


The Optical Mouse: Early Biomimetic Embedded Vision

Richard F. Lyon

Abstract The 1980 Xerox optical mouse invention, and subsequent product, was a successful deployment of embedded vision, as well as of the Mead–Conway VLSI design methodology that we developed at Xerox PARC in the late 1970s. The design incorporated an interpretation of visual lateral inhibition, essentially mimicking biology to achieve a wide dynamic range, or light-level-independent operation. Conceived in the context of a research group developing VLSI design methodologies, the optical mouse chip represented an approach to self-timed semi-digital design, with the analog image-sensing nodes connecting directly to otherwise digital logic using a switch-network methodology. Using only a few hundred gates and pass transistors in 5-micron nMOS technology, the optical mouse chip tracked the motion of light dots in its field of view, and reported motion with a pair of 2-bit Gray codes for x and y relative position—just like the mechanical mice of the time. Besides the chip, the only other electronic components in the mouse were the LED illuminators.

Fig. 1 The Xerox optical mouse chip in its injection-molded dual-inline package (DIP) of clear plastic, with pins stuck into a conductive packaging foam. The bond wires connecting the chip's pads to the lead frame are (barely) visible.



Richard F. Lyon
Google Inc., Mountain View CA, e-mail: dicklyon@acm.org



Fig. 2 The Winter 1982 Xerox World internal magazine cover featuring the Electronics Division and their 3-button mechanical and optical mouse developments, among other electronic developments. The three-button mouse shipped on SmallTalk and Lisp machines, but the 8010 and 6085 office systems used a two-button version [18].

1 Early Mice

At Xerox PARC, wheel mice and ball mice went through several generations in the 1970s, and when Xerox first delivered their commercial workstation product—the Xerox “Star” 8010 Information System—in the early 1980s, it shipped with a ball mouse. But the optical mouse that I first designed in 1980 made it to product a few years later, and displaced the ball mouse in favor of this less expensive and more reliable technology based on a single-chip VLSI sensor with logic (see the chip photo, Figure 1).

The early mechanical mice worked well when they were clean, but tended to gum up over time. They did not have removable balls like the later Apple Macintosh mice, so had to be disassembled and cleaned by a technician. Due to these difficulties, several researchers at PARC had worked on developing no-moving-part optical alternatives. I had the advantage of being able to review several different previous attempts that had resulted in the filing of invention proposals, but no prototypes or patents, as well as the advantage of having the new custom VLSI chip prototyping system available (we had developed this capability in Lynn Conway’s VLSI Systems Group). The previous optical mouse attempts were based on good





Fig. 3 This 1985 product brochure shows the two-button Xerox optical mouse on its special mouse pad. The corner of the mousepad not shown in this image is detailed in Figure 4.

Fig. 4 The mousepad was paper, sold in packs of 25 sheets. The pattern was a hexagonal array of light dots in a dark field, shown here at approximately actual size. Effective mousepads could be made by a copier.



concepts for one-dimensional motion sensing, but the attempted extensions to 2D were not workable.

At the same time as my optical mouse chip, Steve Kirsch independently invented a way to make two one-dimensional trackers work together, using two different LED colors and a mousepad with special colors of orthogonal stripes. In this design, the mouse's coordinate system was in the gridded pad; for this reason, Jack Hawley, maker of the X063X ball mouse for Xerox, called Kirsch's devices "pseudo mice" [28]. My design that Xerox pursued was truly two-dimensional, with coordinates relative to the mouse body, like a mechanical mouse.

Kirsch's mouse was a big success, being adopted for early Sun workstations, but was not really an embedded vision system; it had a few photosensors, but no

Fig. 5 The photodiode is an n-type region in an nMOS process. These and other diagrams are scanned from my 1981 PARC optical mouse report [24].





Fig. 6 The photodiode was incorporated as a dynamic node, going directly into a logic gate, just as other switched dynamic nodes were used in the methodology we had been teaching [9, 10].



imaging array. The Xerox mouse, on the other hand, had a 2D imaging array (4x4 pixels) with simple correlation-based spot tracking, and was an embedded vision system in that sense. Further, it was specifically neuromorphic and biomimetic in the way it incorporated lateral inhibition in its imager.

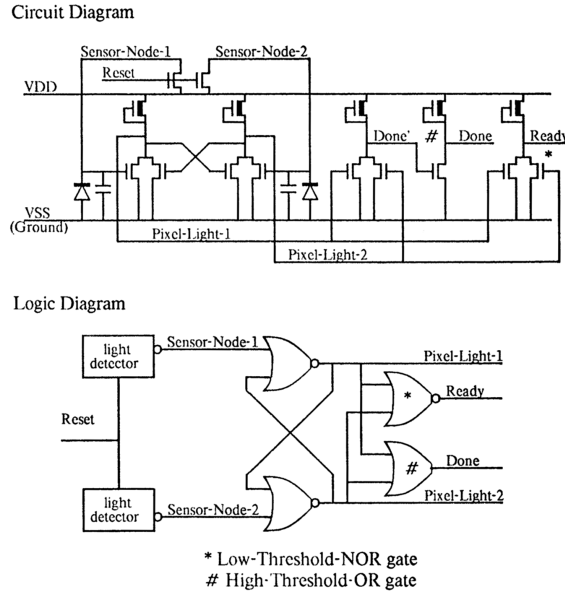
Xerox's optical mouse development was celebrated on the cover of their internal Xerox World magazine in 1982, as shown in Figure 2. The Xerox optical mouse was sold with a range of office workstations, such as the 6085 shown in Figure 3, as well as with Xerox Lisp machines, Tektronics SmallTalk machines, and high-end copier/duplicator products—none of which were high-volume products. Xerox was not successful in their attempts to sell licenses to their optical mouse patents, even after the market for mice exploded on the introduction of the Apple Macintosh with its low-cost ball mouse in 1984.

This chapter reviews the ideas that went into the optical mouse's very application-specific embedded vision system.

2 Image Sensing with Lateral Inhibition

The elementary light detector in an nMOS process is a PN photodiode, with the P region being the substrate and the N region being a diffusion region, as shown in Figure 5 (a "green" area as we taught it in the red/green/blue/black scheme at the time). As shown in Figure 6, a reset transistor back-biases the photodiode to an initial high voltage, and the voltage decays as the diode collects photoelectrons. The photodiode voltage can be used as an input to digital logic (shown in the figure as an inverter), provided the logic is designed to tolerate the intermediate analog values that the voltage will necessarily go through slowly.

Fig. 7 The circuit and logic diagrams for a two-pixel imager. Through the cross-coupled feedback, each pixel inhibits the other. It is essentially a flip-flop, with two known stable states, that can hang for a while in an unstable all-dark state after being reset to that state (like an RS flip-flop with both set and reset asserted at the same time). The gates on the right detect “Ready” when the “Reset” signal has put the imager into the unstable state (Pixel-Light-1 and Pixel-Light-2 both low), and “Done” when light discharging one or both photodiodes has allowed the flip-flop to commit to one of the stable states (Pixel-Light-1 or Pixel-Light-2 high).



The imaging strategy in the optical mouse relies heavily on an engineering interpretation of lateral inhibition as a nonlinear scheme for arriving at a stable image, independent of light level. In the simplest example, as shown in Figure 7, a two-pixel imager uses mutual inhibition in a form that forces the system to decide which one of the two pixels receives more light than the other.

In systems of more than two pixels, each pixel can inhibit, and be inhibited by, pixels within some radius. For example, in the four-pixel imager of Figure 8, pixels at distance 1 and 2 mutually inhibit, but the end pixels, at a distance of 3, do not inhibit each other. This radius-of-inhibition idea has an obvious extension into two dimensions.

The stable images from the four-pixel imager, that is, images where no more pixels can change from “dark” to “light”, are just three: 1001, 0100, 0010. These patterns are ideally suited to imaging and tracking light lines spaced at about 3 pixels, in a dark background, as shown in Figure 9. Figure 8 also shows logic for comparing successive images, by local binary cross-correlation, and keeping track of position by a counter, under control of the timing logic of Figure 10.

When I studied the literature on lateral inhibition in the 1970s, based mostly on the compound eye of the horseshoe crab *Limulus*, it seemed to be based mostly on a linear systems view, with subtraction from neighboring sensor elements resulting in a bandpass or highpass filtering effect. In his 1967 book *Sensory Inhibition*, von Békésy [36] explored the effects of lateral inhibition in various sensory systems, including vision, hearing, touch, and taste, in essentially linear systems terms, for “sharpening” the response to stimuli. Papers argued that the effect could be very

Fig. 8 A four-pixel version of the imager scheme of Figure 7, with timing and motion detection logic. This linear imaging array has each pixel inhibiting neighbors up to two pixels away, but the extreme end pixels do not inhibit each other, illustrating the idea of a radius of inhibition. The done-detect logic is looking for every pixel to either be indicating light, or be inhibited by one that is light; such states are stable.

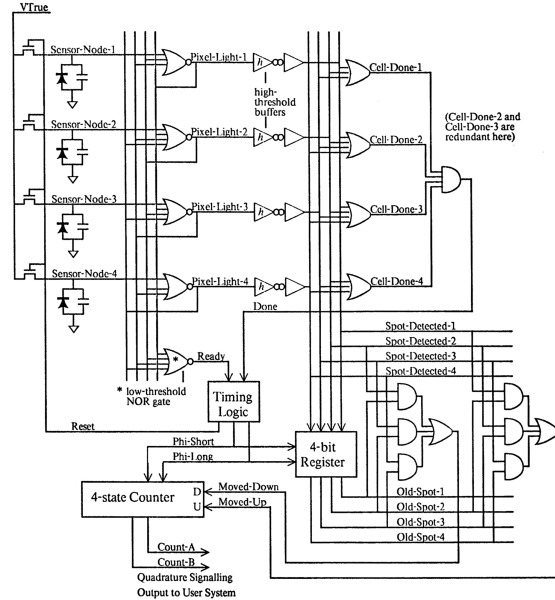



Fig. 9 Light lines in a dark background can be imaged as shown here, and tracked using a linear array of just four pixels, using the circuit of Figure 8.



accurately modeled as linear. For example, while Knight et al. [19] noted that crab eye was quite compressive, with a “graded response to a wide range of light intensities (to a factor of about 10^7 in intensities for *Limulus*)”, they also said that “the response to an intensity decrement is the close mirror image of the response to the corresponding increment. This suggests that we are dealing with a so-called ‘time-invariant linear’ system.” In most subsequent analysis, the locally short-time approximately linear filtering behavior was modeled as if it were accurate, and the extremely nonlinear large-scale long-time adaptive behavior was ignored, so the important property of lateral inhibition as a strategy for normalizing the response across a wide dynamic range was missed. Barlow finally forcefully explained why lateral inhibition is a key part of the strategy for wide dynamic range vision [3]: “If critical limiting factors are emphasized one says that lateral inhibition, colour opponency, and the gain changes of light and dark adaptation, are necessary to transmit information about the light intensities in subdivisions of the visual image, because

the available information has a much wider dynamic range than can be transmitted directly down a nerve fibre in a reasonable time interval.” It may be that I had already inferred that from some of his pre-1981 writings, though they were not as clear on this point.

My original chip layout, done “by hand” on a Xerox Alto with a mouse and the ICARUS IC editor [14], is shown in Figure 11. I had no idea how much light would be needed, or what the contrast ratio of the imaged surface would be, so I used lateral inhibition to make the logical function independent of overall light level. I did no calculations of photodiode capacitance, photon flux, noise margins, or any of the things that I had to learn about decades later, designing image sensors at Foveon. I just wanted to make sure it would work, and generate trackable binary patterns, at any light level. It was very clear that a linearly responding sensor system would be useless, but I saw the way to use the well-studied networks of lateral inhibition for their nonlinear normalizing effects. This approach is now more common, in concepts such as “contrast gain control”—an idea proposed in 1978 by Shapley and Victor [34] for better modeling the cat retina. It is sometimes implemented as “divisive gain control”, dividing by a neighborhood average as a way to control the local gain [32]. In the mouse sensor, the lateral inhibition implements a source of a *race* in time; the first sensor channel to get enough signal inhibits its neighbors, and if there is a near tie, they inhibit each other in a positive feedback loop until one wins.

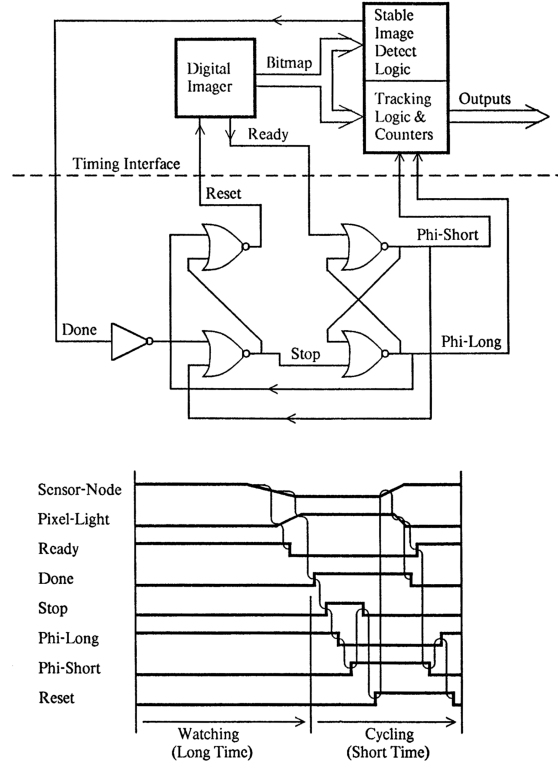
I initially investigated lateral inhibition for robust automatic gain control in hearing models, starting at Xerox before the optical mouse work; I have continued to use the concept this way in my current work on machine hearing [22]. Another interpretation of nonlinear lateral inhibition, rather than as gain control, is as sparse coding—which is essentially what the mouse sensor does. To arrive at most outputs being zero due to inhibition, and only one or a few outputs being active, a competitive or comparative dynamic process operates on initially small differences, resulting in a *winner-take-all* effect. This concept was later used in the silicon retina [21]. Variations on sparse coding and winner-take-all coding have become popular in computer vision in recent decades.

3 Symmetric Mutual Inhibition

The two-pixel and four-pixel examples show symmetric patterns of mutual inhibition: if pixel A inhibits pixel B, then pixel B inhibits pixel A. This is not the only kind of logical inhibition pattern that can be built, but it has the useful property that it leads to a set of stable states that are easy to enumerate, and it is easy to build logic to determine when a stable state has been reached.

Networks with symmetric inhibitory connections came to be known as Hopfield networks, and were valued for these properties [1]. I recognized the value of the final states being stable, and of being able to predict, enumerate, and detect stable states, when working on extending the one-dimensional tracking idea to two dimensions.

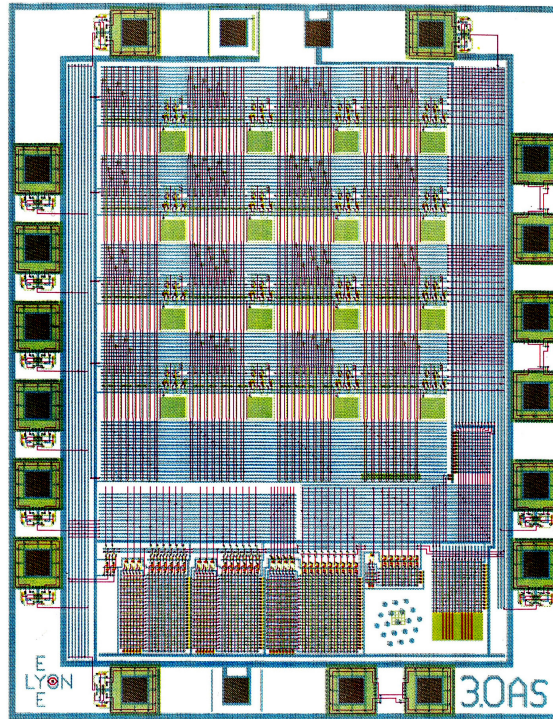
Fig. 10 The “Ready” and “Done” signals from an imaging array with lateral inhibition, such as that in Figure 8, cooperate with this non-overlapping-clock generation circuit to yield a free-running self-timed imaging system. The synchronous digital logic parts of the system use the two-phase non-overlapping clocking methodology that we were teaching for digital system design as part of the Mead–Conway VLSI design revolution [30]. The duration of the “long” clock phase would be whatever time was needed for the imager to reach a stable binary state—faster at high light levels, slower at low light levels.



In this sense, I was building both a Hopfield network, though not of the scale or application envisioned by John Hopfield [16], and also a silicon retina, though not of the scale or generality envisioned by Carver Mead [29]. The imager with lateral inhibition was a nonlinear dynamical system, before that concept was popularized with the notion of chaotic attractors; but the attractors in the mouse imager are stable by design, not chaotic or periodic.

In Figure 8, the column of four NOR gates on the left, with feedback from their outputs to the inputs of others, is the Hopfield network. The rest of the circuitry is for resetting it, determining when it is reset, determining when it reaches a stable state, repeating that sequence indefinitely, and tracking the motion implied by the successive states. Comparable lateral inhibitory connections of limited range are found in real retinas, involving horizontal cells and amacrine cells [11], and in silicon retinas [13].

Fig. 11 The layout of the first-generation Xerox optical mouse chip, which I did quickly toward the end of 1980, had the lateral inhibition, done detection, and image storage logic distributed in the pixel cell array, along with image cross-correlation logic. The arrays at the bottom were essentially programmed logic arrays (PLAs), programmed as the timing generator, the counters, and the logic that converted the sensed move directions, or image cross-correlations, to counter increments. Six of the bond pads are cross-coupled inverter pairs, to debounce the SPDT switch contacts of the three mouse buttons. Of the eleven output pads, four are for the motion encoding and seven are for observing internal timing signals.



4 Metastability

The ability of the Hopfield network to fall into a stable state depends on positive feedback. A multi-stable system also has metastable states, or saddle points, where it can hang for a long time before deciding which side to fall toward. The mouse's done-detection logic waits to see when it has committed, but this strategy only works if the metastable states are static, not oscillatory.


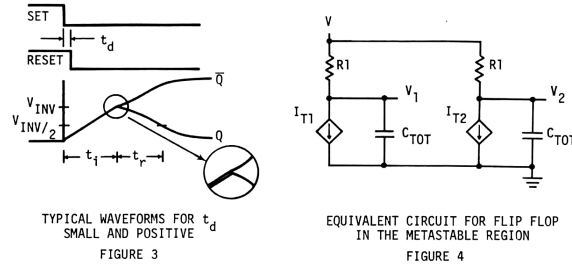


Fig. 12 Chaney and Molnar showed in 1973 that a pair of cross-coupled TTL NAND gates had an oscillatory metastable state.

Oscillatory metastable states were well known to us in the 1970s, having been reported by Chaney and Molnar [7], who showed the oscillatory metastable behavior of a pair of cross-coupled TTL logic gates; Figure 12 shows random samples of a



If the circuit has voltages $V_1 = V_{INV} - \Delta V$ and $V_2 = V_{INV} + \Delta V$ at time t_1 , then the voltage V_1 as a function of time is given by

$$V_1 = V_{INV} - \Delta V \cdot \exp[p(\Delta-1)(t-t_1)] \quad t \geq t_1 \quad (8)$$

where $p = \frac{C_G \cdot \mu(V_{INV} - V_{DD} - V_{TPU})}{C_{TOT} \cdot L^2 \cdot k}$ = the bandwidth of one inverter stage of the Flip Flop

$$A = k \cdot \frac{V_{INV} - V_{TPD}}{V_{INV} - V_{DD} - V_{TPU}} = \text{magnitude of the low frequency gain of one inverter stage}$$

thus the τ_r defined in Section II is given by

$$\tau_r = \frac{1}{p(\Delta-1)} = \frac{C_{TOT} L^2}{C_G \mu} \cdot \frac{k}{k(V_{INV} - V_{TPD}) - (V_{INV} - V_{DD} - V_{TPU})} \quad (9)$$

$$\text{and } V_1 = V_{INV} - \Delta V \cdot \exp\left\{-\frac{t-t_1}{\tau_r}\right\} \quad (10)$$

$$V_2 = V_{INV} + \Delta V \cdot \exp\left\{-\frac{t-t_1}{\tau_r}\right\} \quad (11)$$

Fig. 13 Chaney and Rosenberger showed, at the 1979 Caltech Conference on VLSI, that cross-coupled nMOS NOR gates would exit the metastable point via a simple exponential divergence.

flip-flop output being put into and exiting its metastable state. Charlie Molnar had worked with us at Caltech during the early VLSI developments, and was influential in our thinking about system timing, as described in Chuck Seitz's chapter in the Mead-Conway book [33]. In my initial report, I credited Seitz for the done-detection idea:

Note that we do not use the inhibition NOR gate output itself for done-detection, but a buffered version of it after a high threshold buffer (inverter pair); this is the easiest way to prevent false done-detection during a metastable condition [Seitz 1980]. The buffered signal is not used for inhibition, since that would make it participate in the metastable condition, and because the extra delay would cause oscillatory metastable states.

Chaney and Rosenberger had shown that the metastable state of a pair of cross-coupled nMOS NOR gates, like that of the two-pixel imager of Figure 7, would be a simple unstable equilibrium, which would diverge exponentially toward a stable state, without oscillation [8]; Figure 13 shows their model and analysis from the 1979 Caltech Conference on VLSI. This was the behavior I needed, and I had reasoned that with one capacitance per node, a multinode generalization, such as the one in Figure 8, or the larger one in the full 2D imager, would have similar dynamics in exiting its metastable states. Yet I had no proof.

Hopfield showed in 1984 that with symmetric interconnection weights, such a network can be characterized by an energy function, and that any state change reduces the energy, until it settles into a stable state, a local minimum [17]. Hopfield's differential equation formulation with one state variable per node exactly describes

Inhibition Neighborhoods	Inhibition Radius	Stable Images and how many of each	Total Images
	<1	1	1
	1.1	2 8 8 8 2	42
	1.5	8 8 8 8 4	79
	2.1	2 1 8 8	43
	2.3	4 1 8 4 4 4	25
	2.9	4 4 1 8 4	21
	* 3.0s	4 4 2 8 8 4	30
	3.1	4 4 2 8 8	26
	3.2	4 2 8	14
	3.7	4 2 8	14
	4.3	4 4 8	16

Fig. 14 Inhibition patterns and spot patterns for a range of radii of inhibition. The 3.0s inhibition yields a useful set of 30 stable images with one or two spots each.

the network of interconnected nMOS NOR gates in the optical mouse imager, so we can be sure, in retrospect, that the metastable states of that circuit are non-oscillatory. The cross-coupled TTL gates that Chaney and Molnar analyzed had extra internal (not symmetrically interconnected) nodes with their own state and delay, which is why the metastable states in that case could be oscillatory.

5 Two Dimensions

The idea of a 1D line tracker such as that of Figure 9 has various possible extensions into two dimensions. Extensions that use axis-aligned patterns of lines or dots have problems when the mouse is rotated, which is why pre-1980 attempts at PARC did not get very far. Using the idea of a radius of inhibition in two dimensions, I set about finding a way to track less rigid patterns of light dots in a dark field. An ideal pattern seemed to be a close-packed hexagonal array of dots, allowing the imager at arbitrary angles to treat the dots as being randomly arrayed, with a characteristic distance between them.

After evaluating various radii of inhibition and imager sizes, I found that a 4x4 array with a “3.0 special” radius of inhibition would yield a set of stable images

Fig. 15 How a 4x4-pixel imager might see one, two, three, or four light dots when viewing a hexagonal array. For the inhibition pattern we chose, the three- and four-dot versions are not allowed, so only two of the dots on opposite edges (and not on adjacent corners) will be seen in those cases.






Fig. 16 This programmed-logic array (PLA) converts the cross-correlation bits to counter-control signals. If an image has no light pixel adjacent to one in the previous frame, the “Jump” output is asserted to signal this exception, and the counters are not moved, by not asserting either “Half” or “Full” as the increment size.



containing either one light dot in the central 2x2 area, or two light dots on opposite edges—a total of 30 possible state binary images, as shown in Figure 14. Here “special” means that pixels at a distance of exactly 3.0 pixel spaces will inhibit each other if they are corners, but not otherwise; this scheme eliminates patterns of 3 and 4 dots, but allows patterns of dots on opposite edges, so that motion can be detected, much as with the lines on opposite ends of the linear array of 4. In the 2D case, a pattern of light dots in a hexagonal array, as shown in Figure 15, works well.

When the radius of inhibition exceeds $2\sqrt{2}$, a light spot in an image cannot be adjacent or diagonally adjacent to two different spots in a previous or subsequent image, so we will not have any ambiguity of which direction a spot moved. Further-

Fig. 17 Four-phase quadrature encoding, or 2-bit Gray code, is reported in each dimension. Motion in one direction makes square waves in quadrature (90 degrees out of phase with each other).



more, with the 3.0 special inhibition pattern, since there can only be at most two spots (as shown in Figure 14), there will be at most two motion directions involved in a correlation of one image with the next. Computing the average of two moves is easy, using a half-step bit in a position counter: when an old image and a new image each have two spots, with different apparent move directions, their average can still be represented in terms of half steps. The logic to drive the half and full steps from the detected correlations is shown in Figure 16, to control the output quadrature signals shown in Figure 17.

The clocked logic does not have much state besides the sensor pixels themselves: just one 16-bit old image to be compared with a new image, and a pair of three-bit counters for x and y positions; compare the 4-bit image register and single counter of the one-dimensional version in Figure 8. Only two bits per dimension are taken as output, using a quadrature encoding, also known as a two-bit Gray code. This is the same encoding that ball mice generated by a pair of shaft encoders (optical shaft encoders in some designs). It assures that a device receiving the signals asynchronously will not get an error from two bits switching at not quite the same time.

6 Dynamic Logic

The two-phase nonoverlapping-clock pass-transistor-based dynamic logic used in the Mead–Conway nMOS VLSI methodology was almost a good match to the self-timed optical sensor approach. But since the duration of the long phase was not bounded, dynamic nodes set on the short phase could have decayed away, especially due to light falling on the chip, during that time. This potential problem was anticipated and was easily avoided by adding gated positive feedback to those nodes during the long phase, so that all data was held statically during those times. With this simple addition to the usual two-phase latches, and with the short phase being less than a microsecond, the logic was robust enough that no light shielding was needed.

7 Testing

I tested the first mouse chip by wiring it into the mouse port of my Xerox Alto computer and projecting patterns onto it. When I got the cursor to move in all directions, I declared success. Unfortunately, more rigorous testing was complicated by the fact that I had forgotten to give the chip any electrical input paths. With my team-mate Martin Haerberli, we soon set about designing the next version, with a more compact pixel array better suited to a short optical path, and with inputs that would allow selectively discharging any set of photodiodes, so that all the logic could be tested on a standard electronic chip tester [26]. The resulting product chip layout is shown in Figure 18, and the cover of the magazine that featured it is shown in Figure 19.

8 Going Meta

My manager at the time, Lynn Conway, always had (and ever since has) encouraged me to “go meta” with my ideas, which is why my original optical mouse report included the subtitle “and an Architectural Methodology for Smart Digital Sensors.” A condensed version of the report, with the same title, was created to go with my invited opening talk at the 1981 *VLSI Systems and Computation* meeting at CMU

Fig. 18 The layout of the second-generation Xerox optical mouse chip, completed by Martin Haerberli and Robert Garner after I left Xerox, incorporated more regular arrays outside the more compact photodiode array, which was a more efficient use of space. This new version also incorporated testability features; a second connection to each photodiode allowed it be discharged electrically, simulating light falling on it.






Fig. 19 The redesigned optical mouse chip with testability features was featured—via Charles Bragg’s 1966 painting *Salute*—on the cover of *VLSI Design* magazine in 1982 [26]. The caption reads “There are times when your mouse must be able to see. A single chip may be the solution.”.

[20]. The methodology was basically to combine the Mead–Conway digital design methods, including concepts of self-timed logic, with analog sensors such as photodiodes, in a way that leads to simple and elegant semi-digital designs that could be prototyped on a by-then-standard MOS fabrication service based on Conway’s simplified design rules [25].

The 1981 patent filing (see Figure 20) did not go very meta; it was divided into an imaging array with predetermined stable output states (... each of said cells coupled with several of said neighboring cells in said array to alter the output of said neighboring cells resulting in the recognition of said intermediary pattern as being one of a limited plurality of predetermined such patterns ...), and a cursor control device using it.

In 1982, I presented the optical mouse design at the *Physics of Computation* class that Mead, Hopfield, and Feynmann were jointly running at Caltech. This connection led to my joining the Caltech faculty as visiting associate for 13 years, where I helped Mead and his students with a wide range of neuromorphic vision and hearing chips. The mouse design influenced the development of some of their silicon retina ideas, such as motion sensing chips [35], and the winner-take-all and the address-event schemes for sparse digitization of analog signals [21, 27], as well as silicon cochlea chips [23, 37]. Researchers elsewhere took the ideas in different directions, such as a 60x60 binary smart imager in CMOS [15].

The small size of the optical, as opposed to mechanical, motion sensor allowed researchers to experiment with other configurations, such as following up the sug-

9 Conclusion

The optical mouse was a successful union of ideas in VLSI design, vision, and neural networks. In hindsight, it was a smart camera and an embedded vision system before those concepts were invented. It was a silicon retina before Carver Mead coined that term, and a Hopfield network before John Hopfield invented that concept, and a self-tracker before Gary Bishop came up with visual odometry. It did this all with only 16 pixels, so it was simple, rather than powerful or general. The chip was at the LSI level of complexity, but was a teaching vehicle and popular example for methods of digital VLSI system design that it embodied.

I enjoyed lecturing on this development at universities all around the world, and I still have my viewgraphs, printed on the world's first color laser printer (Gary Starkweather's "Puffin") in 1981, in case anyone would like to hear a rerun.

Acknowledgements I recall fondly the help of a great many people at PARC during the evolution of the ideas in the optical mouse, but I would like to single out just a few. Carlo Séquin taught me about how silicon detects light, giving me the basis to build on. Dan Ingalls provided ideas about two-dimensional tracking. Chuck Thacker reviewed the first optical mouse chip layout and found two errors that would have killed it. Martin Haerberli volunteered to do the second-generation layout, and took over the effort as I was preparing to leave PARC to lead a speech recognition group at Schlumberger Palo Alto Research. Lynn Conway, Bert Sutherland, and George Pake provided the supportive environment that allowed such ideas to take root.

References

1. Yaser S. Abu-Mostafa and J. St Jacques. Information capacity of the Hopfield model. *IEEE Transactions on Information Theory*, 31(4):461–464, 1985.
2. Xavier Arreguit, F. Andre Van Schaik, François V. Bauduin, Marc Bidiville, and Eric Ræber. A CMOS motion detector system for pointing devices. *IEEE Journal of Solid-State Circuits*, 31(12):1916–1921, 1996.
3. Horace B. Barlow. The Ferrier Lecture, 1980: Critical limiting factors in the design of the eye and visual cortex. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 212(1186):1–34, 1981.
4. Ahmed Nabil Belbachir. *Smart Cameras*. Springer, 2009.
5. Gary Bishop. *Self-Tracker: A Smart Optical Sensor on Silicon*. PhD thesis, University of North Carolina at Chapel Hill, 1984.
6. Stuart K. Card. The human, the computer, the task, and their interaction: Analytic models and use-centered design. In David M. Steier and Tom M. Mitchell, editors, *Mind Matters: A Tribute to Allen Newell*, pages 259–312. Psychology Press, 1996.
7. Thomas J. Chaney and Charles E. Molnar. Anomalous behavior of synchronizer and arbiter circuits. *IEEE Transactions on Computers*, 100(4):421–422, 1973.
8. Thomas J. Chaney and Fred U. Rosenberger. Characterization and scaling of MOS flip flop performance in synchronizer applications. In Charles L. Seitz, editor, *Caltech Conference on VLSI*, pages 357–374. California Institute of Technology, 1979.
9. Lynn Conway. The MPC adventures: Experiences with the generation of VLSI design and implementation methodologies. *Microprocessing and Microprogramming*, 10(4):209–228, 1982.
10. Lynn Conway. Reminiscences of the VLSI revolution: How a series of failures triggered a paradigm shift in digital design. *IEEE Solid-State Circuits Magazine*, 4(4):8–31, 2012.

11. Paul B. Cook and John S. McReynolds. Lateral inhibition in the inner retina is important for spatial tuning of ganglion cells. *Nature Neuroscience*, 1(8):714–719, 1998.
12. Hansjürgen Dahmen, Alain Millers, and Hanspeter A. Mallot. Insect-inspired odometry by optic flow recorded with optical mouse chips. In Dario Floreano, Jean-Christophe Zufferey, Mandyam V. Srinivasan, and Charlie Ellington, editors, *Flying Insects and Robots*, pages 115–126. Springer, 2010.
13. Rodney Douglas, Misha Mahowald, and Carver Mead. Neuromorphic analogue VLSI. *Annual Review of Neuroscience*, 18:255–281, 1995.
14. Douglas G. Fairbairn and James A. Rowson. ICARUS: An interactive integrated circuit layout program. In *15th Conference on Design Automation*, pages 188–192. IEEE, 1978.
15. P. Garda, A. Reichart, H. Rodriguez, F. Devos, and B. Zavidovique. Yet another mesh array smart sensor? In *9th International Conference on Pattern Recognition, 1988*, pages 863–865. IEEE, 1988.
16. John J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
17. John J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81(10):3088–3092, 1984.
18. Jeff Johnson, Teresa L. Roberts, William Verplank, David Canfield Smith, Charles H. Irby, Marian Beard, and Kevin Mackey. The Xerox Star: A retrospective. *Computer*, 22(9):11–26, 1989.
19. Bruce W. Knight, Jun-Ichi Toyoda, and Frederick A. Dodge. A quantitative description of the dynamics of excitation and inhibition in the eye of *Limulus*. *The Journal of General Physiology*, 56(4):421–437, 1970.
20. H. T. Kung, B. Sproull, and G. Steele. *VLSI Systems and Computations*. Springer, 1981.
21. J. Lazzaro, S. Ryckebusch, M. A. Mahowald, and C. A. Mead. Winner-take-all networks of $O(n)$ complexity. In *Neural Information Processing Systems*, volume 1, 1989.
22. R. F. Lyon. Machine hearing: An emerging field [Exploratory DSP column]. *Signal Processing Magazine, IEEE*, 27(5):131–139, 2010.
23. R. F. Lyon and C. Mead. An analog electronic cochlea. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(7):1119–1134, 1988.
24. Richard F. Lyon. The optical mouse and an architectural methodology for smart digital sensors. Technical report, Xerox PARC VLSI-81-1, 1981.
25. Richard F. Lyon. Simplified design rules for VLSI layouts. *VLSI Design*, 2(1):54–59, 1981.
26. Richard F. Lyon and Martin P. Haerberli. Designing and testing the optical mouse. *VLSI Design*, 3:20–30, 1982.
27. Misha Mahowald. *VLSI Analogs of Neuronal Visual Processing: A Synthesis of Form and Function*. PhD thesis, California Institute of Technology, 1992.
28. John Markoff. In focus: The mouse that rolled. *Info World*, 5(8):26–29, 1983.
29. Carver Mead. *Analog VLSI and Neural Systems*. Addison-Wesley, 1989.
30. Carver Mead and Lynn Conway. *Introduction to VLSI Systems*. Addison-Wesley, 1980.
31. Alireza Moini. *Vision Chips*. Springer, 2000.
32. Odelia Schwartz and Eero P. Simoncelli. Natural signal statistics and sensory gain control. *Nature Neuroscience*, 4(8):819–825, 2001.
33. Charles L. Seitz. System timing. In C. A. Mead and L. A. Conway, editors, *Introduction to VLSI Systems*, pages 218–262. Addison-Wesley, 1980.
34. R. M. Shapley and J. D. Victor. The effect of contrast on the transfer properties of cat retinal ganglion cells. *The Journal of Physiology*, 285(1):275–298, 1978.
35. John Tanner and Carver Mead. An integrated analog optical motion sensor. *VLSI Signal Processing*, 2:59, 1986.
36. Georg von Békésy. *Sensory Inhibition*. Princeton University Press, 1967.
37. L. Watts, R. F. Lyon, and C. Mead. A bidirectional analog VLSI cochlear model. In C. Sequin, editor, *Advanced Research in VLSI*, pages 153–163. MIT Press, 1991.
38. Agilent Technologies web site. *Company History Timeline – 1999*. <http://www.agilent.com/about/companyinfo/history/timeline.1999.html>, 2013.