as the EOS 20D digital SLR, based on information kindly provided by a colleague in the CMOS integrated circuit industry.

ACKNOWLEDGEMENT

Many thanks to David Greenberg for his valuable tutelage in the operation of these CMOS devices, in particular of the photocapacitor itself, and for his many helpful technical and editorial suggestions on the manuscript.

Thanks to Mitch Alsup for contributing the description of the adaptation of this principle to the sensor used by Canon in its modern digital SLR cameras, as discussed in Appendix A, and for his many

helpful technical and editorial suggestions on the manuscript of that appendix.

Thanks also to Carla Kerr (in Cherokee, *Ka-la Tsu-la Gi-ga-ge*, "Carla Red Fox") for her careful copy editing of this difficult manuscript. This article is dedicated to her recent attainment of citizenship in the Cherokee Nation.

#

APPENDIX A

The CMOS APS sensor of Canon, Inc.

Mitch Alsup, a ULSI (ultra large scale integrated circuit) designer and fellow digital camera enthusiast, was kind enough to provide me with a detailed narrative description of the way the CMOS APS principle is implemented in the sensors made by Canon, Inc. for use in their modern digital SLR cameras, such as the EOS 20D. It uses 5 transistors for every two photodiodes, and is thus often spoken of as a "2.5T" APS implementation. Compared to a 3T implementation, this plan allows more of the cell area to be devoted to the all-important photocapacitors.

The figure on the next page shows the organization of this sensor. Not all details are available, and I have taken some literary license to complete the presentation.

Let me first mention a matter of orientation. As in the illustrative CMOS APS sensor described in the body of this paper, the readout of the array of photodiodes is organized here into units called, according to the custom for such matters, "columns". In this case, the columns run along the long dimension of the sensor array (that way there are fewer ADCs required). But of course in the camera that is the horizontal direction. Thus, a "column" lies horizontally, and a "row" vertically. I have followed this in my schematic drawing for various reasons of convenience and tidiness[10].

The Canon sensor differs from the illustrative one earlier described in several significant ways, notably:

1. In the illustrative sensor, each photodiode is supported by a source follower transistor, which occurs before the access switch. Here, there is one source follower shared by each adjacent two photodiodes. These occur after the access switch transistors (here called the "sense" transistors). The two adjacent photodiodes that share a single follower (with the other supporting circuitry for them) are considered to constitute a "cell".

2. In the illustrative sensor, after a photodiode is "read out", the implicit storage capacitor, C_T, is not reset. Rather, the

[10] One of which is for harmony with Alsup's notation, which I have generally adopted.

interaction between the bias transistor and the follower transistor for the "new" photodetector brings the C_T voltage, v_3, up or down to its new value. Here, after the voltage on the storage capacitor is digitized, the storage capacitor is quickly reset to its standard initial condition by a reset path involving two transistors.

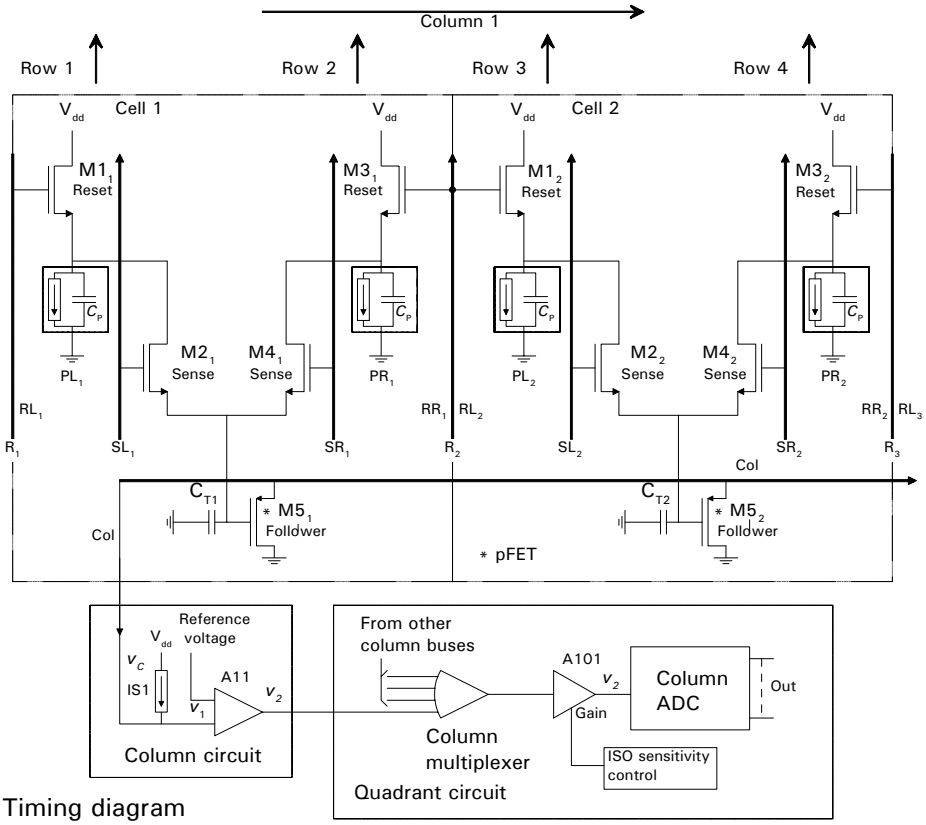

Timing diagram

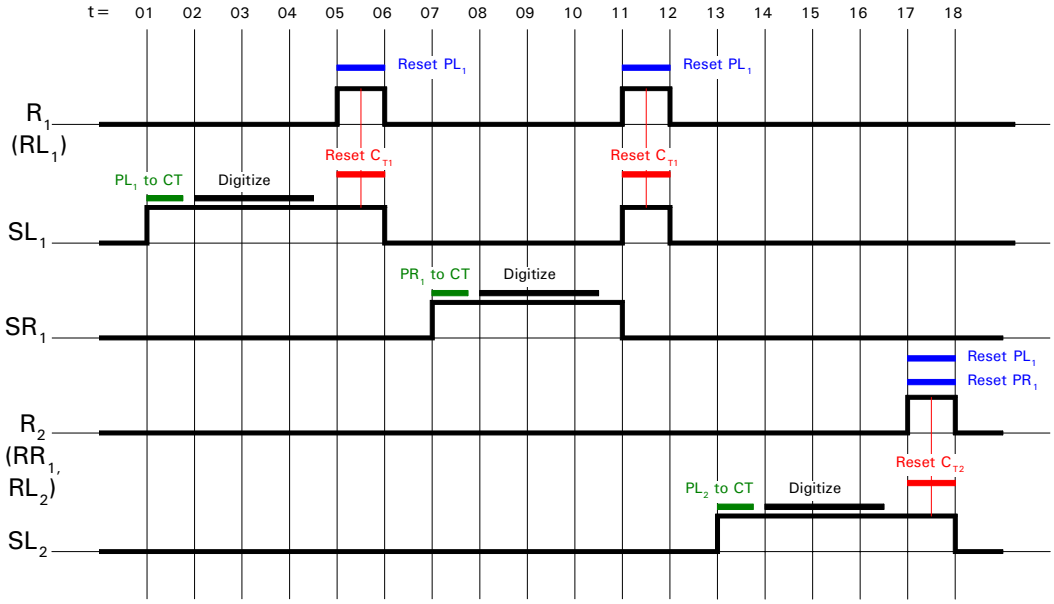

The entire sensor is divided into four quadrants, each of which is served by a separate ADC. The figure shows the first two cells of the first column of the first quadrant. In each quadrant, each column is served by a column circuit, primarily consisting of amplifier A11. The entire quadrant is served by a quadrant circuit, primarily consisting of the ADC, a controlled-gain amplifier (A101), and a multiplexer that allows information from every column in the quadrant to reach the amplifier and the ADC.

Note that for each cell, there are four "bus" leads: the "sense" leads, SL and SR (one pertaining to each of the photodiodes in the cell, PL and PR) and (if we look at the cell in isolation) the reset leads, RL and RR. But, in fact, there is really only one reset lead, R, at each cell boundary. Thus the R lead between cells 1 and 2 is both the right reset lead (RR) of cell 1 and the left reset lead (LR) of cell 2. Of course the first and last R leads in the column have only one role, as there is no adjacent cell.

A column lead, Col, runs through all cells in a given column. It is on this lead that the output of one photocapacitor after another is led to the common circuit for that column.

Note that, as before, capacitance C_T is not a discrete capacitor, but is the intrinsic gate capacitance of transistor M5 (thus its designation).

Sequence of operation

In following the action here, it will be helpful to look at the timing diagram at the bottom of the figure. The colored lines denote the actions that result from the voltage state transitions shown for the various control leads.

As with the illustrative sensor, just before the exposure commences, there is a mass reset operation (not shown on the timing diagram). During this, all the R leads are elevated to V_{dd}. Thus transistors M1 and M3 in each cell, operating as source followers, bring their respective photocapacitors to a voltage nearly V_{dd}—we think of this as resetting the photocapacitors. At the same time, the left sense lead (SL) will be also raised in every cell so as to make M2 conduct. Thus the elevated potential on photocapacitor PL is also passed to C_T, resetting it to a voltage close to V_{dd}.

This voltage on the gate of M5 will "cut it off" (and recall that this happens in every cell). We depend on this initial situation since, when we begin to read out the photodiodes in the various cells, we rely on only one M5 (in the cell currently being read out) exerting any influence on the column bus; any influence by the M5 in other cells would disrupt the process.

All the control leads are now returned to ground potential. All the photocapacitors and C_T capacitors are now "on their own".

Exposure then begins. After the exposure has been completed and the optical shutter closed, the readout cycle begins. (Now we are on the timing diagram!) The photodiodes are read out sequentially along the column, those in the same "row" of all columns being read out simultaneously into their respective column leads. Here's how that happens.

We will start with the readout of the first photocapacitor in the first cell of the illustrated column, PL_1. To access it for readout, at time t = 01 lead R_1 (in its role as RL_1) is raised to V_{dd}, turning on readout transistor $M2_1$. This allows the voltage of photodiode PL_1 to pass on to C_{T1}, and to the gate of the follower transistor, $M5_1$.

Note that, owing to a "division of charge" between the photodiode capacitance, C_P, and C_T, there will be a "capacitive voltage divider" phenomenon at work here. But since C_P is presumably much greater than C_T (owing to the much larger area of the photodiode compared to M5)—typically a capacitance ratio of about 100:1, we should end up with the voltage on the photodiode prior to readout being essentially reproduced at CT1 and the gate of $M5_1$.

We assume that this transfer has settled at some time before t = 02.

The follower transistor, M5, is a pFET, and is operated in a "downward" source follower mode. Its output (at its source electrode) is tied to the column bus, Col. The transistor's "pull up resistor" is actually the current source, IS1, in the common column circuit.[11] (It is in fact the pull up resistor for whichever transistor M5 in the column is active at the particular phase of the readout sequence.)

Note that exactly the same activity is taking place at the same time (or thereabouts) in cell 1 of all other columns of the overall sensor—the entire row.

At the column circuit, the column bus voltage, v_C, is compared to a reference voltage, which is the voltage we would expect to be passed to this point from a photocapacitor that had retained all of its original charge (had no exposure). The difference (v_1) reflects the degree of discharge of the photodiode, the thing in which we are really interested in all this action. It is amplified to v_2 by amplifier A1.

---

[11] Or so I conjecture. I actually do not as yet have any authentic information about the column circuit for this sensor design.

This voltage passes through the ADC multiplexer to the ADC, where it is converted to digital form. Following this, the voltage passed from PL_1 of all the other columns of the quadrant will successively be selected by the ADC multiplexer, amplified, and digitized by the ADC.

After a sufficient time for this to all happen (at time t = 05), reset lead R_1 (in its role as RL_1) is raised to V_{dd}. This causes transistor $M1_1$ to raise the voltage on PL_1 to nearly V_{dd}. Since transistor $M2_1$ is still conducting (as a result of the voltage on lead SL_1), that voltage is passed to storage capacitor Ct, resetting it to nearly V_{dd} so as to be ready to receive its next transfer, which will be from photocapacitor PR.

At time t = 06, leads R_1 and SL_1 are dropped to zero voltage, turning off transistors $M1_1$ and $M2_1$, dismantling this reset path.

Note that in the course of this, photocapacitor PL_1 has been reset to nearly V_{dd}, which doesn't really accomplish anything at this time. It is just a byproduct of using the M1-M2 chain to reset C_T.

Next, the voltage on the photodiode photocapacitor PR_1 will be read out. At t = 07, lead SR_1 is raised to V_{dd}, turning on transistor $M4_1$. This allows the voltage at photocapacitor PR_1 to pass to CT_1. As before, $M5_1$ passes this voltage on to the column bus. In the column circuit, this voltage is amplified as before and passed through the ADC multiplexer to amplifier A101 and then to the ADC to be digitized.

After a brief pause for this to all happen, then (at time t = 11), reset lead R_1 (in its role as RL_1) is raised again to V_{dd}. In the same way we saw just previously, this resets CT_1 to nearly V_{dd} through $M1_1$ and $M2_1$.

Why don't we, for the sake of symmetry, use a path through $M3_1$ and $M4_1$ to reset C_{T1} this time? Because reset lead R_1 is not only RR_1 but also RL_2. Thus reset transistor $M1_2$ (in cell 2) would also be turned on, resetting photocapacitor PL_2—and it hasn't been read out yet.

Now, we perform the same sequence for cell 2, first for photodetector PL_2. At time t = 13, lead R_2 (in its role as RL_2) is raised to voltage V_{dd}, turning on sense transistor $M2_2$. This allows the voltage at photodiode PL_2 to pass on to CT_2, at the gate of the follower transistor, $M5_2$. And so forth.

Now just for awe's sake, recall that in the specific use of this design originally mentioned (the Canon EOS 20D digital SLR camera), there are over 2300 "columns", over 3500 rows, over 4 million cells (and

that many M5s), and of course over 8 million photodiodes and M1s, M2s, M3s, and M4s.[12] In addition, there are over 4600 column circuits (each column has a piece in each of two quadrants of the sensor, each piece with its own column circuit), 4 ADC multiplexers (with over 1150 inputs each), 4 quadrant amplifiers, and 4 ADCs. What an accomplishment!

ACKNOWLEDGEMENT

Great thanks to Mitch Alsup not only for bringing to my attention the Canon sensor architecture but also for his many helpful technical and editorial suggestions on the manuscript.

#

[12] The pixel image developed from the sensor has dimensions 3504 px × 2336 px. However, the actual array of photodetectors is slightly larger than this for various reasons. For one thing, under the color filter array (CFA) approach briefly mentioned in the body of the article, to determine the color of each pixel the demosaicing algorithm considers the outputs of a cluster of photodiodes centered on the pixel of interest. Thus, with respect to the pixels at the boundary of the sensor, there must be photodetectors "in the margin" to complete those clusters.