Analog Integrated Circuits and Signal Processing 1, 93–106 (1991)

Scanners for Visualizing Activity of Analog VLSI Circuitry

CARVER A. MEAD AND TOBIAS DELBRÜCK
Computation and Neural Systems Program, California Institute of Technology, Pasadena, CA 91125

Received May, 1991; Revised June 27, 1991.

Abstract. This paper tutorially describes mixed digital-analog serial multiplexers (scanners) that we use to visualize the activity of one- and two-dimensional arrays of analog VLSI elements. These scanners range from simple one-dimensional devices designed to scan a one-dimensional array onto an oscilloscope, to complete video scanners with integrated sync and blank computation and on-chip video amplifiers. We discuss practical details of design and performance, and we give a source for example scanner layout.

1. Using Scanners in Analog VLSI Design

We use scanners as a diagnostic tool, to observe the behavior of large analog VLSI chips, when the chips have far more nodes than the pins available to us, or when the chip has a topography that maps well onto a display device.

This paper is a tutorial description of the distilled knowledge gained from over 100 chip designs using scanners. We describe how our scanners work and how we use them to multiplex outputs from one- and two-dimensional arrays onto output devices such as oscilloscopes and monitors. In section 2, we discuss one-dimensional scanners and their component parts as prototypical of two-dimensional scanners. In section 3, we discuss two-dimensional scanners and generation of video control signals. In section 4, we discuss the analog parts of the scanners: output representations, current-sense amplifiers, and video drivers. In section 5, we discuss performance of the digital and analog parts of the scanners. Finally, in section 6, we discuss practical design details.

2. A One-Dimensional Scanner

Figure 1 shows a one-dimensional scanner. This scanner multiplexes the output from a one-dimensional array of pixels serially onto a single output line. The pixels generate analog currents that are to be scanned out and displayed. The pixel outputs are multiplexed by MOS switches onto the output line or onto the reference line. At a particular point in the scan, one pixel sends its output into the output line, and all the other pixels feed into the reference line. The reason for this dual-output scheme is that, generally, when we want to scan quickly, we must hold the wires at a fixed potential (virtual ground) while we sense an output current. The dual-output scheme allows us to hold all the wires at the same virtual ground. If we scanned out a voltage, each pixel would need to charge the output wire to the output voltage, and this circuit would run much more slowly than does the current-sense arrangement. The current-sense amplifier in figure 1 shows one method for sensing the current coming out of the scanned pixel.

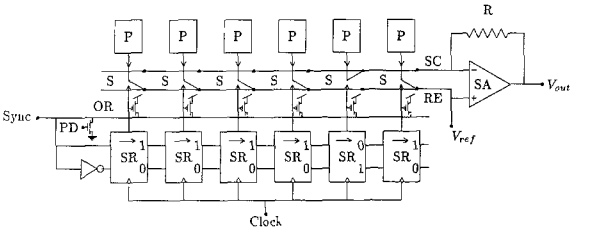

Fig. 1. Schematic of a one-dimensional scanner. Pixels (P) send their output current through switches (S) to either a scan (SC) or a reference (RE) wire. The pixel current is routed through the SC wire and is sensed by the current-sense amplifier (SA). The current-sense amplifier holds the SC wire at approximately V_{ref} and outputs a voltage proportional to the pixel output current. The switches are controlled by the shift-registers (SR). The shift-registers shift in the direction of the arrow. A new bit is loaded into the start of the shift-register at the end of a scan by means of the wired NAND, which computes the OR of low bits in the shift-register (OR). The OR output also acts as a sync input to the oscilloscope. The pulldown transistor on the OR line is biased with the bias PD. The output from the sense-amplifier goes to an oscilloscope.

The blocks at the bottom of figure 1 represent a digital shift-register than controls the scanning process. With each clock cycle, whatever bit is in a given stage is shifted to the next stage. Whatever bit is in the last

stage in the line gets shifted out and disappears forever. Typically, there will be a single register that holds a low bit, and all the other registers will hold a high bit. That one register with the low bit determines the one pixel that will be connected to the output amplifier in that clock cycle. In figure 1, the second pixel from the right is being scanned out. All other pixels are connected to a reference line. The line coming into the bottom of each shift-register stage is the global single-phase clock input. The wired NAND ciruit logically ORs together all the shift-register low values so that a new low bit is generated when only high bits are left in the register. This arrangement is self-initializing and requires no off-chip control.

In the following sections, we discuss the component parts of this one-dimensional scanner.

2.1. Shift-Register Operation

In these scanners, the shift-register is a static digital device. The clock input required is a simple single-phase signal alternating between high and low. Other schemes use a two-phase, nonoverlapping clocking scheme, where there are two clock signals, both alternating between high and low, such that the high periods of the two signals are completely nonoverlapping. The two schemes are shown in figure 2. Since the shift-registers are static devices, they hold their values without refreshing. Hence, we may view a particular pixel continuously by stopping the scan at that pixel. The shift-register's outputs are fully restored to V_{dd} or ground when the clock is high.

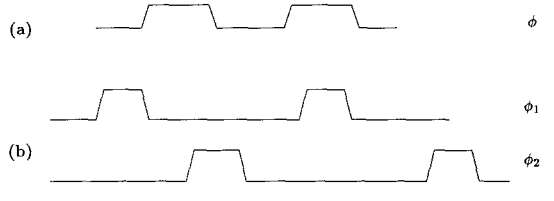

Fig. 2. Two clocking schemes. (a) Single-phase. (b) Two-phase, nonoverlapping. Clock signals go from ground to V_{dd}.

The single-phase CSRL shift-register circuitry is a novel design, and was derived by Massimo Sivilotti [6] from an earlier two-phase version [5]. The circuitry for one stage of a shift-register is shown in figure 3. We see that a stage consists of two parts of cross-coupled inverters, along with n-type pass transistors feeding into the first pair, and p-type pass transistors connecting the first and second pair. (We assume here that we are using an *n-well* implementation; native transistors,

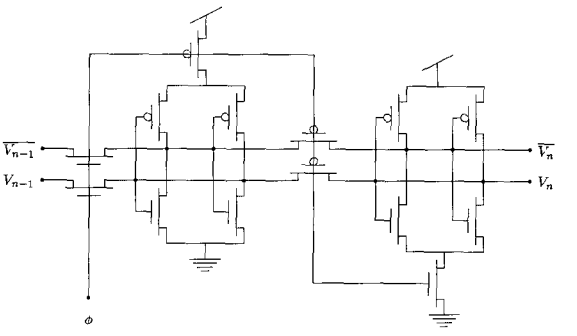

Fig. 3. A single-phase static shift-register stage. The bit in the stage is represented by the complementary pair V_n and $\bar{V}_n$; these signals are fully restored to V_{dd} or ground when ϕ is high.

without a bubble on the gate, are n-channel; and the transistors in the well, with a bubble on the gate, are p-channel. Logic-high is +5 V and logic-low is ground. For a *p-well* implementation, using the same layout geometry, we reverse all potentials with respect to the substrate, and we exchange transistors types.)

The inverters are conventional CMOS inverters except that they have a transistor of one type or the other in series with their power supply. For the left pair of inverters, the p-type power-supply transistor is between V_{dd} and the inverters; thus the inverters will be powered only when the clock input ϕ is low. For the right set of inverers, the situation is reversed: The inverters will be powered when ϕ is high. The bit in a stage is represented by the values V_n and $\bar{V}_n$; the two signals are fully restored and complementary when the clock is high.

When the cross-coupled inverters are powered, they have two stable states corresponding to V_n high or low and $\bar{V}_n$ complementary to V_n. When they are not powered, we can load bits onto their inputs. That loading is the function of the pass transistors connecting this stage to the previous stage and connecting the first pair of cross-coupled inverters to the second.

It is easy to see how this arrangement works, as long as you already believe that it works. The key requirement to making these single-phase shift-registers function correctly is that the pass transistors must be sufficiently weaker than the power transistors that the driven stage cannot change the state of the driving stage. The correct sizing restriction will ensure unidirectionality of information transfer. Analytically, and also empirically, a sizing restriction that works over the current range of MOSIS parameters is that the pass transistors feeding into the cross-coupled inverters must have a saturation current at most one-fourth that of the power transistors or those in the cross-coupled inverters [6].

If this sizing restriction is obeyed, then the shift-register will function correctly with any clock rise- or fall-time. In fact, a sinusoidal oscillator may be used as the clock input.

This fully static, single-phase, shift-register stage is a novel design. Other static digital shift-registers use two-phase, nonoverlapping clocks that are inconvenient and unnecessarily complicated to generate, and are particularly prone to clock-skew problems. Other single-phase designs are either not fully static, or have a high transistor count [3].

2.2. *Generating a New Bit*

The scanners discussed here have at most one low bit in the shift-register. Whenever that low bit falls out the end of the shift-register, a new low bit is generated and is loaded in at the beginning. The rest of the time, high bits are loaded into the shift-register. The register is thus self-initializing. This trick is managed by a wire that computes the logical NAND of all the outputs of the shift-register stages. Conceptually, we can think of this circuit as an OR gate for low-going inputs. If there is a low bit in any stage, then the OR line is pulled high, and only highs are loaded into the shift-register. If there is no low bit, then the OR line gets pulled low by a pull-down transistor, generating a new low bit. In addition to generating a new bit, the NAND output acts as a synchronizing signal (sync) to trigger the oscilloscope sweep.

2.3. *Multiplexing the Pixel Output*

We use the low bit marching along in the shift-register to switch the output from the selected pixel (or column, as we shall see in section 3) onto the output line. The outputs from all the other pixels are connected to the reference line. The switch, shown in figure 4, is a conventional analog multiplexer, consisting of four transistors, two of each type. We use this complementary arrangement since a particular type of transistor will not pass signals well when the signals are far from that type's bulk potential, that is, native transistors will not pass signals near V_{dd}, and transistors in the well will not pass signals near ground.

When the switching transistors are turned off, whatever charge is under the gates of the transistors gets pushed out from the channel and adds to the current we are observing. This charge injection sometimes causes unwanted transients in the pixel output, particularly at low current levels. The use of complementary

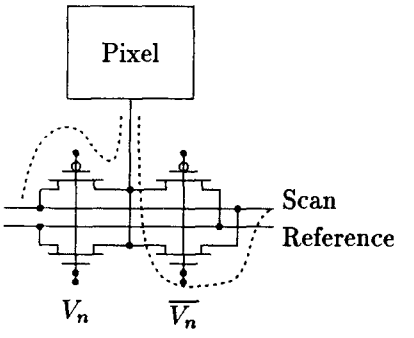

Fig. 4. Pass-transistors used to switch pixel output onto output line or onto reference line. The dotted lines show the flow of current when V_N is low.

pass transistors somewhat ameliorates this charge-injection problem, since one type injects electrons and the other type injects holes [7]. More sophisticated techniques, not described here, may further alleviate this problem [9].

2.4. *Sensing the Output*

In general, it is difficult to scan rapidly if the pixel generates a voltage that must charge the output line. For a one-dimensional scanner that will be viewed on an oscilloscope, a voltage scan is feasible as long as the scan rate is no higher than a few kHz [8]. A two-dimensional video scanner, however, would require pixels that could drive a long metal wire at several MHz. For this reason, we have chosen schemes that sense a current generated by the pixels. By using a single, fast, feedback amplifier, we can hold the output line at a fixed voltage, *virtual ground*, thus allowing much higher scan rates.

Figure 1 shows the prototypical current-sense amplifier, arranged in a negative feedback configuration. Any deviation from V_{ref} at the scanout line is multiplied by the gain of the amplifier and fed back to the scanout line through a feedback element. Thus the feedback amplifier senses the error signal at the scanout line and uses it to correct the voltage at the scanout line. Under this condition, the current through the feedback element is equal to the pixel output current. If the current were not equal, then the voltage on the scan line would continue to change. Since, in the circuit shown, the output voltage is fed back through a linear resistor, the output voltage is proportional to the pixel output current. The proportionally constant is set by the feedback resistance R. In section 4, we discuss this topic at greater length.

3. Two-Dimensional Scanning

In this section, we discuss all digital parts of two-dimensional scanning, including selection of rows and generation of video control signals. The horizontal portion of a two-dimensional scanner is identical to a one-dimensional scanner. The added vertical portion performs row selection. This scheme is shown schematically in figure 5. In the simplest case, the vertical scan will connect a pass-transistor switch inside each pixel in a row of the array. Each pixel in that row will send its current to the horizontal scanner, which will select one of the pixels for output. Section 4 discusses row selection in more detail.

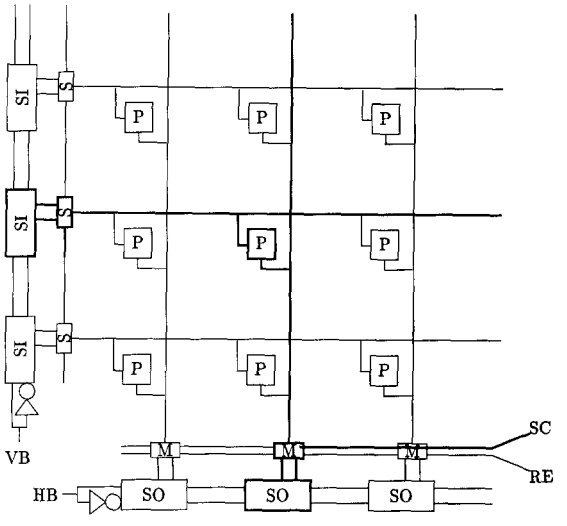

Fig. 5. Two-dimensional scanning. Row selection is done by vertical shift-registers (SI). Column selection is done by horizontal shift-registers (SO). Output from row of pixels (P) selected (S) by output from vertical shift-registers is multiplexed, by switches (M) that are controlled by the horizontal shift-registers, onto either scanout wire (SC) or onto reference wire (RE), as for one-dimensional scanner. New bits are loaded into shift-registers as needed (HB and VB). Bold lines show the center pixel in the array selected for scanout.

3.1. Making a Raster-Scan Picture

Because of the proliferation of personal computers, many varieties of multiscanning monitors are available at modest cost. These monitors are designed to display output from a variety of different display-adaptor cards; two familiar examples are the EGA and VGA graphics standards for IBM and clone machines. In contrast to the requirements for National Television Standards Committee (NTSC) video (the standard format for broadcast television in the United States), the timing requirements for multiscanning monitors are not nearly as stringent, and there is no need to mix the video brightness, color, and timing information into a single signal. In fact, making a video signal for a multiscanning monitor is relatively simple, so we shall confine our discussion here to the production of such a signal. The resulting signal will not generally be understood by a regular television monitor or video-tape recorder, but is much easier to generate.

In the simple mode (the only one we discuss here), there are five inputs to the monitor: three analog inputs for the three colors—red, green, and blue (RGB)—and two digital inputs for horizontal and vertical synchronization (sync). The RGB signals will be discussed in section 4.

Figure 6 shows how the electron beam moves across the monitor face. The beam starts at the upperleft corner of the screen and scans across the screen. After the end of each line, the horizontal-sync signal starts the horizontal retrace. After the last line, the vertical-sync signal starts the vertical retrace. The sync signals tell the monitor when to start the horizontal and vertical retrace.

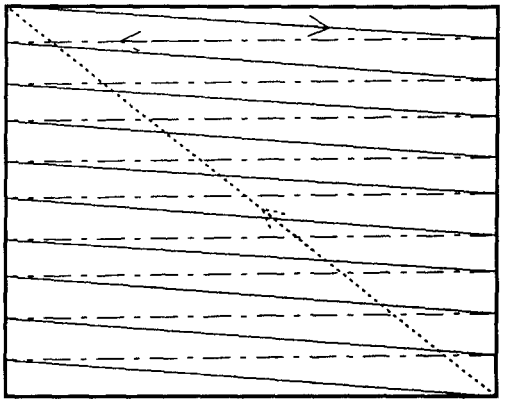

Fig. 6. Electron beam movement on monitor display. Horizontal sync occurs after each line; vertical sync occurs at end of frame.

We generate two additional digital signals that are used to blank the video signal during the horizontal and vertical retrace periods. These are not separate inputs to the monitor; they are used only in our video circuitry. (We could generate a single blank signal, and thus save one pin, but generating separate horizontal and vertical blank is slightly simpler and is a help in debugging.)

The detailed timing requirements for these monitors are given in the owner manuals; what follows is a slightly simplified summary for a particular monitor, the NEC Multisync model II. A video frame must consist of between 200 and 700 lines. The vertical sync must

run at 60 ± 15 Hz. The blank period must be some fraction of the display period, and the sync pulses must occur somewhere near the beginning of the blank period. For a horizontal scan line, the blank period must be 33% ± 7% of the display period, and for a vertical scan line, the blank period must be 20% ± 10% of the vertical display period. Adherence to these values will almost always produce an acceptable picture, although adjustment of the scan rate will sometimes help produce a more pleasing aspect ratio or will fill the monitor screen more fully.

3.2. Generation of Sync and Blank Signals

To generate the horizontal and vertical sync and blank signals, we take advantage of the timing already extant in the horizontal and vertical shift-registers by extending the shift-registers that do the row and column selection (the display interval), to include the blank intervals as well. This scheme is shown in figure 7. The horizontal and vertical scanners act as before, with the addition of extra shift-register stages that encode the sync and blank signals. The sync signal is generated

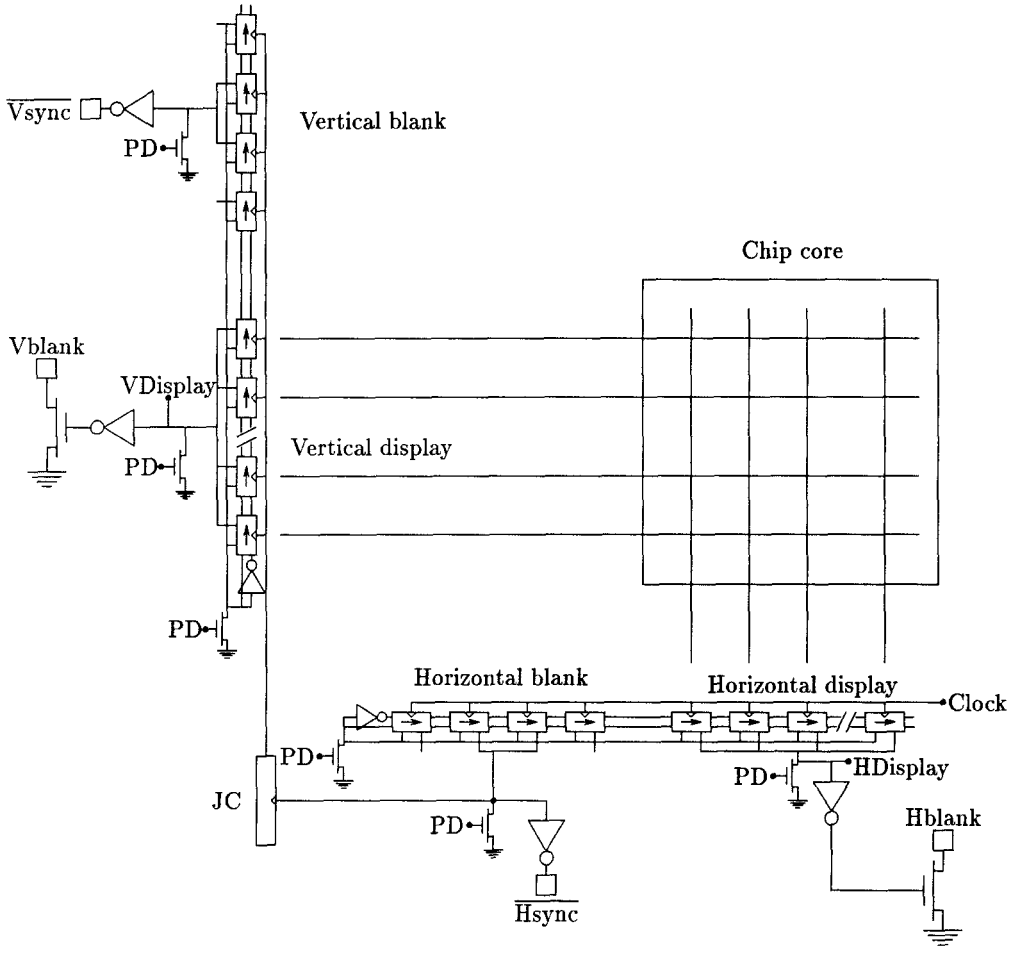

Fig. 7. Video scanner with integrated sync and blank signals. Shift-register stages are shown as boxes with arrows showing direction of shift. Clock inputs to shift-registers are shown as carets coming into top of horizontal shift-registers and into right of vertical shift-registers. Wired NAND connections are shown as lines extending outward from shift-register. Lines extending horizontally into the chip core are row select lines; lines extending vertically into chip core are column output wires. JC is the Johnson row counter. Th pulldown transistors (PD) act to pull down the wired NAND lines; these transistors are all biased with a common pulldown bias. Sync and blank output pins are shown as small boxes. The actual scanner is larger; only a few sections are shown for clarity. Actual sizes, for a 50×50 pixel chip, are as follows: Horizontal blank section: 11 stages. Horizontal sync stages: stages, 2, 3, and 4 of horizontal blank section. Vertical blank section: 10 stages. Vertical sync stages: stages 2 and 3 of vertical blank section. Transistor W.L ratios: Sync output inverters: 116:2. Blanking pulldown transistors: 232:2. These large transistors are capable of driving off-chip loads.

by a wired NAND line that is pulled high when the bit is in the shift-register stages representing the sync signal. Similarly, the blank signal is generated as the NOT of a wired NAND encoding the display period. Thus, there are two additional wired NAND lines for each dimension of scanning. Because the display interval and the sync interval do not overlap, only one additional row of wired NAND transistors is required.

3.3. Counting Rows

Generally, a chip with a complex pixel will have far less than 200 rows of pixels, so we will choose to scan each row several times to fill the rows of the video image with the required number of lines of video. We use a Johnson counter (figure 8) to count the number of monitor scan lines that display the same row of pixels. The clock for the Johnson counter is the horizontal sync signal; the buffered output of the Johnson counter is the clock for the vertical scanner.

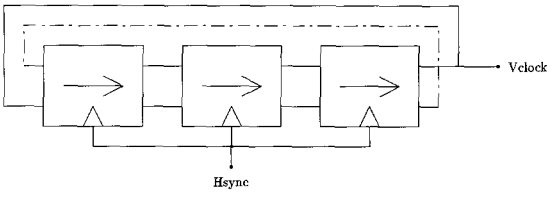

Fig. 8. A three-stage Johnson counter. The last stage of the counter feeds back into the first stage with an inversion. Hence, one possible sequence of states is as follows: {0, 0, 0}, {1, 0, 0}, {1, 1, 0}, {1, ,1, 1}, {0. 1, 1}, {0, 0, 1}. (There is also a parasitic state consisting of the sequence {1, 0, 1}, {0, 1, 0} that is not encountered in practice.) The vertical clock is the buffered output from the last stage; the clock input comes from the horizontal sync signal. An N-stage Johnson counter counts $2N$ lines. Counters with $N \leq 2$ do not have parasitic states.

3.4. Crystal Oscillator Circuit for Generating Clock

For video scanners we save ourselves the inconvenience of generating a clock signal off chip by using the crystal oscillator arrangment shown in figure 9. Using this circuit, we generate the main clock signal with an off-chip crystal and a few resistors and capacitors.

3.5. Scanning a Hexagonal Array

Some arrays are best laid out hexagonally. A hexagonal arrangement imposes an additional constraint on scanning, since alternate rows of the array must be delayed

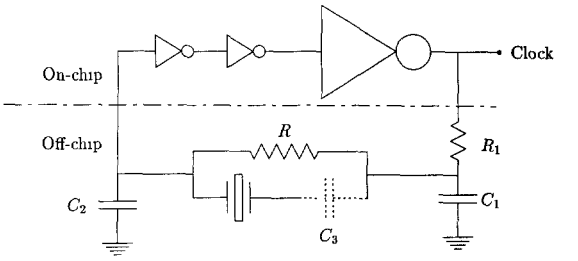

Fig. 9. Crystal oscillator circuit. The crystal is driven by and drives the three on-chip inverters. The use of three inverters provides enough gain that the oscillator does not balance. The last inverter is larger than the others to provide the drive for the off-chip components and for a large on-chip fanout. The off-chip resistor R biases the inverters into their high-gain region. At the resonant frequency, the crystal is a short circuit. The inverters provide a 180° phase shift, plus some internal delay. The R_1C_1 combination provides additional phase shift, so that the crystal can make up the required 360° phase shift. Capacitor C_2 stabilizes the oscillation. Values of components: $R \approx 10\ M\Omega$, $R_1 \approx 500\ \Omega$, $\frac{1}{2}\pi R_1 C_1 = f$, where f is the crystal frequency, and $C_2 = C_1/4$.

For frequencies above 4 MHz, R_1, C_1, and C_2 may be unnecessary, but $C_3 = 1pF$ may need to be substituted to prevent overdriving the crysal, which results in oscillation at an overtone frequency. If C_3 is not used, the crystal should be driven by C_1 directly, or from the clock output if R_1 and C_1 are not used. Transistors in first two inverters have W:L ratios of 14:2; last inverter has a W:L ratio of 116:2.

by one half of a clock phase to map the chip topography accurately onto the monitor screen. One scheme that we have used is shown in figure 10. In this scheme, an additional OR line on the vertical scanner encodes whether the row being scanned is even or odd. This signal is used to select either the first half-phase or the second half-phase from the horizontal shift-register stage as the source of the horizontal multiplexer pass-gate signals.

4. Output Representation, Current Sensing, and Video Drivers

In this section, we discuss the analog parts of scanners, including typical transformations between internal pixel representation and output current, the virtual ground scheme, current-sensing amplifiers, and video amplifiers for monitoring interface.

4.1. Pixel Output Representation and Row Selection

Figure 11 shows pixel output ciruits. In each case, we suppose that the internal pixel representation is a voltage

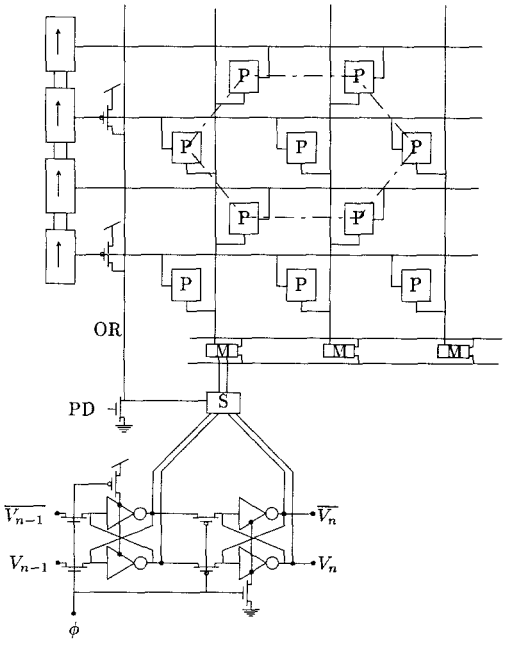

Fig. 10. Scanning of a hexagonal array. In the chip layout, alternate rows of the array are shifted by half of a pixel-width. The shift-register is as before. Inverters shown have a power-supply transistor in series with either V_{dd} or ground. Signals from each horizontal shift-register stage pass through MOS switch (S), where either the first half-phase or the second half-phase output from the shift-register is taken to be the switching signal for the output multiplexer switches (M). An odd-row wired NAND (OR) encodes whether an odd or even row is being scanned, and controls which half-phase is taken to be the multiplexer control signal. For the array shown in the figure, the first half-phase is taken for the odd rows of the array, and the second half-phase is taken for the even rows. Only one horizontal shift-register stage is shown in the figure.

that is converted into an output current that we sense using a current-sense amplifier. In figure 11(a) and (b), the internal voltage representation is converted into a current using a single transistor. In (c), the internal representation is a differential voltage ΔV that is converted into a current with a transconductance amplifier.

If the scanout requires that a bias voltage be given to each row of the array (as in figure 11(c), we use a pass-gate arrangement similar to the switches used in the horizonal scanner. The horizontal driver circuit we use for supplying the bias voltage to the row of pixels is shown in figure 12. The time available during horizontal blanking is sufficient to allow the voltage to settle to the correct value.

4.2. Current-Sense Amplifiers

To sense the current output by the pixels, we use a current-sense amplifier. If the amplifier is off-chip, we use the arrangement shown in figure 13(a). This amplifier senses a bidirectional current, and the sensitivity is set by the resistance R. The feedback will keep the negative input to the opamp very near V_{ref}; under this condition the current through the feedback resistor must equal the input current, and hence the output voltage will be $V_{\text{ref}} + IR$.

For the off-chip current-sense amplifiers, we successfully use a TL074 opamp; the only disadvantages are the requirement for a +12V/−12V power supply and the limitation to the pixel rate to less than 1 MHz. Other opamps may be faster, but are more difficult to stabilize. We often use the extra amplifiers (the TL074 comes in a quad configuration), in follower configuration, to supply the reference voltage V_{ref} or other reference voltages used in the core of the chip.

On-chip, we use unidirectional sense amplifiers like those shown in figure 13(b) and (c). A simple analysis, based on the subthreshold exponential relationship between gate voltage and drain current, shows that the output voltage is logarithmic in the current I:

$$V_{\text{out}} = \frac{V_{\text{ref}} + kT/q \ln (I/I_0)}{\kappa}$$

I_0 is the leakage current, and $\kappa \approx 0.7$ is the back-gate coefficient. Above threshold, the output voltage is related to the square root of the input current:

$$V_{\text{ref}} = \frac{V_{\text{ref}} + V_{\text{T}} + \sqrt{I/I_0'}}{\kappa}$$

V_{T} is the threshold voltage, and I_0' is a constant with units I/V^2.

It is intuitively clear that, the tighter the feedback amplifier clamps the sense line, the more speedup the feedback arrangement will provide. We quantify the speedup by defining a natural time scale that is the open-loop time constant of the pixel output line.

$$\tau_{\text{in}} = \frac{C_{\text{in}}}{G_{\text{fb}}},$$

where C_{in} is the capacitance of the line and G_{fb} is the source conductance of the feedback transistor. This time

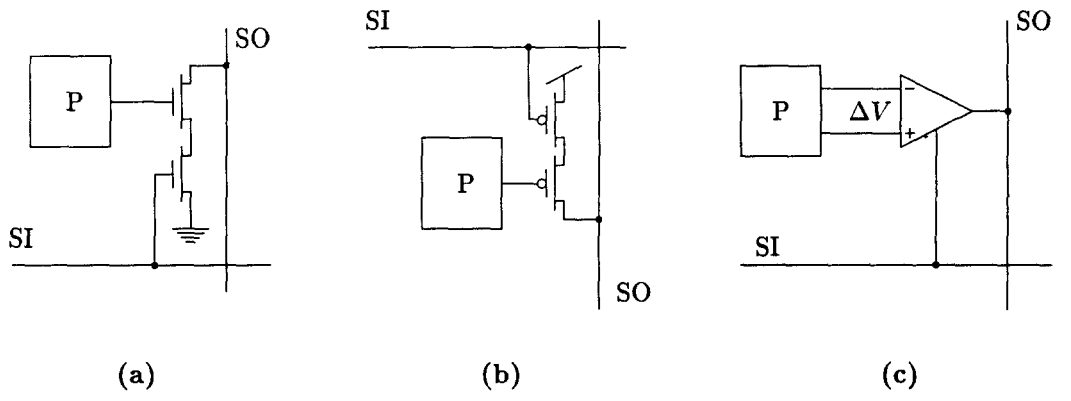

Fig. 11. Scanout-selection examples. The scanin (SI) and scanout (SO) lines connect to the vertical and horizontal scanners, respectively. Pixels (P) supply either a single-ended voltage ((a) and (b)) or a differential voltage (c). SI lines select one row of pixels. In (a), SI goes high to select the pixel. In (b), SI goes low to select the pixel. In (c), SI biases the transconductance amplifier to select the pixel.

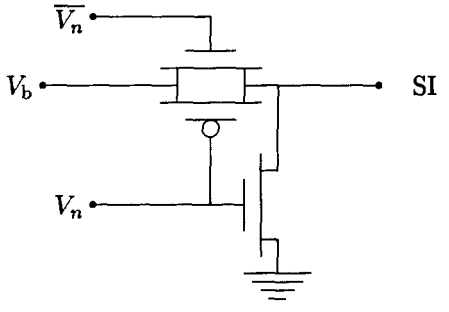

Fig. 12. Vertical scanner output circuit used to generate a voltage to bias a row of transconductance amplifiers. V_b is the desired bias voltage, SI is the scanin line, V_n and $\bar{V}_n$ are logic signals coming from the shift-register. V_n goes low to select the row.

constant represents the speed of a current-sensing arrangement in which we measure the voltage at the source of the feedback transistor, holding the gate and drain of the feedback transistor at a constant voltage. A simple linear analysis, similar to those appearing in the literature (for example, Bult and Geelen [1]), shows that the speedup provided by using a fast feedback amplifier with voltage gain of A is just κA; in other words, the resulting first-order time constant of the output is

$$\tau_{\text{out}} = \frac{\tau_{\text{in}}}{\kappa A}$$

where $R \approx 0.7$ is the back-gate coefficient. However, if the feedback amplifier is too slow, then the output signal will ring. To prevent ringing, we must ensure that the feedback amplifier is at least $4\kappa A^2$ times faster than the input node. The implication of this result is that, although a high-gain feedback amplifier is desirable for maximum speedup, to prevent ringing, we may have to settle for a gain of 100 or less. In our amplifiers, we use minimum-length transistors, and we run the bias of the amplifiers high enough that the gain is actually reduced by above-threshold effects.

The transconductance amplifiers shown in figure 13(b) and (c), are simple single-stage transconductance amplifiers [4]. Since the feedback amplifier must run much faster than the input node to prevent output ringing, we build our feedback amplifiers using ring transistors, where the inside of the ring is the drain of the transistor. The ring-like arrangement allows very wide transistors to be placed in a small area, with minimal side-wall capacitance. The resulting amplifier has relatively low gain but high transconductance.

We must take care with the logarithmic sense amplifiers shown in figure 13 to ensure that the amplifiers are operating in their proper voltage ranges. A simple transconductance amplifier with a native-type bias transistor and differential pair will only operate correctly when the output voltage is above $V_{\text{min}} = \min(V_+, V_-) - V_b$. V_+ and V_- are the amplifier input voltages, and V_b is the bias voltage. This condition is satisfied for the circuit in figure 13(b). However, the amplifier in figure 13(c) will output a voltage closer to ground that V_{ref}. Hence, the amplifier must be able to operate when the output voltage is substantially below both of the inputs. In this case, a well-type differential pair must be used. In addition, if the pixel requires that V_{ref} be held near ground, the back-gate effect on the exponential feedback transistor would require that the outut voltage be below ground. Hence, we usually put this feedback transistor in its own well and tie the source of the transistor to the well. For the logarithmic sense amplifier in figure 13(b), this arrangement is not possible because the feedback transistor is in the substrate.

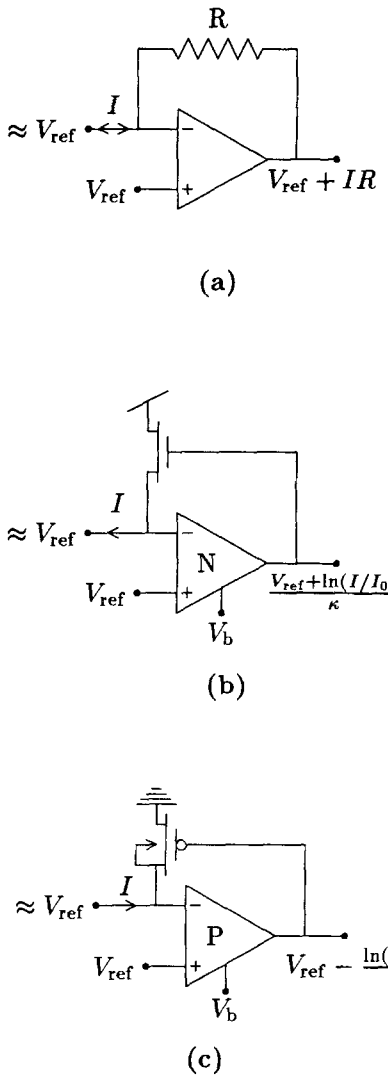

Fig. 13. Three current-sense amplifiers (a) a bidirectional off-chip linear arrangement; (b) and (c) unidirectional on-chip logarithmic amplifiers. For (b) and (c), output voltages are shown for the subthreshold case. Both (b) and (c) are on-chip simple transconductance amplifiers [4] biased with bias V_b. N and P refer to the type of the amplifier: N means that the differential pair should consist of native devices, P means that the differential pair should be constructed from transistors in the well. In (c), we show the feedback transistor in its own well, with the well tied to the source of the transistor. This modification will prevent the back-gate effect from requiring an output voltage below ground for a large current I and a V_{ref} near ground.

If we use these current-sense amplifiers in a video driver, it is essential that we introduce the minimum amount of stray capacitance. Hence, we buffer the output of the current-sense amplifiers, using a voltage follower, before driving a large capacitive load. For the same reason, schemes that place the feedback element off-chip have difficulty at video rates, due to the pad, package, and circuit-board capacitance.

In section 5 we discuss performance of the curent-sense amplifiers.

4.3. Video Amplifiers

The RGB lines are AC coupled inside the monitor. Hence, their DC levels do not matter. The total brightness range is generally about one volt from black level to full saturation. The brightness of the image is given by the contrast in the video signal between the video level and the blank periods surrounding the sync signals. It is up to the user to provide the video blanking.

The transition from the internal video signal, as computed by the on-chip sense amplifiers, to the off-chip video signal that drives the monitor is often the most painful part of getting a chip running. The monitor RGB inputs are terminated with a standardized 75-Ω load—a load that requires very large transistors to drive directly. We use smaller output transistors and off-chip amplifier arrangements like those shown in figure 14 to do the impedance matching, and to provide maximum flexibility in interfacing to a particular monitor. These video drivers have been evolved to use the minimum number of off-chip components, and function with a minimum amount of trial-and-error component twiddling.

The on-chip current-sense amplifier drives an on-chip voltage follower, which drives the gate of a large on-chip output transistor Q_1. If this transistor is of the same type as the feedback transistor in the logarithmic current-sense amplifier, then the current flowing in the output transistor will be proportional to the input current to the current-sense amplifier. The output transistor has both source and drain coming out to pads, giving us the flexibility to do level and gain adjustment off-chip. The output transistor is incorporated into the off-chip video driver in an inverting mode. The source of the output transistor is tied to the appropriate rail, and the drain is pulled to the opposite rail with a resistor. To drive the monitor input, the output of this inverting amplifier is then either amplified with an inverting amplifier or followed with an emitter follower. The on-chip blanking transistors pull down on the appropriate place in the video driver circuit to blank the video signal. For the inverting configuration, the blanking transistor pulls down on the output of the inverting amplifier. For the noninverting, emitter follower, the blanking transistor pulls down on the input to the emitter follower. The four possible combinations of output

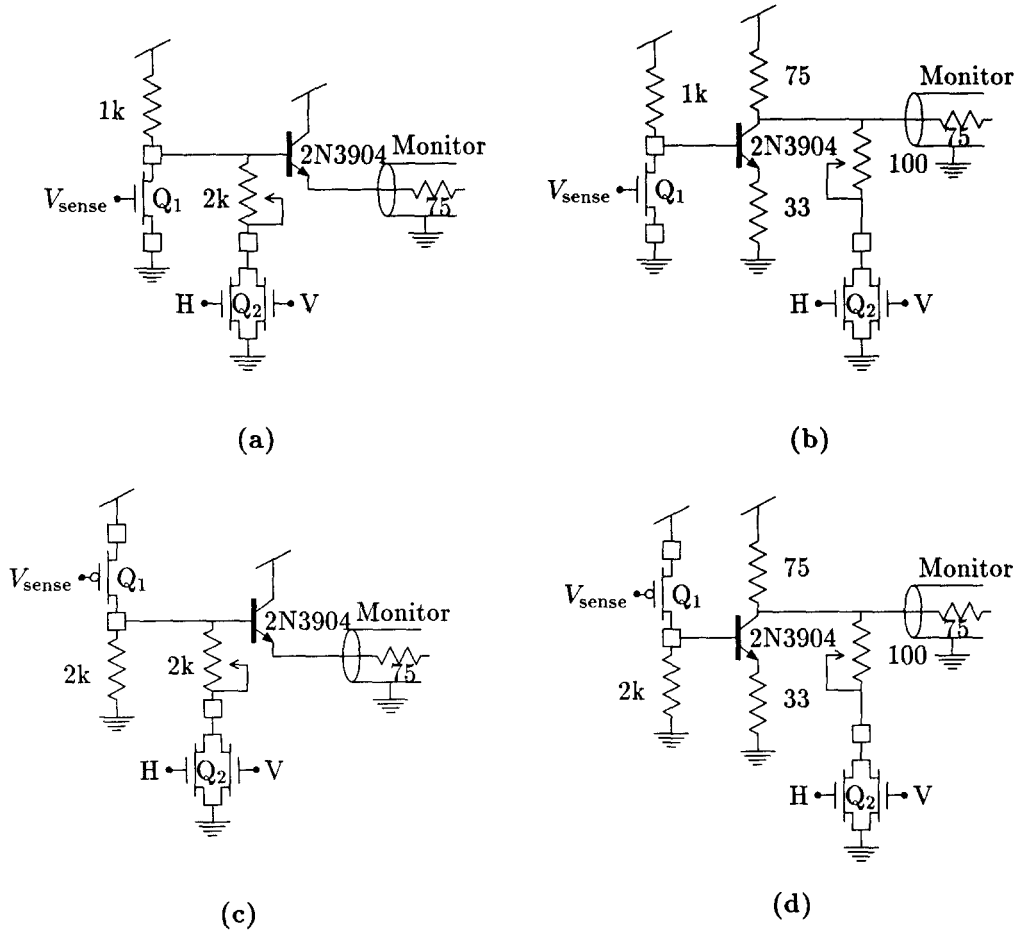

Fig. 14. Video-drive amplifiers. (a) Driver used for inverting internal signal with n-type output transistor. (b) Driver used for noninverting internal signal with n-type output transistor. (c) Driver used for inverting internal signal with p-type output transistor. (c) Driver used for noninverting internal signal with p-type output transistor. The Q_1 and Q_2 FETs are on-chip, and all other components are off-chip. V_{sense} comes from voltage-follower driven by output of curent-sense amplifier. Bonding pads are shown as small boxes. H and V are the on-chip horizontal and vertical blank signals; they drive the blanking FETs Q_2. the 75-Ω resistor shown in the monitor coax is inside the monitor; we do not supply it. The potentiometers are used to adjust the blanking level, and hence, the brightness and contrast of the video image. FET W:L ratios: blank transistors, 232:2; output transistors, 950:2; blank transistor ohmic resistance at V_g = 5 V: n-FET, 140 Ω; p-FET, 240 Ω.

transistor type and inverting-noninverting amplifier type shown in figure 14 cover most possible situations.

The preceding discussion has assumed that we are generating only a single color. If we want to generate a white picture, we connect the output from the video-driver to all three RGB inputs, or duplicate the output driver for each color. If we are scanning out more than one signal from the chip, we may want to generate a separate color with each channel. In this case, the blanking circuitry shown in figure 14 is slightly more complicated, since we cannot directly connect the separate video signals to a common blank.

In section, 5 we discuss the performance of the video drivers.

5. Performance

In this section, we give measurement results on the power consumption and speed of a typical scanner.

5.1. An Example Scanner

Figure 15 shows a photograph of a silicon retina equipped with the two-dimensional video scanner. The chip has 68 rows and 43 columns of pixels, and each row is counted 6 times. Figure 16 shows the sync and blank outputs from this video scanner. This chip was fabricated in 2-μm p-well technology through MOSIS.

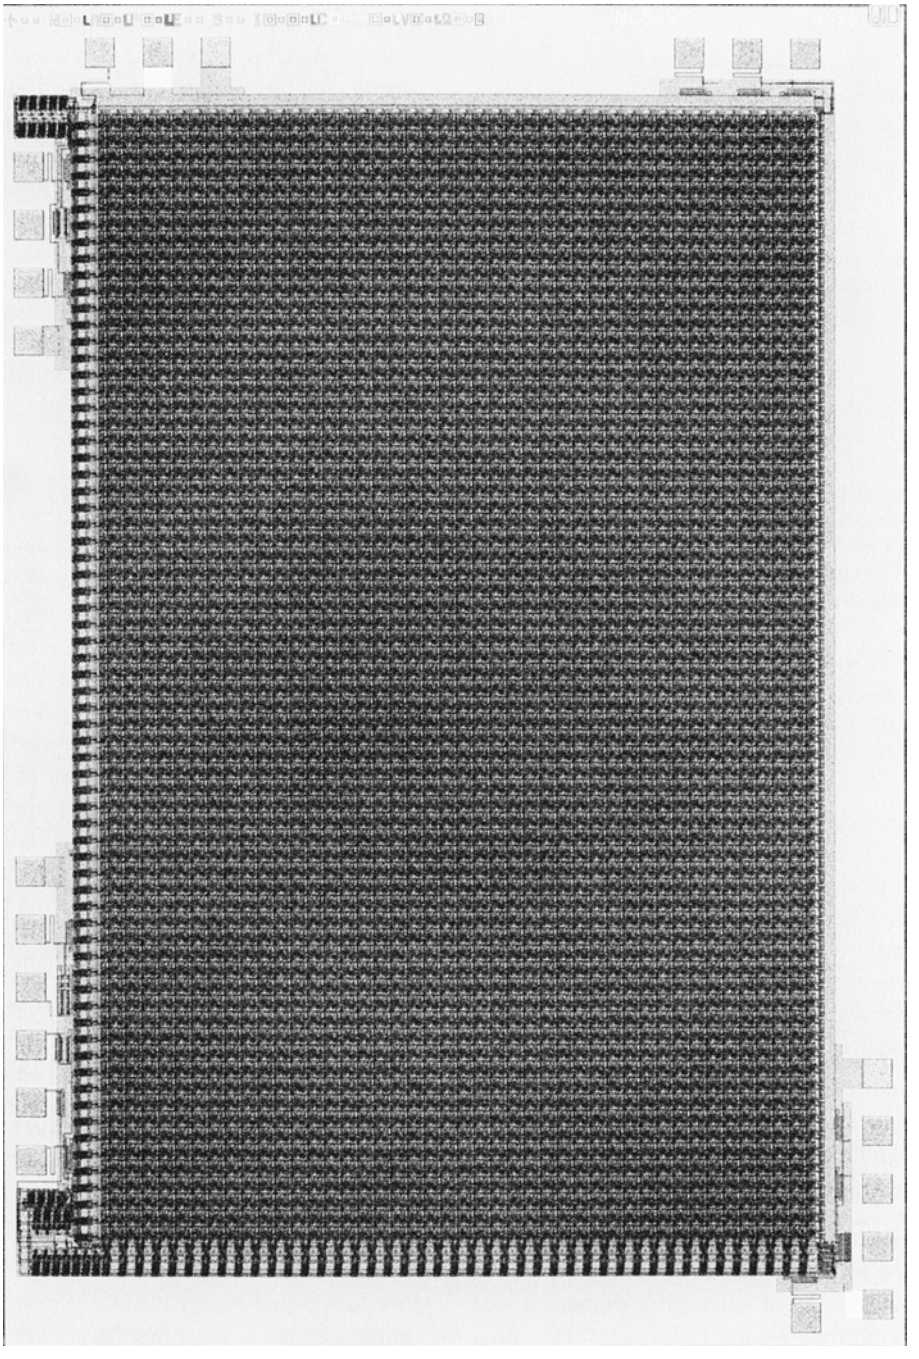

Fig. 15. Photomicrograph of a fabricated silicon retina with scanner. The chip measures approximately 4600 by 6800 microns. There are 43 columns and 68 rows of pixels. The horizontal sync and blank shift-registers and the Johnson counter are at the lower left of the chip. At the upper left are the vertical sync and blank shift-registers. At the lower right are the sense amplifiers and output pads. The pads at the upper right of the chip provide bias voltages for the core of the chip.

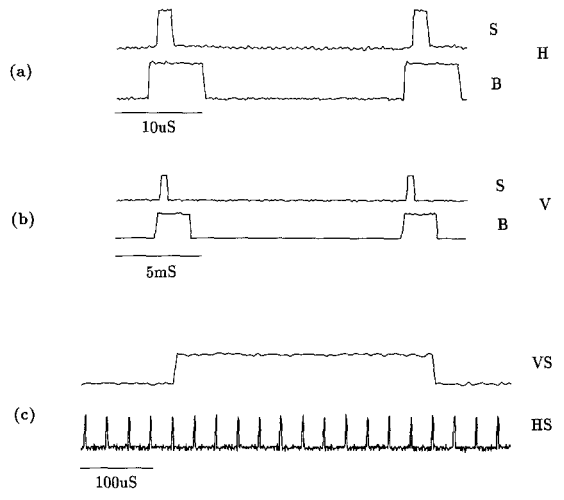

Fig. 16. Sync (S) and blank (B) signals from the video scanner shown in figure 15. (a) Horizontal (H) sync and blank signals. (b) Vertical (V) sync and blank signals. (c) Vertical and horizontal sync signals shown together. This chip has 68 rows and 43 columns. Each row is counted six times, for a total of 408 lines of video.

The behavior of n-well chips is similar. The timing shown in this figure produced excellent framing over a range of $\pm 10\%$ clock frequency.

5.2. Power Consumption and Speed

The digital parts of the scanner in figure 15 consume 6 mW of power when running at a 1.8 MHz clock rate with a 5 V supply voltage. At a clock frequency of 1.8 MHz, the vertical retrace frequency is 60 Hz. This scanner is capable of running at up to 11 MHz, far in excess of the required speed for chips of this complexity.

5.3. Current-Sense Amplifier and Video-Driver Performance

The on-chip logarithmic current-sense amplifiers shown in figure 13 have a usable bandwidth that is dependent on the pixel output current and on the voltage gain of the feedback amplifier. We use ring transistors with a W:L ratio of 20:2 in our feedback amplifiers, yielding a gain $A \approx 40$ and an effective speedup κA of about 25. When the pixel output current is very small, the current-sense amplifier will not stabilize the input line fast enough to drive the output at video rates. For moderate pixel output currents, on the order of tens of nanoamperes, the natural time constant of a typical pixel output line of one picofarad capacitance is less than one microsecond, and the bandwidth of the system begins to be limited by the bandwidth of the video-driver circuitry. As an example, the maximum measured bandwidth for the current-sense amplifier in figure 13(b) driving the video driver in figure 14(a) is 8 MHz. We have built complete systems with bandwidth in excess of 5 MHz, more than sufficient for display of arrays of more than 100 by 100 pixels.

6. Practical Design Details

In this section, we discuss practical design details and give a source for example layout.

6.1. Technology

We have fabricated and tested these scanners using a wide range of fabrication processes. The scanner design itself is generic, and we use the same layout for both *n- and p*-well technologies. Monitors generally trigger on the low-going edges of the sync signals, but we have found that the same sync generation output circuitry used for n-well technology works for p-well also. Th monitors care about the sign of the video-signal voltage, however, so we must pay some attention to the actual technology when designing output circuitry. For example, the video-driver circuits shown in figure 14 all utilize n-type blanking transistors that pull down to ground. The identical generic CIF layout, fabricated in the complementary technology will result in a p-type blanking transistor that pull up to V_{dd}. In this case, we can modify the video-drivers in figure 14(b) and (d) so that the blanking transistors pull up on the input to the inverting video-driver. The drivers in figure 14(a) and (c), however, would require additional off-chip active components to function correctly.

6.2. Power-Supply Separation

We have found it very beneficial to separate the power supply for the scanner and the core of the chip. In fact, we often have four power-supply pins: one for the digital scanner, one for the clock driver, one for the video circuitry, and one for the core of the chip. This added flexibility is provided at the cost of a few pins and often lets us isolate problem areas having to do with power utilization and latchup. In addition, the sensitive analog circuitry is isolated from the relatively large clock noise generated by the digital scanner.

6.3. Layout for Scanners

The value of exploratory design has been greatly enhanced by the availability of fast prototyping through the MOSIS fabrication service. It is no longer necessary that even relatively complicated projects be group efforts, involving months of design, and elaborate negotiations between foundry and customer. Instead, projects now can be designed by individuals in a few days.

Using any one of several readily available and affordable design tools, a designer can lay out the required circuitry for the chip core. This circuitry is integrated with a scanner frame using a simple silicon compiler that places pixels in rows and columns and abuts the scanner cells. To verify the layout of the chip, a layout extractor is used to generate the netlist of a small version of the circuit. Using a netlist comparison program, this netlist is compared with a netlist extracted from a schematic of the chip. When the design verifies correctly, a full-sized version of the chip is compiled and is electronically mailed to MOSIS. Six to eight weeks later, the packaged project comes back, ready to test.

Relatively complicated parts of designs, such as the scanners described here, are nonexploratory and exist only for circuit testability purposes. They still require considerable design and debugging, and a single mistake can kill the entire scanner operation. The aim of this article is to avoid a time-consuming duplication of the effort that has gone into development of these scanners. Hence, we are making the layout and schematics for these scanners available via anonymous ftp. To access these sources, you can ftp to the internet host **hobiecat.cs.caltech.edu** (preferably during off-hours), log in as **anonymous** and give your name as password. The scanner directory is **/usr/ftp/pub/scanners**. A README file in that directory will provide further details.

7. Summary and Conclusion

In this article we have given a tutorial description of how to design both one- and two-dimensional scanners, with emphasis on practical aspects of design and on the production of useful output images or signals. In our laboratory, we have used scanners on over 100 chip designs over the past three years. In fact, scanners act as our eyes into the microcosm of large arrays of VLSI elements, allowing us to observe the collective nature of the computation that the chip is performing, or the distribution and global effect of circuit offsets. For many circuits, such as retina or cochlea models or visual motion models, scanners provide insights into the spatio-temporal structure of the network operation that would be lost in a static view of a single node. For certain auditory processing chips, the use of a two-dimensional video scanner becomes essential in the process of converting the chip operation into a signal that can be understood and processed by our own visual system into a coherent picture of the computation. In these circuits, a global view is essential in understanding the circuit operation.

We do not regard these scanners as useful for chip-to-chip communication. The well-known pitfalls of sending analog information off-chip, with the attendant corruption by digital noise and the problem of matching references and other values between analog chips, is an obvious concern. Of equal concern, however, is the fundamental problem of temporal aliasing and bandwidth utilization. Even if we could noiselessly send analog values from one chip to another (or several others), these signals would still be discretely sampled in time—not a desirable quality for an analog, continuous-time, processing system. A serial sampler, such as a scanner, would not be biased in its sampling, but by the same token, it would waste its bandwidth on uninteresting events. That is, a scanner would send information from a location on the sending chip even if nothing were changing there. More sensible, in this regard, is a scheme that "spends" bandwidth on locations with interesting events [2]. These "event-driven" scanners remain to be proved practical, as do many aspects of multichip analog design; in the meantime, serial scanners will be indispensible adjuncts to exploratory analog VLSI design.

8. Acknowlegments

Our thanks to the Office of Naval Research for their support under grant NAV N00014-89-J-1675, and, of course, to the MOSIS fabrication service. Thanks also to Mohammed Ismail, Shih-Chii Liu, Lyn Dupré, Xavier Arreguit, Buster Boahen, and anonymous reviewers for helpful comments and discussions.

References

1. K. Bult and G.J.G.M. Geelen, "A fast-settling CMOS op amp for SC circuits with 90-dB gain," *IEEE J. Solid-State Circuits*, vol. 25(6), pp. 1379–1384, 1990.
2. M.A. Mahowald, "Evolving analog VLSI neurons," in T. McKenna, J. Davis, S. Zornetzer (eds.), *Single Neuron Computation, Proceedings of ONR Meeting, Woods Hold, MA*, Academic Press, to be published 1991 or 1992, Ch. 15.

3. M.S. McGregor, P.B. Denyer, and A.F. Murray, "A single-phase clocking scheme for CMOS VLSI," *Proc. Stanford Conference on Very Large Scale Integration,* MIT Press: Cambridge, MA, 1987, pp. 257–271.
4. C.A. Mead, *Analog VLSI and Neural Systems,* Addison-Wesley: Reading, MA, 1989.
5. C. Mead and J. Wawrzynek, "A new discipline for CMOS design," in H. Fuchs (ed.), *1985 Chapel Hill Conference on Very Large Scale Integration,* Computer Science Press: Chapel Hill, NC, 1985, pp. 87–104.
6. M. Sivilotti, "Wiring considerations in analog VLSI systems, with application to field-programmable networks," Ph.D. thesis, Dept. of Computer Science, California Institute of Technology, Pasadena, CA, June 1991.
7. M.A. Sivilloti, M.A. Mahowald, and C.A. Mead, "Real-time visual computations using analog CMOS processing arrays," in P. Losleben (ed.), *Proceedings of the Stanford Conference on Very Large Scale Integration,* MIT Press: Cambridge, MA, 1987, pp. 295–311.
8. L. Watts, "One-dimensional resistive networks and scanners," internal laboratory memorandum, Physics of Computation Laboratory, California Institute of Technology, Pasadena, CA, 1990.
9. G. Wegmann and E.A. Vittoz, "Analysis and improvements of accurate dynamic current mirrors," *IEEE J. Solid-State Circuits,*" vol. 25(3), pp. 699–706, 1990.

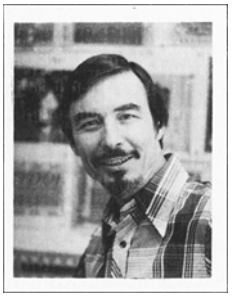

Carver A. Mead, Gordon and Betty Moore Professor of Computer Science has taught at the California Institute of Technology, for more than 30 years. He has contributed in the fields of solid-state electronics and the management of complexity in the design of very large scale integrated circuits, and has been active in the development of innovative design methodologies for VLSI. He has written with Lynn Conway, the standard text for VLSI design, *Introduction to VLSI Systems.* His recent work is concerned with modeling neuronal structures, such as the retina and the cochlea using analog VLSI systems. His new book on this topic, *Analog VLSI and Neural Systems,* has recently been published by Addison-Wesley. Professor Mead is a member of the National Academy of Sciences, the National Academy of Engineering, a foreign member of the Royal Swedish Academy of Engineering Sciences, a Fellow of the American Physical Society, and a Life Fellow of the Franklin Institute. He is also the recipient of a number of awards including the centennial medal of the IEEE.

Tobias Delbrück, born March 16, 1960 in Pasadena, is a graduate student in the Department of Computation and Neural Systems at the California Institute of Technology. His main scientific interests lie in the use of novel analog VLSI technology for artificial vision.