N. A. Patrin

Performance of Very High Density Charge Coupled Devices

Abstract: The results of an experimental study of the performance of very high density 8, 128, and 256 bit CCD shift register structures are presented. The primary topics discussed include transfer efficiency, frequency response, impact of "fat zero" operation, and observed temperature dependency.

Introduction

If any particular aspect of the original charge-coupled device (CCD) [1,2] can be said to be responsible for the considerable current interest in this device, it is probably the potential for achieving very high density information storage with present day silicon process technology. The basis of this potential stems from the simplicity of the device structure, which involves the simple linear chain of coupled MIS capacitors depicted in Fig. 1.

Associated with, and in fact stemming from the structural simplicity of the device, is operational simplicity. The sequence of events involved in the operation of the device can be traced conveniently by assuming the substrate to be n-Si (p-channel) and considering the three-phase electrical timing sequence shown in Fig. 1. At some time t_1 when all of the ϕ_1 electrodes in the chain are at a negative potential $-V$, a localized depletion region will exist under each of these electrodes, or a potential energy "well" will be established at the surface of the Si under these electrodes. If the time duration during which the ϕ_1 electrodes are held at $-V$ is short compared with the time required for thermally generated minority charge (holes) to invert the n-Si surface, the regions under the energized electrodes will remain in full depletion and the potential wells will remain empty. If a packet of minority charges, $Q \approx C_{ox}(V - |V_T|)$, where V_T is the MOS threshold voltage, is introduced in the vicinity of the first ϕ_1 potential well, this charge will be captured by the well and held at the surface directly under the ϕ_1 electrode. When $V_{\phi 1}$ is reduced to ground, this charge packet is released. For the electrical timing sequence shown in Fig. 1, $V_{\phi 2}$ is at $-V$ prior to the time that $V_{\phi 1}$ goes to ground, while $V_{\phi 3}$ is still at ground. Thus the packet of charge released by the ϕ_1 well is captured by the right-adjacent ϕ_2 well, provided that the two potential wells are coupled during the phase transition. If the necessary coupling between adjacent MIS capacitors is achieved, a packet of charge introduced at the start of the chain is made to shift directionally along the chain. Since the achievement of directionality requires that one adjacent potential well be collapsed while the other is fully formed and empty, each packet of charge, or each bit of information, requires a minimum of three electrodes in this simple structure.

For a structure defined by, for example, 0.2-mil line widths, line spacings and channel widths, the channel area per bit is of the order of only 0.25 mil^2. For 0.2–0.4 mil spacing between adjacent channels, the average area per bit is two to three times this figure, providing an average array bit density of one to two million bits per

Figure 1 Schematic representation of a portion of a three-phase CCD and associated phase timing.

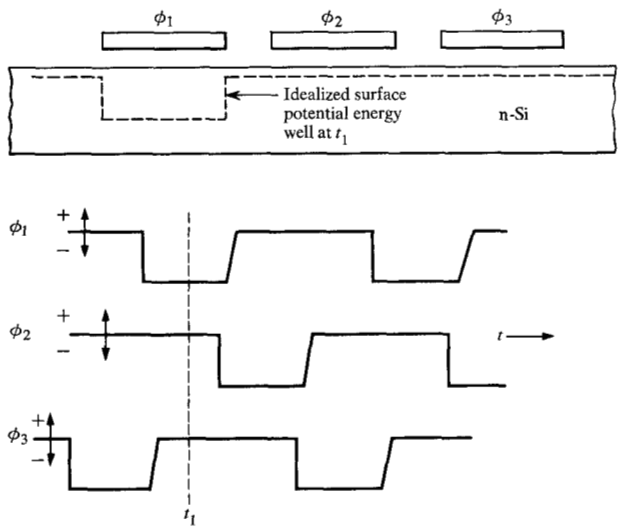

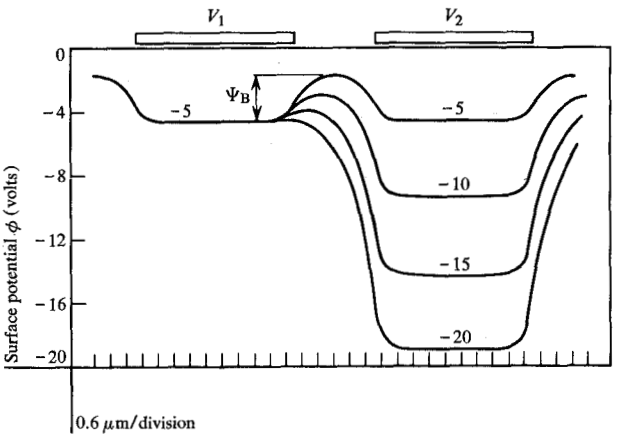

Figure 2 Surface potential variation in a plane bisecting a linear chain of 6 × 6 μm capacitors with 3-μm edge separation. Substrate is 7.2 ohm-cm, n-Si.

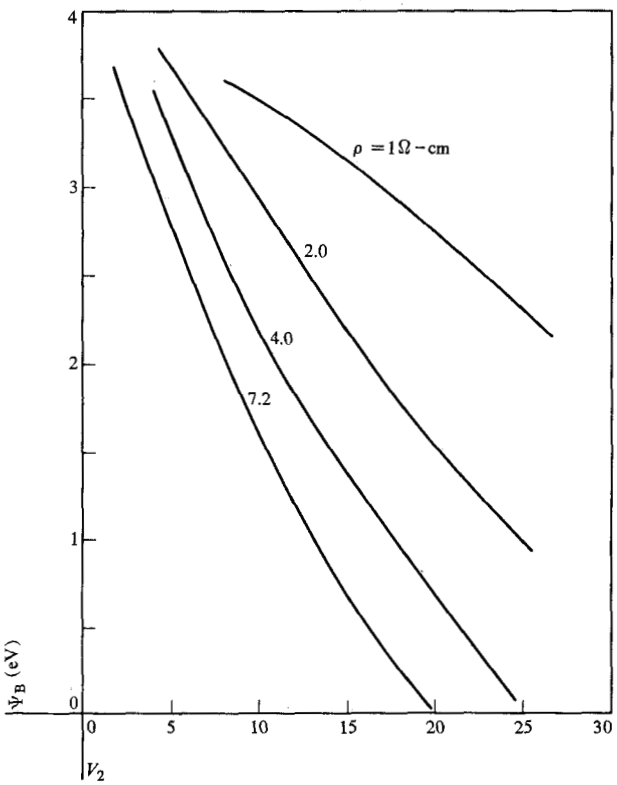

Figure 3 Dependence of the potential energy barrier between adjacent 6 × 6 μm MIS capacitors on the potential of one electrode (V_2) when the other electrode is at −5 V (V_1). Electrode edge separation is 3 μm. Substrate is n-type Si.

square inch. It is the achievement of such high density with present day processing and photolithography, and the achievement of much higher densities by means, for example, of electron beam line definition, that promises to make the CCD a potential near-term embodiment of large scale integration. In this paper we present some of the experimental data that have been obtained from a study undertaken to assess the performance of such very high-density devices.

Achieving coupling between adjacent potential wells

The proviso required for transfer of charge packets along the MIS chain, whatever the bit density, is that sufficient coupling between adjacent elements can be established. For the finitely separated electrode structure shown in Fig. 1, the achievement of coupling depends on the extent of potential-energy fringing in the vicinity of the interelectrode gap. It is expected that this fringing will be enhanced by increasing electrode potentials and by decreasing the substrate doping levels. A computer analysis was carried out to quantify the dependence of fringing, and the achievable coupling, on substrate resistivity, electrode potentials, and the gap size in a three-dimensional linear MIS array [3]. Representative results of this analysis are shown in Figs. 2 and 3.

The data given in Fig. 2 depict the surface potential variation in a plane that bisects an MIS chain comprised of 6 μm × 6 μm electrodes and 3 μm edge spacing. The substrate (n-Si) resistivity is 7.2 ohm-cm and the effective oxide thickness of the insulator is 330 Å. The potential variation shown is for the case where V_1 is held fixed at −5 V (−4.7 eV potential energy well) while V_2 is varied between −5 and −20 V (surface potential varies between −4.7 and −19.1), with all wells being empty. A similar surface potential profile is obtained for the case where V_1 is at, e.g., −20 V and the corresponding potential energy well contains sufficient positive charge ($Q \cong 14.4\ C_{ox}$) to bring the surface potential to −4.7 V. When V_2 is at −5 V, a value that may be considered to be due to an applied substrate bias, a potential energy barrier ψ_B of approximately 2.9 eV decouples the two adjacent wells. The value of ψ_B decreases monotonically as V_2 is made more negative until, at a V_2 of approximately −19 V, ψ_B disappears (threshold coupling). If V_2 were held at −20 V, some of the charge in the V_1 well would transfer to the V_2 well. This would lead to a reduction of the V_2 surface potential from −19.1 V toward ground and an increase of the V_1 surface potential from −4.7 V to more negative values, with the result that a finite ψ_B would again appear, cutting off the transfer. To transfer the balance of the charge held by the V_1 well would now require the reduction of the V_1 potential at a rate not exceeding the ability of the charge to flow into the V_2 well. It should be noted that the condition that must be satisfied to transfer the full V_1 well charge completely into the V_2 well is that, at the end of this transfer, the surface potentials under the two electrodes must be equal and that $\psi_B = 0$. The fact that such a sur-

face potential condition is not possible for finite gaps simply means that a completely full well cannot be propagated along the chain, the gap effect thus imposing a "skinny one" operation mode. If a given well in the chain did indeed contain a "full one," then during the first transfer operation, part of the charge would transfer properly, part could spill backwards into the preceding well, part could be lost to the substrate, and part of the charge could be left behind (finite substrate bias).

The dependence of the potential energy barrier between potential energy wells on electrode potentials and substrate resistivity for a given gap of 3 μm is indicated in Fig. 3. The indicated decrease of ψ_B and the associated enhancement of coupling with increasing substrate resistivity and increasing "downstream" electrode potential merely supports the qualitative expectations cited earlier. Since the use of large resistivities and large electrode potentials attacks the gap effect limitation by depending on an increase in longitudinal potential energy fringing, this approach also leads to an increase in transverse fringing and thereby enhances the probability of encountering crosstalk between adjacent channels. The avoidance of crosstalk forces the use of either an isolation diffusion or a larger separation between adjacent channels, with a consequent sacrifice in bit density.

The potential energy profiles discussed above were derived under an assumption of zero oxide charge. If finite (positive) oxide charge is introduced into the model, the gap-effect decoupling is aggravated in p-channel devices, and alleviated in n-channel devices. In fact, intentional oxide doping has been suggested [4] for n-channel structures. Although the n-channel oxide doping approach may be generally viable, the inhomogeneity of this doping, when coupled with modest variations in interelectrode gaps across a large array, could lead to a significant variation in threshold coupling across the array and thereby to significant variations in performance throughout the array.

A straightforward alternative to the gap-effect solutions cited above is to simply eliminate the gap by using a slightly modified Si gate process. This was the approach taken in the design and fabrication of the pSAG CCD structures used in our study, as well as in the design and fabrication of the dual 480-bit arrays described by Vogl and Harroun [5]. The virtual elimination of gaps allowed the use of the 1-2 ohm-cm substrates nominally preferred for support-circuit design, led to successful device operation with phase line potentials as low as -3 to -5 V, and also led to prevention of crosstalk between 0.2-mil spaced channels without resorting to isolation diffusion.

Device structure

A cross-sectional view of the overlapped electrode devices used in this study is presented in Fig. 4. Although

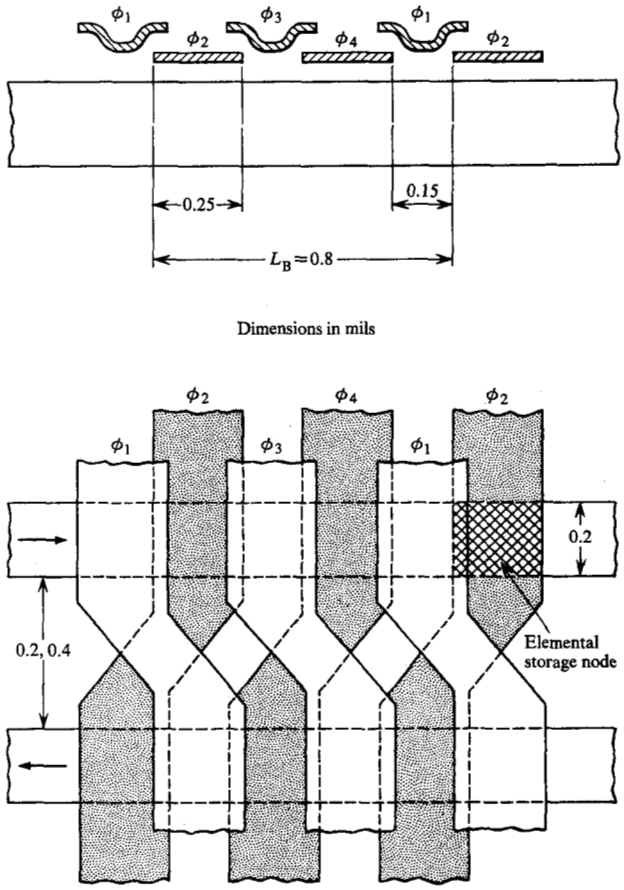

Figure 4 Schematic representation of the horizontal and vertical structure of the very high density devices tested.

overlapped electrode structures can be designed to be operable in a three-phase mode, the natural driving sequence is either two- or four-phase. In the structure shown, each spatial bit has associated with it four independent electrodes. The ϕ_2, ϕ_4 doped polysilicon electrodes define the storage well (node) locations. The ϕ_1, ϕ_3 aluminum electrodes serve as transfer/isolation gates between the storage nodes. By using the ϕ_1, ϕ_3 electrodes merely as gates, and not requiring storage under those elements, the Si surface area under the aluminum electrodes can be smaller than the surface area under the ϕ_2, ϕ_4 electrodes, thereby reducing bit size. At the same time the potentials applied to ϕ_1, ϕ_3 can generally be smaller than those applied to ϕ_2, ϕ_4, thereby reducing driving power as compared with conventional four-phase operation.

The quasi-four-phase timing of the pulses is shown in Fig. 5. Consider some time t_0 when charge is stored under a ϕ_4 electrode, and this charged well is isolated by the grounded adjacent ϕ_1, ϕ_3 electrodes. Since information is stored only in every other storage node, the ϕ_2 electrodes can be grounded at t_0. By energizing ϕ_1 and ϕ_2 at

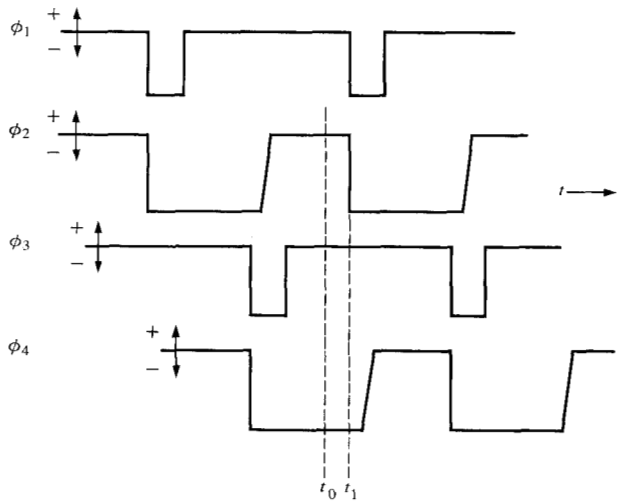

Figure 5 Four-phase timing used in device operation.

t_1, while ϕ_3 is kept grounded, positive coupling between ϕ_4 and the right-adjacent ϕ_2 is achieved. By collapsing the ϕ_4 potential to ground at t_1^+ the charge can be made to complete its transfer from the ϕ_4 to the ϕ_2 node. The ϕ_1 potential can now be brought to ground, thereby completely isolating the ϕ_2 node. In the half-cycle described, information has been transferred through one half a spatial bit. In the next half-cycle, information is transferred through the ϕ_3 channel and onto the next ϕ_4 node.

By driving the ϕ_1 and ϕ_2 lines with a ϕ_1' driver, and driving ϕ_3 and ϕ_4 with a ϕ_2' driver, the devices can be made to transfer charge in a two-phase mode. The timing of ϕ_1' and ϕ_2' is identical to the timing of ϕ_2 and ϕ_4 shown in Fig. 5. To achieve charge flow directionality in this mode, the surface potential under the even electrodes (polysilicon) must always be greater than the surface potential under the associated odd electrode (aluminum). A simple way of achieving such surface-potential asymmetry under equally driven electrodes is to build in a thicker insulator under the Al electrodes. The obvious advantage of the two-phase driving mode lies in the reduction of the required number of drivers. The basic disadvantage of the simplistic two-phase structure described here is associated with the fact that the storage nodes are not allowed to fill to a level above the surface potential level of the associated transfer element. This leads to practical signal levels that are approximately 25 percent of the levels in comparable four-phase operation. Although the devices studied have successfully been operated in a two-phase mode, the data reported here are limited to four-phase operation.

Device size

From the device dimensions given in Fig. 4 it follows that the area of an elemental storage node is 0.05 mil^2. The total channel area per bit is 0.16 mil^2. In the devices described here, the spacing between adjacent channels is 0.4 mil to allow the crossover of phase lines between channels while still satisfying a 0.15-mil linewidth ground-rule. The crossover was used to provide opposite directivity in adjacent channels. With such relatively large spacings the average area per bit is 0.48 mil^2, giving an array bit density of 2×10^6 bits/in^2. If directivity alternation between channels is not needed, the spacing can be reduced to 0.2 mil, while still providing adequate isolation between channels in the 1-2 ohm-cm substrates used, and thereby decreasing the average area per bit to 0.32 mil^2. Shift-register configurations of 128 and 256-bit chains of these elements have been experimentally built and successfully operated.

Transfer efficiency

The transfer of a charge packet from one potential energy well into another in a surface CCD is produced by a combination of thermal diffusion and drift due to charge gradients and fringing fields, while it is hindered by trapping at interface states with the result that complete transfer, in principle, is never realized in finite transfer times. If the amount of charge actually transferred from the i'th potential well into the adjacent $i+1$'th well, $Q'(i)$, is related to the charge $Q(i)$ in the i'th well prior to transfer by

$$Q'(i) = \eta_i Q(i) ,$$

then η_i defines the efficiency of this transfer operation [1]. Aside from being a function of physical parameters such as surface mobility, channel length and time constant(s) associated with interface trapping, η_i is also a function of the charge levels involved in transfer.

If a large charge packet, $Q_1(0)$, denoting a "1" bit, is launched into a device in which thermal charge generation and charge loss to the substrate are negligible, the overall transfer efficiency after transfer through N_B bits would be given by

$$\eta_0(N_B) = \eta_1\eta_2\eta_3 \cdots \eta_{\theta N_B} < 1 ,$$

with θ denoting the number of transfers per bit. If $Q_1(0)$ were to follow a sufficiently long series of bits containing zero charge, so that it would not have any charge to pick up, then the charge $Q_1(\theta N_B)$ representing the "1" after θN_B transfers would be

$$Q_1(\theta N_B) = \eta_0(N_B) Q_1(0)$$

and the lost or lagging charge would be dispersed among the succeeding bits. If a second large charge packet $Q_2(0) = Q_1(0)$ is launched immediately after $Q_1(0)$, then at least some of the charge left behind by Q_1 could be picked up by Q_2. Thus $Q_2(\theta N_B)$ could be larger than $Q_1(\theta N_B)$ and in fact could be equal to $Q_2(0)$.

Such optimization of $Q(\theta N_B)$ caused by the pickup of charge left behind by preceding bits is the basis for using a "fat zero", or bias charge, mode of CCD operation. In the fat zero mode (FZM), a finite amount of charge $q(0)$ is launched for every zero bit [6]. If η_i' is used to denote the transfer efficiency of $q(i)$ from the i'th to the $i+1$'th potential well, and if n_i' and $q(i)$ are such as to satisfy

$$\eta_i' < \eta_i ;$$

$$(1-\eta_i')q(i) \approx (1-\eta_i)Q(i) ,$$

then the charge left behind by a $Q(i)$ during the i'th transfer would be compensated at least partially by the charge picked up in the $i+1$ well even if Q were to follow a long series of "zeros." Thus the effective transfer efficiencies per transfer, and thereby the overall transfer efficiency, would be increased. If $\eta_i' \ll \eta_i$ then the value of bias charge needed to provide the efficiency improvement could be very small compared to Q.

For an input pattern of (1110011) imbedded in a sea of zeros, the observed output patterns for typical 8, 128, and 256 bit unrefreshed devices operated at 100 kHz are shown in Fig. 6. Although these devices were operated in FZM, the amount of bias charge used was so small as to be imperceptible in affecting the amplitude of an output zero. The bias charge was, however, perceptible in its effect on the output amplitude of the first one, increasing this amplitude by an amount between two and five percent in the long devices. Qualitatively, results such as these indicate that (a) a steady state one (1_{ss}) or zero (0_{ss}) is achieved on essentially the second consecutive like bit; (b) the difference between the value of a 1_{ss} and a worst-case one (1_{wc}: first one after a sufficiently long series of zeros) is essentially equal to the difference between the worst case zero (0_{wc}: first zero after a long sequence of ones) and 0_{ss}; and (c) the 1_{ss} signal in a 256-bit device is equal to the 1_{ss} signal in a 128 and an 8-bit device.

These observations combine to imply extremely low dispersion for the long structures, as well as charge conservation. No discernible amount of charge is lost to (or contributed by) the substrate. These qualitative observations also imply extremely high transfer efficiencies.

The actual FZM transfer efficiencies typically observed under 100 kHz clocking are listed in Table 1. The overall transfer efficiency was calculated from

$$\eta_0 = 1_{wc}/1_{ss}$$

and the average transfer efficiencies per bit, η_B, and per transfer, η_t, were derived from η_0 via

$$\eta_B^{N_B} = \eta_0$$

$$\eta_t^{\theta N_B} = \eta_0 ,$$

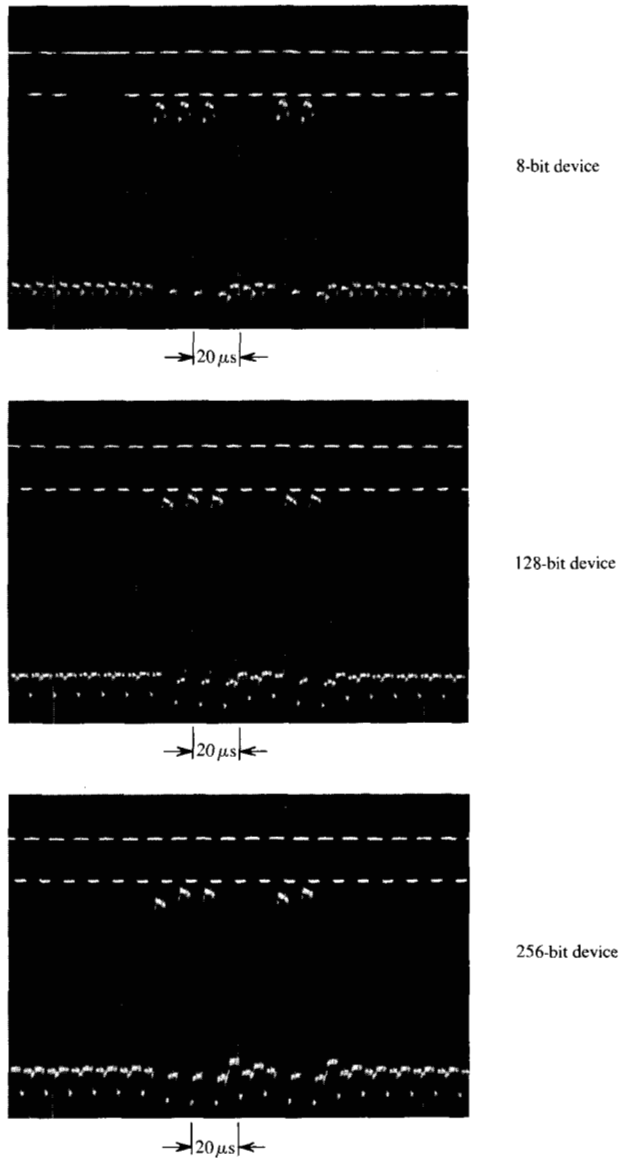

Figure 6 Oscilloscope traces showing the 100 kHz output from 8, 128, 256-bit very high density devices under common operating conditions. Vertical/Horizontal scaling = 400 mV/20 μsec. per division. Output is measured at the source-follower node of an on-chip integrating amplifier.

Table 1 Typical transfer efficiencies for tested devices.

	Number of bits N_B		
	128*	128	256
Overall transfer efficiency, η_0 (%)	98	95	90
Transfer efficiency per bit η_b (%)	–	99.96	99.96
Transfer efficiency per transfer η_t (%)	99.99⁺	99.98	99.98

*Best available devices

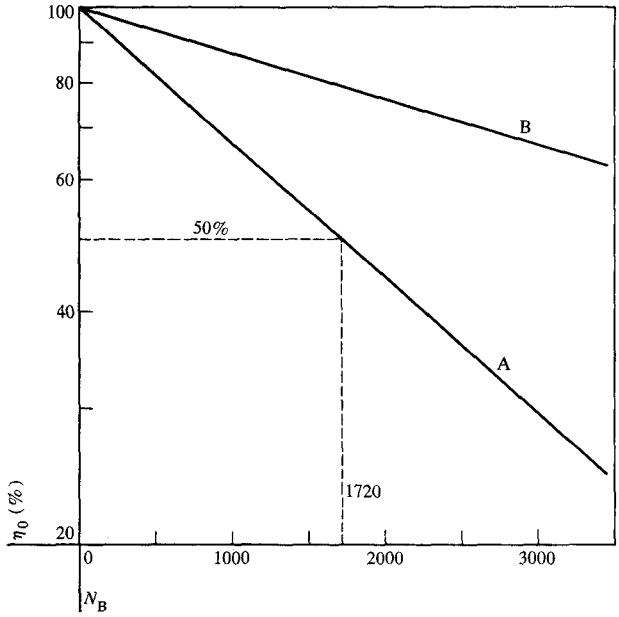

Figure 7 Expected variation of overall transfer efficiency (η_0) as a function of the unrefreshed number of bits (N_B). Line A is based on the nominal transfer efficiences listed in Table 1. Line B is based on transfer efficiencies observed in the best devices.

with $\theta = 2$ characterizing these devices. By extrapolating the nominally observed values of η_0 (e.g., 95 percent over 128 bits) to a larger number of bits (Fig. 7) it is expected that 1024-bit unrefreshed devices would be characterized by an overall transfer efficiency of 67 percent. The best efficiencies observed (98 percent over 128 bits) imply the attainment of 70-75 percent overall efficiency for 2048-bit devices.

The practical implication of achieving very large values of η_t and η_B is of course dependent on the device application being considered. In a digital memory application for which access time minimization is an important consideration, the use of long shift registers is generally undesirable, and the ability to fabricate 1024- or 2048-bit unrefreshed devices may seem academic. However, the kind of short register signal purity that the efficiencies in Table 1 imply would allow for the use of simple, compact, sense-refresh amplifiers after, e.g., 32-64 bits without significantly sacrificing average bit density. Furthermore, if the level of performance indicated by the efficiencies in the table is substantially higher than may be required in a particular application, one can reduce the phase voltages from those used here (ϕ_1, $\phi_3 = -12$ V; ϕ_2, $\phi_4 = -15$ V) and thereby obtain an associated reduction in drive power requirements.

Frequency response

An important performance characteristic that can determine the application of surface CCD structures is the clock frequency range over which these devices can be operated. Whereas it is expected that the low frequency limit would generally be imposed by thermal charge generation rates and a consequent tendency of empty wells to fill with minority charge, the high frequency limit is expected to be imposed by those physical device parameters that determine the minimum transfer time windows that have to be used if gross transfer inefficiency is to be avoided. A convenient way of describing the frequency response of a CCD involves the specification of the worst-case signal discrimination under digital operation.

The observed room temperature dependence of the normalized $(1_{wc} - 0_{wc})$ difference on frequency for a typical unrefreshed 128-bit device is presented in Fig. 8 with the normalization being made with respect to the value of a 1_{ss} at 500 kHz:

$$D_{wcn}(f) \equiv \frac{1_{wc}(f) - 0_{wc}(f)}{1_{ss}(500\ \text{kHz})}.$$

The relative constancy of D_{wcn} for frequencies between approximately 1 kHz and 1 MHz is associated with the fact that 100-kHz transfer efficiencies given in Table 1 were found to characterize device performance over a broad frequency range. That such an association is at least approximately valid can be verified from the observation that 0_{ss} was essentially zero and that 0_{wc} was essentially equal to $(1_{ss} - 1_{wc})$ in the range of frequencies where D_{wcn} was relatively constant. Under these conditions we have

$$D_{wcn} \approx 2\eta_0 - 1\,.$$

The decrease of D_{wcn} for frequencies below approximately 1 kHz is associated with the growth of 0_{ss}, and thereby to the growth of 0_{wc}, due to the pickup of thermally generated minority charge during the time $\tau = N_B/f$ that the zero bits travel through the register. The 500 Hz data point in Fig. 8 indicates that such room temperature well filling is fairly modest for $\tau = 0.25$ sec. Similar data on 256-bit devices indicate that 1-sec storage/delay times are nominally realizable at room temperature. The effect of temperature on device operation is discussed in the next section.

The decrease in the D_{wcn} response indicated by curves FZM and SZM for frequencies greater than 1 MHz is associated with the fact that the times allowed for transfer are no longer sufficient to maintain the transfer efficiencies given in Table 1. In fact, at a clock frequency of 4 MHz, the allowed 60-70 ns transfer time window leads to an overall 128-bit transfer efficiency of approximately 50 percent under SZM operation and a consequent loss of worst-case signal discrimination. The introduction of a finite bias charge that is only barely perceptible in its effect on the value of 0_{ss}, however,

raises D_{wcn} to approximately 0.77 at 4 MHz and apparently extends the usable response to 10 and possibly 20 MHz.

Temperature dependence

At the high operating frequencies where room temperature performance becomes transfer time limited, the effect of temperature on CCD operation is dependent on the dominant incomplete transfer mechanism. For the case where drift-aiding fringing and interface trapping are negligible, the reported proportionality between Si surface mobility and inverse temperature [7] renders the diffusion coefficient temperature independent and thereby has been shown to cause transfer time limited operation to be somewhat insensitive to temperature variation [8]. For another case in which device design is such that drift-aiding fringing plays a dominant role in the determination of required minimal transfer times, while interface trapping is still negligible (e.g., buried-channel CCD [9]), the decrease of mobility with increasing temperature may degrade high-frequency transfer by effectively reducing the importance of drift-aiding fringing. For yet another case in which minimum transfer times are dicated by interface trapping, the acceleration of detrapping with increasing temperature may obviate the need for fat zero mode operation for performance optimization at high frequencies.

At clock frequencies substantially lower than those at which transfer time limited performance degradation occurs, the effect of temperature on device operation has been mentioned in the previous section to involve the filling of zero bits with thermally generated charge during the time $\tau = N_B/f$ that these bits propagate through the structure. The observed variation of the normalized $(1_{wc} - 0_{wc})$ difference with temperature shown in Fig. 9 clearly displays this low-frequency dependency. By progressively increasing the clock frequency from 1 kHz to 100 kHz, and thereby decreasing τ from 128 to 1.28 ms, the 128-bit device can be operated to progressively higher temperatures. In the temperature range 20 to 120 °C, the critical temperature, T_c, at which worst case signal discrimination is lost can be seen from the figure to be related to τ via

$$T_c(\tau_1) \approx T_c(\tau_2) + 42 \log_{10}\left(\frac{\tau_2}{\tau_1}\right) \quad (°C),$$

with

$$\tau_i = N_{B_i}/f_i,$$

and with N_{B_i} denoting the bit capacity of a register being clocked at a frequency f_i. For the 8V substrate bias used in obtaining the data of Fig. 9,

$$T_c(\tau = 0.128 \text{ sec}) \approx 42 °C.$$

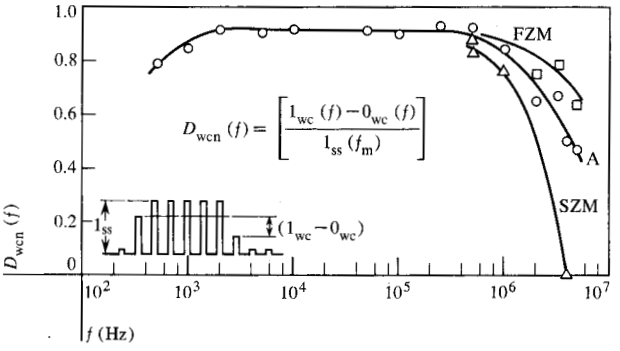

Figure 8 Room temperature frequency response of a very high density, unrefreshed 128-bit device. Curve A represents the overall response under "fat zero" operation; curve FZM denotes the response of the CCD only, with the sense amplifier response factored out; curve SZM represents overall device response with "skinny zero" operation. ϕ_1, $\phi_3 = -12$ V; ϕ_2, $\phi_4 = -15$ V; substrate bias, 3 V.

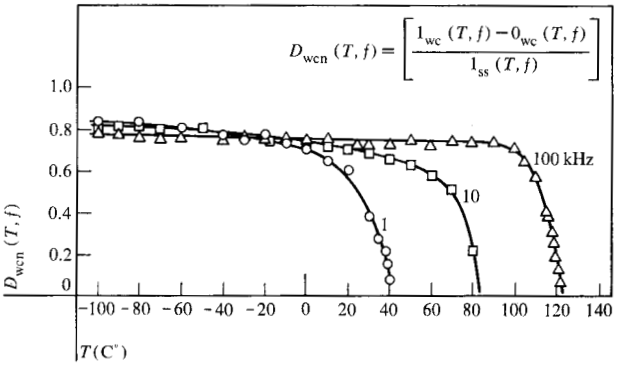

Figure 9 Temperature response of a very high density 128-bit device at three clock frequencies. ϕ_1, $\phi_3 = -7$ V; ϕ_2, $\phi_4 = -15$ V; substrate bias, 8 V.

By decreasing the substrate bias toward nominally used values of 2 to 3 V, the curves shift toward higher temperatures. From a more detailed analysis of the data in Fig. 9 it can be shown that the average "leakage current," or average zero-bit charge buildup rate, follows the temperature dependency of the intrinsic carrier concentration n_i, and that the empirical expression for T_c given above provides a first-order approximation to this dependency in the stated temperature range. If we combine the observed proportionality between zero-bit growth and n_i with the observed dependency on substrate bias, it appears that the dominant leakage mechanism involves charge generation within the field-induced depletion region.

Aside from exhibiting the general temperature dependency of the CCD structures studied, the data of Fig. 9 can be used in conjunction with the data of Fig. 8 to describe the dependence of performance on ϕ_1, ϕ_3 voltage excursions. Whereas the same ϕ_2, ϕ_4 voltage

amplitudes and the same 1_{ss} levels were used in these two experiments, the low-temperature D_{wcn} values in Fig. 9 are smaller by approximately 0.1 than the corresponding D_{wcn} values in Fig. 8. Although a portion of this difference is associated with the different values of substrate bias and the use of a skinny zero vs a fat zero, it is also related to the different ϕ_1, ϕ_3 amplitudes used in the two experiments. By lowering the substrate bias to 3 V and by launching fat zeros, the low-temperature 100-kHz curve in Fig. 9 moves upwards to D_{wcn} values of approximately 0.86. By increasing ϕ_1, ϕ_3 amplitudes to -10 V, D_{wcn} is further increased to approximately 0.9. Further increases in ϕ_1, ϕ_3 have little perceptible effect for clock frequencies below 1 MHz.

Summary

The very high density p-channel CCD structures described here have been found to be characterized by transfer efficiencies that are among the highest observed to date for surface devices. On the basis of observed performance at clock frequencies up to 5 MHz, it appears that the 128- and 256-bit structures can provide usable worst case digital signal discrimination up to 10 to 20 MHz, and to higher frequencies for shorter registers. Limitations associated with the particular on-chip charge-integrating sense amplifier used here prevented testing at these higher frequencies.

The temperature-limited low frequency operation was found to be characterized by nominally attainable 1 sec input-output delay times at room temperature (e.g., 256 bits at clock frequencies under 500 Hz), with this attainable delay being inversely proportional to the temperature variation of the intrinsic carrier concentration.

In the devices tested the fat zero or bias charge mode of operation was found to be relatively insignificant in improving performance at clock frequencies below 1 MHz, and essential in extending the operation of long devices beyond 4 to 5 MHz.

The excellent digital signal discrimination that has been observed over wide frequency and/or wide temperature ranges allows for the use of simple, compact sense-regeneration amplifiers that have negligible impact on overall density. An exemplary array, employing refresh after every 32 bits, would provide average access times as low as 16 μs when clocked at 1 MHz. Although individual 32-bit devices would be expected to be functional up to 20 MHz, and thereby be characterized by sub-microsecond average access times, the 100^+ watt/in^2 power densities required to drive these devices at such high frequencies [10] renders the exploitation of the performance potential of these devices somewhat impractical in large arrays. For situations in which 20^+ μs access and 2 to 3 MB/in^2 densities are acceptable tradeoffs, however, the type of surface CCD structures described here appear to present a means of attaining, e.g., 32 kBit storage, including supports, on nominal 150-mil chips.

Acknowledgments

It is a pleasure to acknowledge J. J. Chang and A. Lo for originally suggesting the basic device structure; M. L. Joshi and the members of his staff for excellent processing support; J. J. Chang and B. Agusta for their support of this work; and L. G. Heller and H. S. Lee for their participation in many helpful discussions pertaining to this study.

References

1. W. S. Boyle and G. E. Smith, *Bell System Tech. J.* **49,** 587 (1970).
2. G. F. Amelio, M. F. Tompsett and G. E. Smith, *Bell System Tech. J.* **49,** 593 (1970).
3. N. A. Patrin, "Three Dimensional Potential Energy Variation in a Linear MIS Array," presented at the IEEE Device Research Conference, Ann Arbor, Michigan, June 1971.
4. R. H. Krambeck, *Bell System Tech. J.* **50,** 3169 (1971).
5. N. G. Vogl and T. V. Harroun, *1972 IEEE International Solid-State Circuits Conference Digest,* Lewis Winner, New York 1972, p. 246.
6. C. N. Berglund and H. J. Boll, *1970 IEEE International Electron Devices Meeting, Digest p. 16*; W. E. Engler, J. J. Tieman and R. D. Baertsch, *1971 IEEE International Solid-State Circuits Conference Digest,* Lewis Winner, New York 1971, p. 164.
7. L. Vadasz and A. S. Grove, *IEEE Trans. Electron Devices* **ED-13,** 863 (1966).
8. H. S. Lee and L. G. Heller, *IEEE Trans. Electron Devices* **ED-19,** 1270 (1972).
9. R. H. Walden, R. H. Krambeck, R. J. Strain, J. McKenna, N. L. Shryer, and G. E. Smith, *Bell System Tech. J.* **51,** 1635 (1972); C. K. Kim, J. M. Early and G. F. Amelio, *IEEE NEREM Record* **14,** 161 (1972).
10. N. A. Patrin, A. Bhattacharyya, M. L. Joshi and J. J. Chang, *IEEE NEREM Record* **14,** 157 (1972).

Received October 20, 1972

The author is located at the IBM System Products Division Burlington Laboratory, Essex Junction, Vermont 05452.