

XEROX
BUSINESS SYSTEMS
Systems Development Department

To: Distribution

Date: October 8, 1979

From: Ron Crane

Org: SDD/SDT

Subject: Dandelion Memory System

Filed: [IRIS]<Workstation>MCTL>MC.TLmemo.press

Dist.:WS Design group

This memo covers several aspects of the memory system for the Dandelion processor. The memory system reads and writes 16 bit words for both the processor and the display. The low bank is shared between the two. Error correction is performed on all words delivered to the processor. The memory cycle time is 411 nanoseconds (nS). The low 64K words are located on the memory control card, with up to 128K words located on the storage card. If 64K memory chips are used instead of 16K chips, these numbers can be multiplied by 4.

Memory Functions

This section provides a description of each of the functions of the memory system as viewed from the processor. Both a system level diagram and a detail block diagram of the memory system are on the following page.

Read

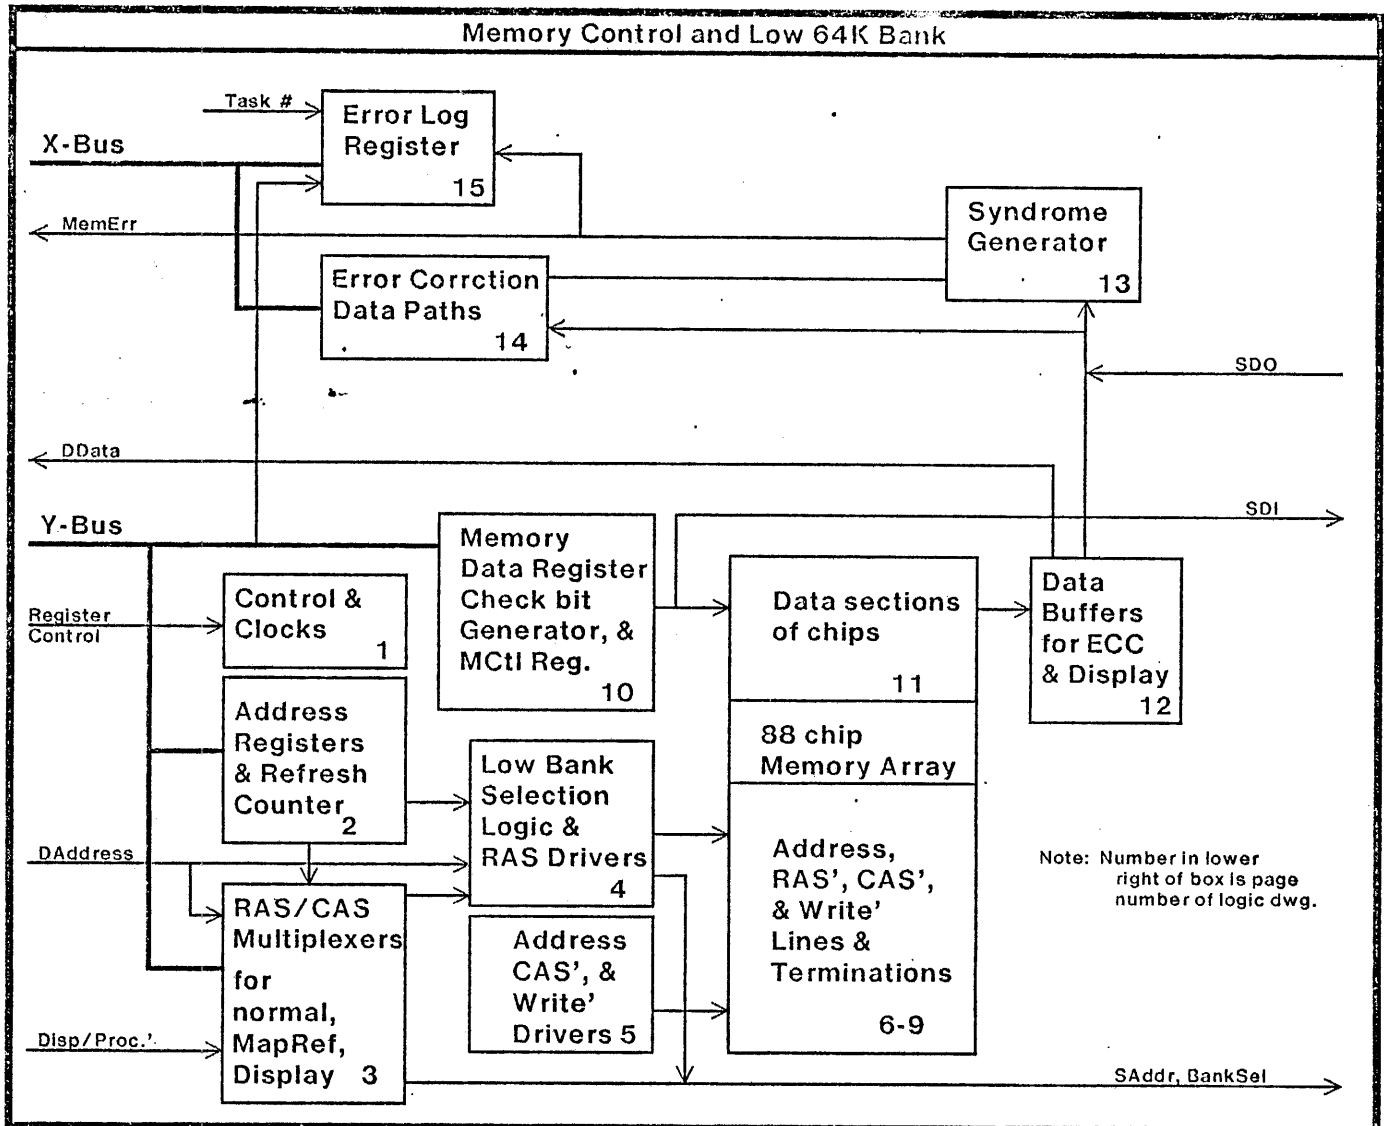
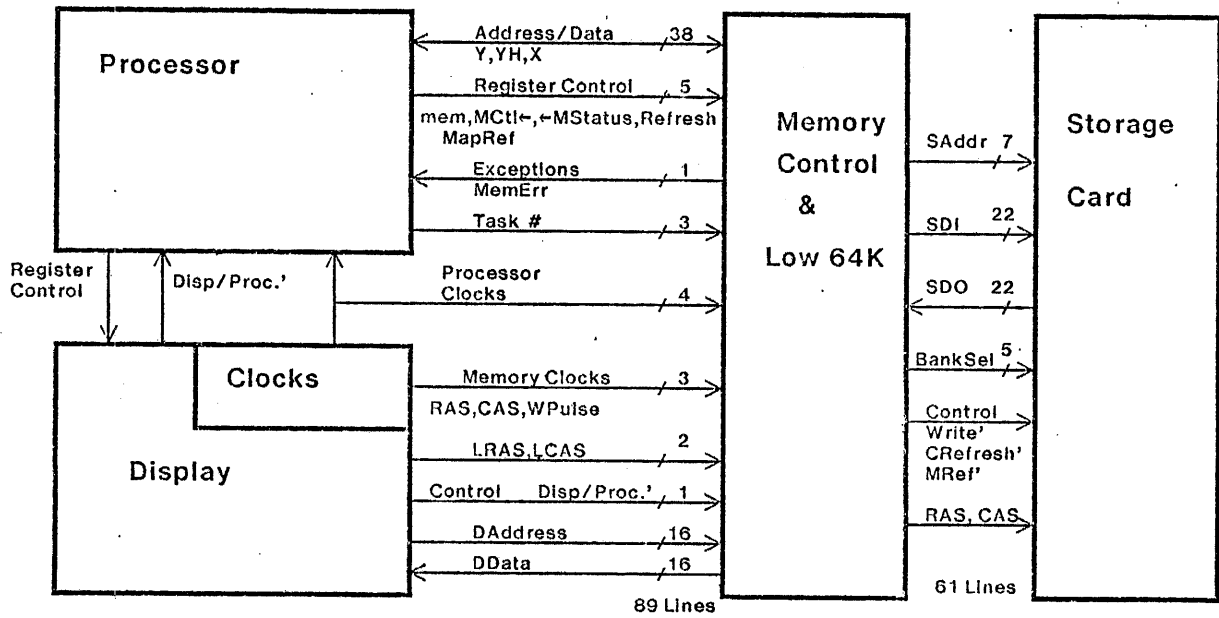
A read operation is started by placing the memory address on the Y bus and asserting 'mem' in the first cycle of a click. Bank selection is done by the 2 low order bits present on the YH bus and Y0,1, while the word selection within the chip is done by the low 14 bits on the Y bus. The data can be read back to the X bus at the end of the third cycle by asserting 'mem' during the third cycle of a click. The address must meet a bit dependent setup time, because the RAS signal actually latches the most significant 7 bits of the address in a bank during the first cycle of a click. All data read back is error corrected unless the correction inhibit bit is set in the MCTL register.

Write

A write operation starts just like a read operation with the address sent out in the first cycle of a click. The data to be stored must be delivered to the memory during the second cycle of a click, by asserting 'mem' in the second cycle, and placing the data on the Y bus. Error correction check bits are always calculated and stored automatically by the memory system.

Map Reference

A map reference memory read is just like a regular read, except that the 22 bits supplied by the Y and YH busses are shifted around to facilitate indexing into a page map, which in combination with microcode, provides a 22 bit virtual memory system. The low 8 bits supplied are discarded (since they are the word location on the page), and the high 14 bits (virtual page number) are moved to the low 14 bits used for the address into real memory. The high 4 bits are 0100, thus fixing the location of the 16K map between 65K and 80K in real memory. In microcode, MapRef should only be selected during cycle 1, and never during cycles 2 or 3, since nothing will happen. A diagram of this is included in the logic drawings.



Refresh

The memory controller contains circuitry to facilitate memory refresh. In particular, the counter is included on the board and all banks of memory are refreshed at the same time. The CAS (column address strobe) signal is suppressed during refresh, thus minimizing the current used in refresh. Refresh is initiated by asserting the refresh control signal from the processor during cycle 1 of a click. Refresh should not be asserted during cycle 2 or 3 of a click or when the display is using the low bank, since nothing will happen in these cases. While the 16K chips require only 128 refresh cycles every 2 mS, it is recommended that microcode implement 256 refresh cycles every 2 mS until it is clear that 64K memory chips can be delivered with 256 cycles every 4 mS.

Display Lockout

The low 64K of the memory is shared between the display and the processor. The display has priority. System timing is cyclic, with 5 clicks per round. When actually scanning a line, the display consumes clicks 1 through 4, leaving click 5 for the processor. Thus, only one click out of 5 is available for use by memory refresh, display handling, and cursor microcode. About half the bandwidth remains in the 5th click for emulator use after memory refresh, display, and cursor tasks have been subtracted.

Lockout occurs only if the display is outputting a line from memory and access to the low bank is attempted. Accesses to the high bank(s) are not affected. Lockout does not occur during retrace intervals (horizontal and vertical), or during any other period of display inactivity (such as partially or completely shutting off the display). By convention, time critical hardware tasks using the first 4 clicks must never attempt access to the low (display) memory bank since a lockout could occur causing extra delay. See the display controller description for exact details of display timing.

Lockout is implemented by generation of a wait signal in the processor whenever a bank 0 (low 64K bank) access is attempted and the display is already using the low bank. The processor suspends the microcode which started in that click, and continues as normal the arbitration of what runs in the next click. In this manner, lockout in one click does not hold up operation in the following click.

Error Correction

Since soft errors can occur in the memory (alpha particles from the package, etc.) error correction circuitry is included in the memory system. Six check bits added to the 16 bit word provide single error correction and double error detection (SEC-DED). No explicit indication of single errors is provided, although the status of any particular operation can be read from the Status & Errors register after an operation. Error correction can be disabled, and the check bit positions in memory selectively set by writing into the Mctl register and reading the MStatus register.


Double Error Logging

A double error signal is available and also latched on a per task basis in the MStatus register. Thus, a task, upon entering a critical data transfer phase, could clear its particular bit, perform the task, and then check to see if its bit was set (double error). If an error did occur, its effect would be limited to events in that interval, over which some corrective action might be taken.

General Comments on Memory Programming


Details of memory programming are contained in the next figure, describing the memory registers. As mentioned earlier, neither MapRef nor refresh should be asserted in cycles other than cycle 1. In addition, refresh should not be asserted whenever the display is using the low bank.

MAR← Memory Address Register




MAR← mem during first cycle of click.
 Action: Contents of YH(6,7),,Y(0,15) is used as memory address. Access is started.

MDR← Memory Data Register




MDR← mem during second cycle of click.
 Action: Contents of Y Bus go into memory location specified by contents of MAR as loaded during first cycle of click. No write occurs if the low 64K bank is selected and it is already being used by the display.

← MD Memory Data



←MD mem during third cycle of click.
 Action: Memory data to X-Bus is single error corrected if MCTL bit 15 is set. The status of a given read operation can be found by looking in MStatus before the next memory read (←MD) is done. The occurrence of both single and double errors are indicated here. This operation gives the contents of the memory cell specified during cycle 1, independent of whether a write was specified during cycle 2.

MapRef Map Reference




MapRef during cycle 1 of click. (not during cycle 2 or 3)
 Action: This action is the same as a MAR← except that the physical address is derived differently. An access is started in the 65K - 80K bank of memory. The location accessed is specified by the 14 bit page number.

Refresh

Refresh during cycle 1 of click. (not during cycle 2, 3, or when display is using low 64K bank)
 Action: A RAS only cycle is initiated in all memory chips. Row Address is supplied from an internal 7 bit counter which is incremented once per occurrence of refresh. DO NOT USE refresh if the display is using the low bank of memory during that cycle. No refresh will occur.

MCTL← Memory Control Register

MCTL← during any cycle.




Action: Normally this register is set to 0. A-F can be set to one to test syndrome bits and error indications. Individual bits of the error log can be cleared by setting bit 4 and using Pt0-2 to specify the bit to be cleared. Bit 15, inhibit correction, affects only the data being read. Check bits are always generated and stored in memory during writes.

← MStatus Memory Status

← MStatus during any cycle.

Action: This register is loaded every time memory data is read by the processor (←MD). High byte has status of most recent memory access. Low byte latches any occurrence of double error on a per task basis. Register is 0 if no errors logged.



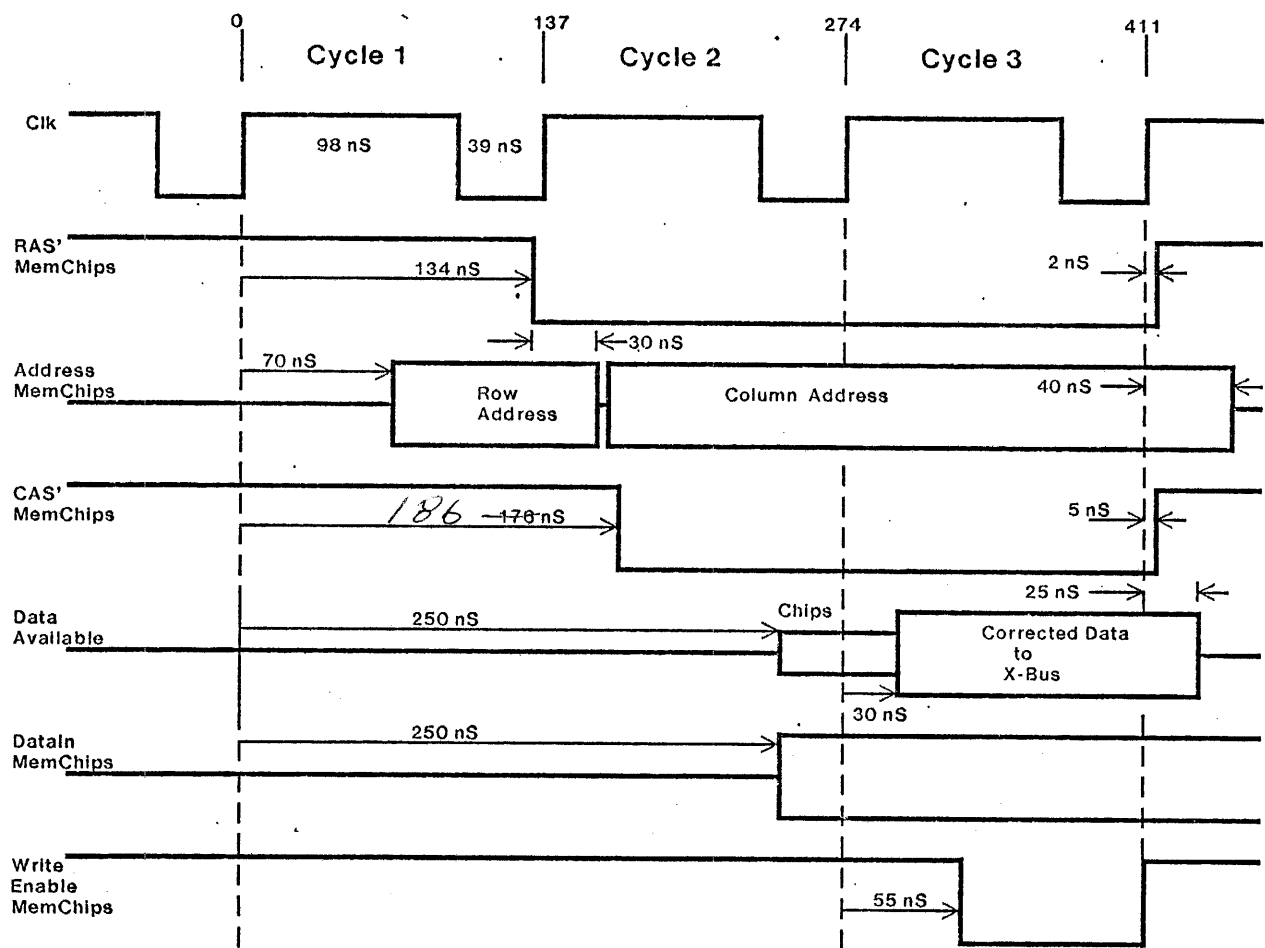
Dandelion Memory Registers

Memory Timing

Typical memory timing is shown based on measurement of processor port of the first stitchweld card. Both the processor and display ports will be described.

Processor timing is shown below. The memory address must be valid on the Y and YH busses early enough that the proper bank is selected and address lines valid for RAS' (row address strobe). The column address bits are latched by the RAS' signal. The CAS' (column address strobe) signal occurs 42 nS after the RAS' signal and latches the column address in the memory chips. Data becomes valid at the output of the chips at a maximum of 150 nS after RAS' or 100 nS after CAS', whichever is later. (Because 16K chips are used, 1 of the 7 bits used for RAS must come from the low byte. The contents of the low byte are often the result of an arithmetic operation computing the next address (high byte is held fixed). The maximum settling time of the high nibble of the low byte is too long if a carry from the low nibble occurs. Consequently, bit 12 (instead of bit 8) of the low byte is used during RAS. Consistent juggling occurs for map references so that this is invisible to the microcoder. This affects only the maximum run of sequential page mode accesses as described below.)

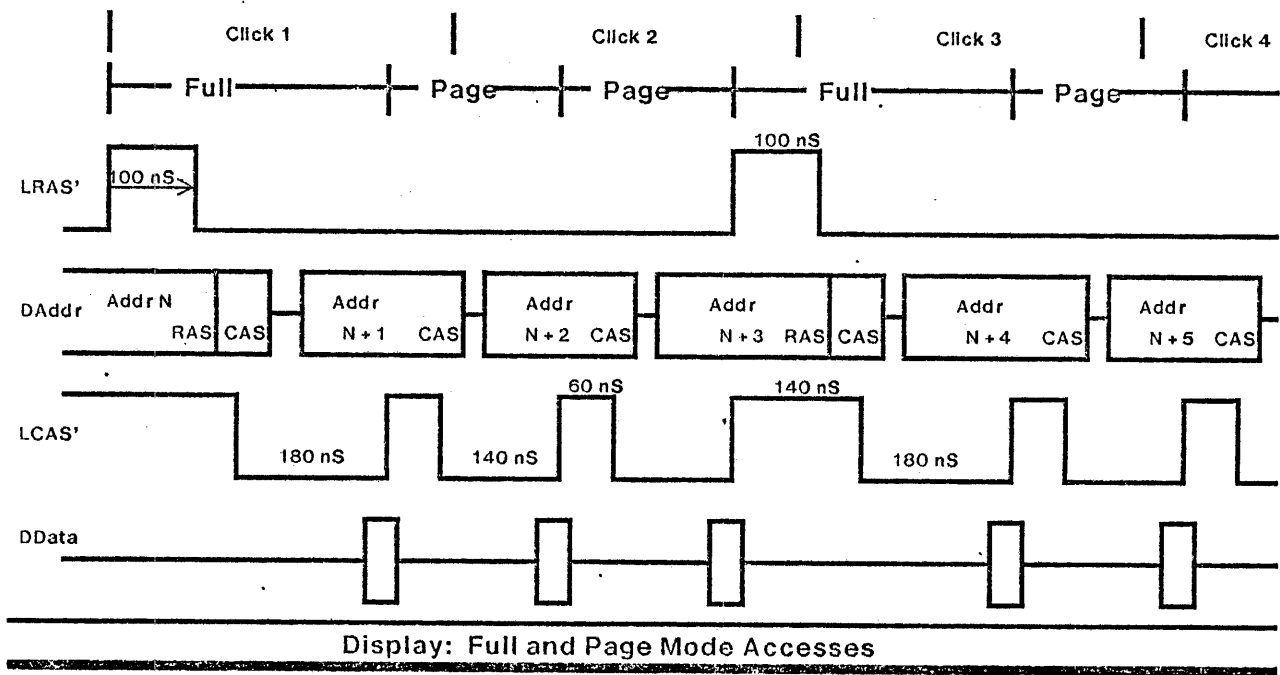
When writing into memory, the data to be written must be supplied during the second cycle of a click. The data is actually written in the latter half of the third click. Notice that up until the presence of the write pulse, all signalling is identical to a read cycle. The memory chips latch and hold the old data on their outputs during a write pulse if it occurs more than 150 nS after the RAS' signal. Thus, it is possible to write into a location and read data from it, all in the same memory cycle.



Normal Memory References through Processor Port

The display port supports both full and page mode accesses. The data delivered to the display port is not error corrected. The full access cycle time is 280 nS and the page mode access time is 200 nS. While the full access time is smaller than that specified in the data sheets (320 nS) for continuous operation, it is the average that is important, and the average cycle time in this case is 342 nS (6 full accesses per round, counting click 5). A page mode access occurs when the RAS' signal goes low and the CAS' signal cycles several times, strobing several different column addresses (low 7 bits) into the memory chips while retaining the same row address. (Because bit 12 is used during RAS, the maximum number of sequential page mode accesses between full accesses is 7, since bit 12 will change on every 8th access. The insertion of full accesses at the appropriate times is handled by the display controller.) The display controller generates all the timing and address signals for the display port.

In normal operation, the display controller will seize the low bank of memory for 4 clicks of every round. It will start with a full access which is aligned on a click boundary, and then proceed with page and full accesses until the end of click 4. The other page or full accesses will not necessarily be synchronized with any click or cycle boundaries. They are packed so as to maximize the number of accesses during the 4 clicks the display has the memory.



Memory Interface Signals

There are 6 groups of signals for the memory control card. They are power, data/address, error, register control, processor clocks, and memory clocks.

Power (29 watts typical)

+ 12 volts
+5 volts
-5 volts
Ground

Data/Address Busses

Y Bus	16	data must meet 15 nS setup time, address meets setup based on bit position (Y0-7 65nS, Y8-11 11nS, Y12-15 36nS)
X Bus	16	data available 45 nS min. before end of cycle 3
YH Bus	8	data must meet 80 nS setup time (includes 2 extra lines)
Display Data	16	lines
Display Address	16	lines
Task#	3	lines (used for error logging)

Error Indication

MemError available approximately 45 nS before end of cycle 3

Register Control setup time = 70 nS

The control section of the processor must supply signals to load registers and enable data output to the bus.

mem		When ANDed with cycle 1,2, or 3 produces the following:
c1	MAR←	Loads memory address reg. from Y & YH busses
c2	MDR←	Loads memory data reg. from Y bus & starts write
c3	←MD	Gates memory data to X bus
Mctl←		Control reg. with check bit inversion, bank select, and ECC enable bits
←MStatus'		Gates MStatus register to X bus
MapRef	c1 only	Does a MAR←, but with juggled bits.
Refresh'	c1 only	Causes a refresh operation and increments refresh counter
Disp/Proc.'		Goes high when display port is using low 64K.

Memory clocks (LH - low-high transition, HL - high-low transition)

RAS'	HL 121 nS into cycle 1 re: qualified clock. LH at end of c3
CAS	LH 24 nS into cycle 2 re: qualified clock. HL at end of c3
LRAS'	Same as RAS, except when low bank is used by display
LCAS	Same as CAS, except when low bank is used by display
WPulse	LH 40 nS after qualified clock, HL 19 nS before qualified clock

Processor clocks

Cycle1'	Low during cycle 1 of click.
Cycle2'	Low during cycle 2 of click.
Cycle3'	Low during cycle 3 of click.
ppClock	137 nS period with 39 nS pulse width (LH - 7 nS, HL - 91 nS)

Memory Banks & Standby Power

The system (including the storage card) contains a total of 12 - 16K memory columns. To minimize power consumption, only one column at a time is cycled for normal memory accesses, and two at a time if the display is using the low 64K while the processor is using one of the higher banks. During refresh, all banks receive a RAS only cycle. Cycling all banks continuously in refresh cycles causes a drain of 3.75 amperes from the + 12 volt supply while cycling only 2 banks with normal memory references consumes 1 ampere.

System Parts Cost

The memory card can be broken up into control and error correction logic, and memory chips. Below are the chip counts and cost estimates for both the memory control card (64K) and storage card (128K).

Component	Mem. Control + 64K		128K Storage	
	IC's	\$	IC's	\$
Memory logic @\$ 1	77	77	26	26
Memory chips @\$ 5	88	440	176	880
Bypass Caps @\$.20	165	33	202	41
PC Board	-	100	-	100
Total	165	\$650	202	\$1047

Probability of Single and Double Errors

The following calculations yield probabilities of errors due to independent random processes in each chip. They do not include correlated events such as power line transients or static discharges which could affect all of the chips at the same time. A memory with 22 bits/word is assumed.

The hard failure rate is assumed to be .04%/1000 hours. The mean time to a single hard chip failure is about 13 months (9470 hrs.) for a 192K system using 16K memory chips.

The soft error rate for the chips is assumed to be 1%/1000 hours. Following are the probabilities of 0, 1, and 2 soft errors in a 22 bit word in a 10 hour period. 10 hours was selected as the interval over which errors could accumulate, with the system being reset after 10 hours. It is expected that most systems would be rebooted at least once in 10 hours. The mean time between single errors is 38 intervals and the mean time between double errors is approximately 36,200 intervals. It should be pointed that these probabilities are those that one would expect to measure with a program which continually scans through all memory cells looking for an error. If a program is confined to a small segment of memory, it would perceive a proportionately smaller probability of soft error.

Prob.(0 errors in 22 bit word in 12 bank system in 10 hr. interval)	= .9736
Prob.(1 single error in 22 bit word in 12 bank system in 10 hr. interval)	= .0263
Prob.(1 double error in 22 bit word in 12 bank system in 10 hr. interval)	= 2.76×10^{-5}

Error Correction Logic

The error correction logic generates 6 check bits which are stored with each 16 bit word to provide single error correction and double error detection. 6 chips are used to generate the check bits. 25 chips are used to generate the syndrome bits, correct the data bits, and provide the capability of recording errors on a per task basis. Details of the error correction logic, including correction code tables, are included in the logic drawings.

RAS-CAS Multiplexer

Between the leading edges of the RAS and CAS pulses, the contents of the address lines to the memory chips must change from the row address to the column address. This transition must occur after the 20 nS hold time requirement for the row address and before the column address setup requirement of 10 nS after the CAS pulse (i.e. setup time = -10 nS). Since CAS' follows RAS' by only 42 nS, this leaves a 32 nS window in which the transition can take place. The design center for typical chips is 30 nS which allows for a 10 nS delay in the RAS' buffer chip or 22 nS extra delay in the delay line, multiplexer and address driver circuits. The RAS-CAS multiplexer is switched by a delayed version of RAS, with the delay generated on board to minimize possible skew.

The delay circuit is implemented using an inductor, capacitor, and resistor (single element delay line terminated in its characteristic impedance of 360 ohms) feeding a gate. This is used in preference to a packaged delay line because it costs less, takes less space, and will be easier to adjust when converting to a printed circuit board which will probably require adjustment of the delay. To minimize noise pickup, the delay components should be located next to the gate receiving the signal. (Delay in seconds $D = (LC)^{.5}$ Impedance in ohms $Z = (L/C)^{.5}$; L-Henries, C-Farads)

Memory Array Line Termination

It is important to terminate the lines driving memory chip inputs to prevent damage due to undershoot on the high-low transitions. The memory chips require that their inputs never go below -1 volt, to prevent forward biasing some internal parts of the chip and causing damage. Undershoot (and overshoot) result from the transmission line behavior of the signal lines in the array. This behavior becomes evident whenever the signal risetimes are comparable to the propagation time through the line (5nS and 3nS respectively for this system). Either series drive or shunt termination can be used with the lines. Both of these are shown in the following figure. The series drive consumes less power, since there is no steady state current flow, but it has a much longer propagation time.

Because speed is important in this system, the direct drive with matched termination was selected. The Schottky TTL drivers have an asymmetrical output capability; they sink more current than they can source. For this reason, the termination resistor is terminated to a +2 volt source instead of ground. To obtain this voltage without excessive dissipation, half of the address lines are driven with the true value of the address while the other half are driven with the complement of the address. All of the termination resistors are tied to a common capacitor tied to ground. This amounts to a voltage divider with half of resistors tied high and the other half tied low, since each true value has a complement.

The last figure gives results of tests on the D0 96K storage card. This is of interest because it provides a reasonable estimate of what impedances might be expected for the address and control lines on a PWA. The impedance of the trace loaded with chips is in the 50 to 60 ohm range. Using a 62 ohm resistor to 1.9 volts results in a source current requirement of 25 mA and a sink current requirement of 29 mA. The effect of a 20 mA source current on chip temperature rise is tabulated in the bottom box of the figure. The dissipation in the S241 is the largest, since it contains 8 drivers. The extra junction rise should not cause trouble, however, since the 54S241 is rated to 125°C with the same .5 watt internal dissipation (i.e. a 55°C rise above the commercial part spec.). To ease the dissipation slightly, the termination resistor values are selected to produce about .4 volt undershoot, a value slightly larger than the characteristic impedance of the line.



Power Supply Considerations

The memory system requires +12 volts, +5 volts, and -5.2 volts, all with a tolerance of $\pm 10\%$. The -5.2 volt supply must never go positive with respect to ground when the +12 volt supply is on. While the chip manufacturer does not require it, they suggest that the -5 volt supply be the first to come on and the last to go off. The supplies should be bypassed at entry to the board and at every logic chip and every third memory chip.

Propagation Delays and Line Termination



Purpose is to minimize ringing and undershoot on the signal line.

Series Drive

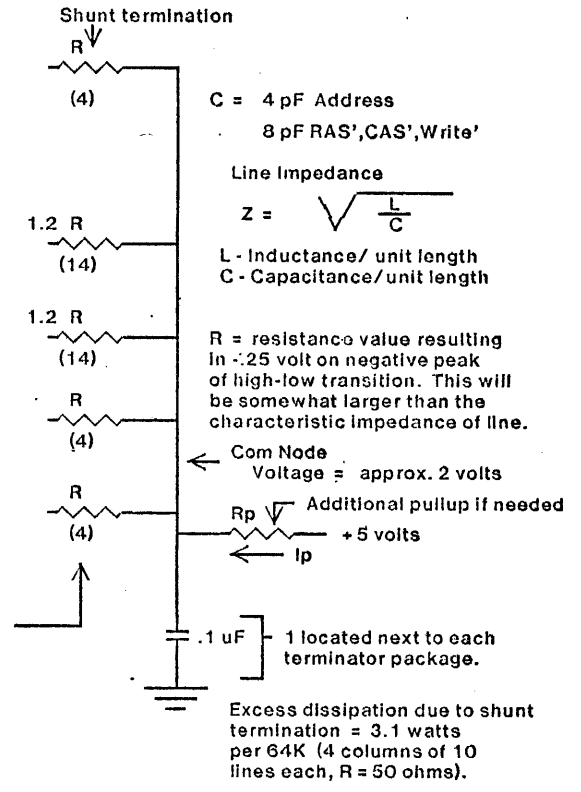
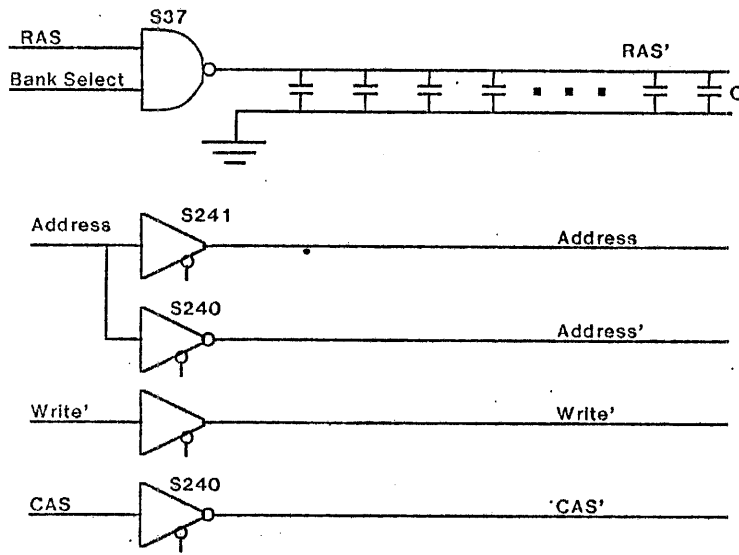


Note that point B takes the longest to see a full signal swing. This termination technique takes less drive current. Resistor value, R, should be same or slightly less than line impedance.

Direct Drive with Matched Termination



Termination resistor requires large drive current in logic high state. Terminating resistor to +2 volts instead of ground eases this problem considerably. Resistor R, should be same or slightly more than line impedance.




Additional Pullup to reduce Dissipation in S240, S241, & S37

S240, S241 have max source current of 15 mA. A voltage source of 4 volts with a 25 ohm series resistor makes a reasonable model of the S240 output when in the high state. Thus, logic high output voltage must be $4 - .015 \times 25 = 3.625 \text{ V}$ or greater. Terminator current must not exceed 15 mA and consequently the common node voltage supplied to terminators may have to be raised by use of an additional pullup resistor. Below is a table of terminator resistor values, R , and associated pullup resistor parameters. The 64 mA sink spec is also met by the following values.

Term R	Com Node V	Additional Pullup		
		R_p	I_p	R_p Diss.
47	2.92	2.2	.94 A	1.95 W
75	2.5	6.8	.37	.92
100	2.13	22	.13	.36
121	1.82	Infinite	0	0

Number of each type of line for 64K bank.



Columns must be paired to balance current from address lines.

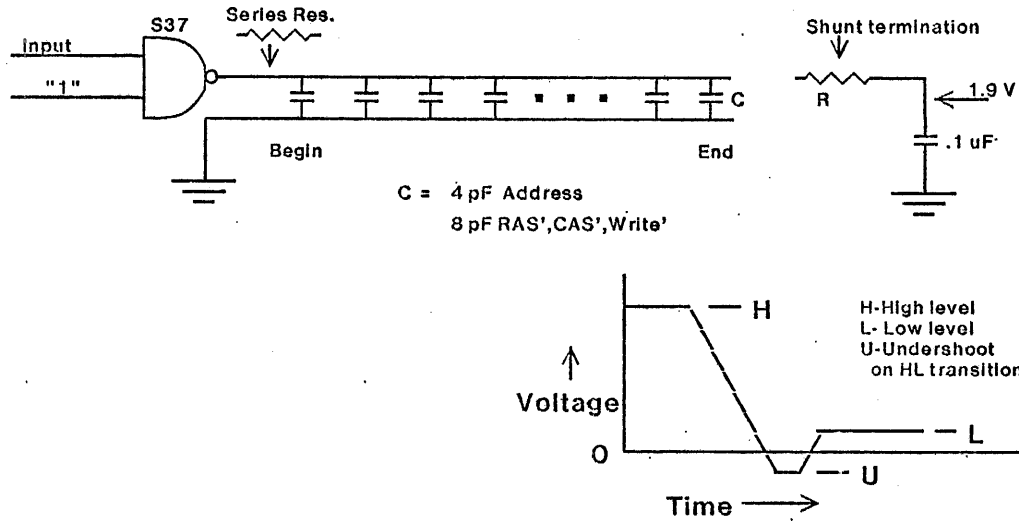
Number of 1's & 0's contributing to the capacitor voltage under various conditions

No Access	Write	Refresh
14 - 1	14 - 1	14 - 1
14 - 0	14 - 0	14 - 0
8 - 1	6 - 1	8 - 1
4 - CAS	4 - CAS	4 - RAS
	RAS	
	Write'	
22 + 1.33	20 + 1.33	22 + 1
	+ .25 + .85	
<u>23.33</u>	<u>22.43</u>	<u>23</u>

Capacitor voltage =
 $\frac{23}{40} \times V_H = .575 \times 3.3V = 1.89V$
 This can be increased by use of additional pullup shown above.

Shunt termination of the memory drive lines can speed up the memory. See 96K storage card tests. Supplying 2 volts for the terminating resistors can be accomplished by inverting the signal to half of the address lines. This guarantees a level half way between logic high and logic low. Connection of terminators for RAS, CAS, and write moves this level up somewhat and causes some variation during different operating conditions of the memory. This variation is not too great as can be seen from the calculation in the box above. Blasing the termination voltage upward with an extra pullup could be used to reduce the power dissipation in the driver chips when in the high state. With R greater than about 120 ohms, this should not be necessary.

Memory Line Drive and Termination



Tests on D0 96K Storage Card 6-24-79										
CAS Line										
	S37	S37	18 Mem. Chips			Voltages			Extra Supply	
Series	Input	Output	Begin	End	Termination	H	L	U	Current	
LH *	0 nS	4 nS	9 nS	8 nS	27 ohm series	4 V	.2	-.4	0	
HL	0	6.5	12	9.5						
Shunt	LH	0	4.5	5.5	8	62 ohms to 1.9 V	3.5	.3	-.3	.01 A
	HL	0	4.5	5.5	8					
						R = 100 ohm to 1.9 V				-.75 V
Addr. Line										
	S37	S37	18 Mem. Chips			Voltages			Extra Supply	
Series	Input	Output	Begin	End	Termination	H	L	U	Current	
LH *	0 nS	4 nS	8 nS	7 nS	R = 27 ohm series	3.8V	.2	-.7	0	
HL	0	5	9	7.5						
Shunt	LH	0	4	5	7	R = 62 ohm to 1.9 V	3.45	.3	-.4	.01
	HL	0	4	4.5	7					
<p>* Time is from 1.3 V point on input to .8 V for HL and 2.8 V for LH. 18 chips in row 9" long Risettime and falltime for shunt terminated line = 5nS. Propagation time for 10" line with 6 loads LH- 3 nS HL- 2 nS. (TTL loads) Extra supply current is that due to the 62 ohm shunt termination. Undershoot is measured at the end of the line.</p>										
Power Consumption 108 memory chips MK4116-2										
<u>Conditions</u>					<u>+12 V</u>	<u>+5 V</u>				
RAS cycling 310 nS -L & 100 nS -H					1.41 A	.07 A				
CAS cycling 310 nS -L & 100 nS -H					.07	.07				
RAS & CAS cycling L- 310 nS H- 100 nS					2.13 A	.14 A				
Quiescent current					.07	0				

Effect on IC package dissipation						
Package	Thermal Resistance junction to ambient	Internal Dissipation	Extra dissipation for 20 mA @ 3.5 V out	Extra Junction temp. rise	Total temp. rise	
S241 20 pin	80 °C / watt	.5 W	.24 W	19.2 °C	59.2 °C	
S37 14 pin ceramic packages	100 °C / watt	.18 W	.12 W	12 °C	30 °C	