A Fast Access Bulk Memory System using CCD's/ A Recorder Buffer Memory using CCD's

K. Siemens C. R. Robinson J. Mastronardi

Bell-Northern Research, Ottawa, Ontario, Canada

ABSTRACT The 8K CCD shift register memory developed at Bell-Northern Research has been applied to two prototype memory structures. A disc replacement system using parallel storage and a block addressable serial storage buffer have been developed to study optimization of design techniques to achieve a wide range of data handling rates, system flexibility, and low power bulk memory systems.

Bell-Northern Research has developed programs to assess the feasibility of using CCD shift registers in bulk memory systems. An 8K CCD memory device was used in two systems, a disc replacement that could be used on a PDP-11 Unibus, and a block addressable buffer memory capable of handling a wide range of data rates. The first memory system, with a capacity of 64K words (1 M bit) was designed as a direct replacement for the DEC RS64 disc memory. It was used as a vehicle both to study the problems and system parameters associated with a CCD memory system and to compare operational performance with that of the disc memory.

The second system, which was contracted for by NASA Langley, required particular optimization techniques in applying CCD's to achieve wide data rates and low power operation. A discussion of these systems follows a description of design considerations using the 8K CCD shift register memory.

CCD DESCRIPTION

The designs of the memory systems discussed in this paper were influenced by the configuration and operational requirements of the 8K CCD that was developed at Bell-Northern Research. The charge coupled device configuration determined the design of the addressing network, the clock timing, and the read/write sequence. The 8K CCD is structured as 32 recirculating dynamic shift registers (tracks) of 256 serial data bits each. In any system addressing scheme the 8 least significant address bits form a virtual address of the data location within each track and is determined by the count on an 8 bit counter. The 0 address is arbitrarily set the first time data is entered and is matched to that of the counter. The next 5 bits are applied to the address decoder inputs on the device; this permits random access to any track within one bit time. Bits of greater significance control access to the chip itself through the decoded chip enable (CE) inputs.

Two clock phases, $\emptyset_1$ and $\emptyset_2$, are required to transfer data along the shift registers, while a third clock, $\emptyset_3$, controls peripheral circuitry on the chip. A typical CCD timing chain, considered to begin with the rising edge of $\emptyset_2$, is shown in Figure 1. It is important to note that a READ operation always precedes WRITE. The CE and Address which are set up for the READ during one timing chain (output data becomes valid during $\emptyset_3$) is stored on chip in preparation for the WRITE command and DATA INPUT which occur during $\emptyset_2$ of the next cycle (see Figure 1).

*This work was supported by the National Aeronautics and Space Administration under Contract NAS1-13507. Contract Monitor: A. Fripp, Langley Res. Center

*This work was supported by the National Aeronautics and Space Administration under contract NAS1-13507. Contract monitor: A. Fripp, Langley Research Center.

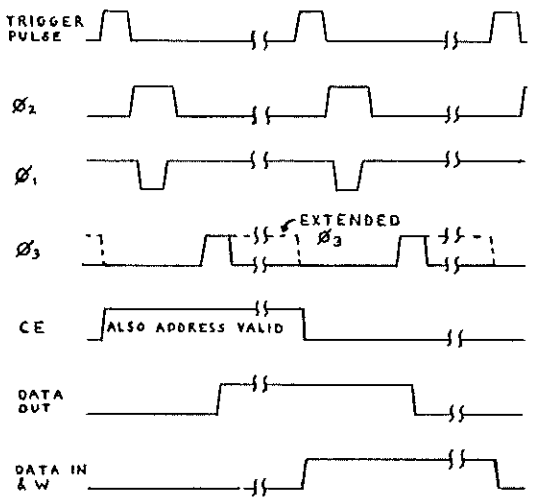

Fig. 1 Typical CCD Read/Write Cycle Waveforms.

Clock and control signal loading and level variation are critical parameters in CCD operation. The 8K clock phases ϕ_1 and ϕ_2 have a worst case loading of 700 pf. each. Rise and fall times are required to be at least 25 nsec, but not greater than 50 nsec for high transfer rates. To meet the requirements of high capacitive loading and high speed, MH0026 MOS clock drivers were chosen for both memory systems with a limitation of 4 CCD's per driver established. The ϕ_3 clock has a capacitive loading of 150 pf. so that a driver can operate a greater number of devices. Since the address, CE, W, and data input lines have relatively small capacitive loads, 10 pf. or less, simple drivers such as TTL open collector gates or TTL-to-MOS drivers can be used to drive a larger number of devices.

A typical range of operation for the 8K CCD is shown on the Schmoo plot in Figure 2. The range remains relatively constant throughout the frequency specification for the device; that is, 10 KHz to 1 MHz. As a greater number of CCD's are driven by the same clock driver, the clock rise and fall times increase and the Schmoo plot tends to shift toward a higher substrate voltage requirement. Since the CCD has a broader range of operation at 12 volts V_{DD} than at 10 volts, a better yield of devices is obtained at the higher operating point, despite a power penalty.

It must be kept in mind, however, that the majority of the power dissipation is in the drivers and not in the CCD itself. Driver power dissipation is derived from the standard power calculation:

$$P = P_{AC} + P_{DC} = CV^2f + VI \text{ (Duty factor)}$$

which contains a substantial DC component.

It is shown later how to manipulate the drivers and the timing chain to reduce driver power dissipation.

The power supply operating point for the disc replacement memory was selected as

$$(V_{DD}, V_{BB}) = (12, -5)$$

For the buffer memory low power operation was mandatory so that its operating point was selected as

$$(V_{DD}, V_{BB}) = (10, -3)$$

A device operating window of $\pm 5\%$ around the operating point was used over the required frequency range in selecting components; see Figure 2.

CCD DISC REPLACEMENT MEMORY

The first system that was designed to incorporate the 8K CCD was a disc replacement system with a capacity of 1 M bit. A block diagram of the CCD memory system is shown in Figure 3. The memory emulates the RC11/RS64 disc system and interfaces directly on the PDP-11 Unibus. The CCD Memory controller is much simplified compared to the RC11 yet retaining the same essential features.

The 1 Megabit memory is organized as 64 K words consisting of 16 parallel data bits and 1 parity check bit. An increase in word size requires only the addition of sufficient memory devices in parallel with the existing memory and extension of the address buses providing

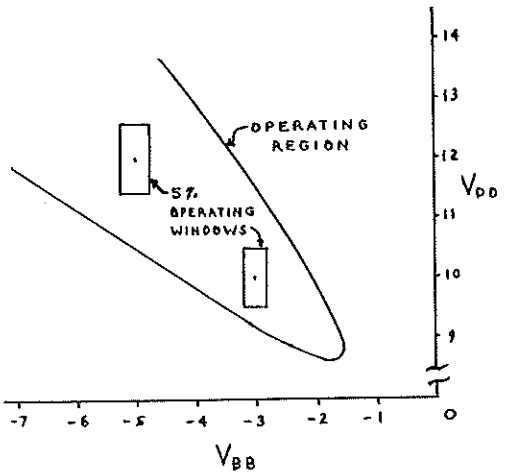

Fig. 2 Typical Operating Range of the 8K CCD.

for extremely flexible system design. Software access is permitted to any 32 word sector. Since all data is recirculating in 256 word loops maximum access time to commence transfer is 255 clock periods. Therefore at an operational clock rate of approximately 800 KHz, the average latency time is about 150 μsec.

Three modes of operation are required:

1. Idle mode: to maintain data the CCD's are idled at 20 KHz and chip enable is inhibited.

2. Race mode: When a data transfer request is enabled the idle mode timing pulse is cleared and the virtual address of the data position on a track is incremented. This address is then compared with the 8 least significant bits in the disc address register in the controller. In case of mismatch the timing chain is triggered, data shifted one bit, the virtual address counter again incremented and a fresh comparison made.

3. Transfer mode: When the virtual address matches the start of transfer address that was set up in the controller, a WAIT condition is set up before the CCD timing chain is again generated.

From this point there are some differences between the WRITE and READ cycles. (see Figure 1).

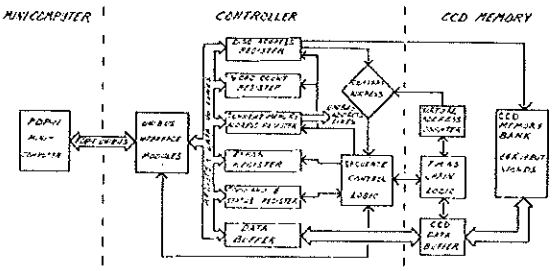

Fig. 3 Block Diagram of the Disc Replacement Memory

During a write cycle, the address match signal commands the INTERRUPT CONTROL MODULE to request (NPR) word transfer. When bus control is granted, the MASTER CONTROL MODULE brings the contents of the core location (as specified by the Current Memory Address register) into the controller data buffer. When this has been completed, bus control is relinquished and a TRANSFER REQUEST pulse is sent to the CCD memory logic. Upon receiving this signal, the memory logic reads the contents of the data buffer and triggers the CCD timing chain, thereby storing the data in the CCD memory. When the WRITE operation is finished, the CCD memory returns a TRANSFER REPLY pulse back to the controller. This pulse is then used to increment the three parameter registers in the controller; the current memory address register, the disc address register, and the word count register. After an appropriate delay to allow for propagation of the increment pulse, the RACE signal is again asserted and the cycle is repeated. The virtual memory address is now in step with the current word address pointer so that the WRITE cycle will repeat until the word count register overflows, indicating that the required number of transfers has been made. Since the controller addresses 32 word sectors, the number of transfers will always be a multiple of 32. At the end of a transfer cycle, the memory goes back to its slow speed IDLE mode.

During a READ cycle, when address match is detected,a TRANSFER REQUEST

signal is generated which triggers the CCD timing chain to enable a data shift and READ. As soon as the output data from the CCD becomes valid, it responds with a TRANSFER REPLY pulse which the controller uses to latch the data word into its data buffer. Once this is accomplished, the UNIBUS is requested for an NPR word transfer. When BUS CONTROL is granted, the contents of the data buffer are transferred into the core location specified by the current memory address register. When the transfer is acknowledged, the bus is released. At the same time, the three parameter registers are incremented. After allowance for propagation delay, the memory logic is again signalled to go back to the RACE mode. As with the WRITE cycle, the read cycle continues until the word count register overflows. Transfer rate for the CCD disc replacement memory is governed by the PDP-11 UNIBUS and is approximately 300K words per second.

Since the CCD memory is a prototype system, only simple error detection schemes, latency error and parity error were incorporated. A latency error signal is raised if during the data transfer mode an address match is completed but a TRANSFER REQUEST signal is not received with 50 μsecs. The CCD timing chain must then be triggered to keep the dynamic memory refreshed, causing RACE mode to be reestablished.

During a WRITE cycle, a parity bit is generated from the data input and is stored as a 17th bit of the word. When data is read, the output data is used to generate a parity bit which is compared with the stored 17th bit. If they do not match, a parity error signal is asserted.

PERFORMANCE RESULTS

The three modes of operation of the CCD disc replacement memory demonstrate a capability not possible with magnetic disc memories. Not only can the memory stop temporarily at an address location to wait for a transfer request from the processor, but it can race to its start address at the beginning of a data transfer. The low speed during idle mode is used to conserve power.

Operating power of the CCD memory including controller and memory logic and drivers was measured at approximately 100 watts. This compares favourably with the RC11/RS64 disc which nominally requires 250 watts and this ratio would improve as the number of equivalent discs increases. This design, however, was not optimized with respect to power consumption, since standard TTL, components and open collector gates were used for board interface and input drive lines (other than clocks) to the CCD's. Modular construction of the memory boards resulted in redundancy of open collector drivers. It is expected that a redesign using low-power Schottky TTL,CMOS devices, and TTL-to-MOS drivers for all CCD input lines would reduce the operating power of a 1 M bit system to less than 25 watts. Other techniques of saving power consist of manipulating the CCD timing chain and switching off the clock driver's power supply to reduce their D.C. power dissipation as will be discussed later.

Functional tests of the system included writing and retrieving text files from the CCD memory. Software programs on the PDP-11 were written to write data test patterns into the CCD memory. Every CCD chip that was used was previously tested for pattern sensitivity and frequency of operation on a Macrodata MD-104 memory tester. A continuous program of READ/COMPARE following the initial WRITE of test data was carried out with zero errors detected after weeks of continuous operation in which stored data was read and compared.

BLOCK ADDRESSABLE BUFFER MEMORY

Experience in the development of the CCD disc replacement memory enabled considerable design improvements to be applied to a different CCD memory structure. A buffer memory capable of accepting serial data with rates from 150 KHz to 4.8 MHz was developed under contract with NASA Langley. A low power memory with a randomly-accessible memory block structure was specified. A block diagram of the memory is shown in Figure 4. A typical use for this memory would be to accept bursts of data arriving from sources which may be clocked at widely differing rates, temporarily storing the data in separate blocks, then transmitting at a fixed rate selected blocks of data to a permanent store. Since serial data must

be stored in each memory block in a first-in-first-out basis, a CCD shift register memory is particularly suited for this application.

The system described is a feasibility model representative of a larger operation memory. Some of the problems that had to be overcome in the memory design were:

1. Accepting data rates over the range 150 KHz to 4.8 MHz with a CCD specified to operate only to 1 MHz.,
2. Synchronization of incoming data bursts with the virtual start address of partially filled memory,
3. Blocks of memory to be randomly accessible.
4. Power consumption for 100 K bits of data to be less than 2 watts in idle mode and under 6 watts in access mode.

In order to reduce the CCD clock rate and also maintain synchronism between independantly accessed storage blocks four way multiplexing of the input data was used with high speed RAM's operating as a variable length input buffer to the banks of parallel CCD's. Thus data presented to the system at extremes of 4.8 MHz to 150 KHz is fed to the CCD memory at 1.2 MHz and 37.5 KHz respectively. Input data bursts of integrals of 1 K bits were used and all non-accessed CCD blocks were triggered at 1/4 the data rate to ensure that both stored data was kept alive and at the end of the data input the idle CCD's would have been clocked an integral number of 256 transfers to bring all blocks back into virtual address synchronism again.

Input data is collected in 4 bit nibbles from where it is written in parallel into 4 high speed RAM's. This action proceeds with the RAM address starting at zero and incrementing it after every 4 parallel bits have been written in. When the CCD virtual address reaches zero the RAM address is reset to zero and the data contained at that address is transferred in parallel to the 4 CCD inputs. Fresh input data then replaces this data at address zero following which both RAM and CCD addresses are incremented and the READ/WRITE cycle repeated. The size of the RAM buffer is set at 4 x 256 bits and is determined only by the track length of the CCD and the level of multiplexing.

Special provisions are made to minimize the power dissipation. Low power Schottky TTL and CMOS logic are used with silicon on sapphire CMOS static RAM's to meet the necessary speed/power criteria. Specially selected CCD's with lower V_{DD} (10 volts cf. 12 volts nominal) were used with optimised timing waveforms. The major source of power dissipation lies in the clock drivers when their output is held low. Refering to Figure 1 the most important technique in saving power is a modification of the usual ϕ_1, ϕ_2 and ϕ_3 clock timing chain. The nominal 150 nsec. ϕ_3 width was expanded to a maximum consistent with the cycle time and in addition, this pulse was used to gate off the power supply to ϕ_2 which is held "low" over this period. Thus ϕ_1 and ϕ_3 drivers are held "high" in their low dissipation state while ϕ_2 driver dissipation is reduced by turning off the supply for a large part of each cycle, this is particularly beneficial during low speed or idle operation. Figure 5 shows a comparison of driver power dissipation as a function of frequency. All CCD address lines etc. are held high when the CCD is not being addressed and clocks to CCD's not containing data are held static.

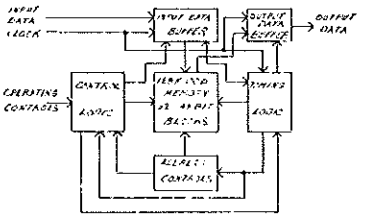

Fig. 4 Block Diagram of Buffer Memory.

Conclusion:

Bell-Northern Research has demonstrated the feasibility of applying CCD shift registers in both parallel and serial organized memory structures. The disc replacement memory was used primarily as a vehicle for studying the problems associated with driving large numbers of CCD shift register memories that have highly capacit-

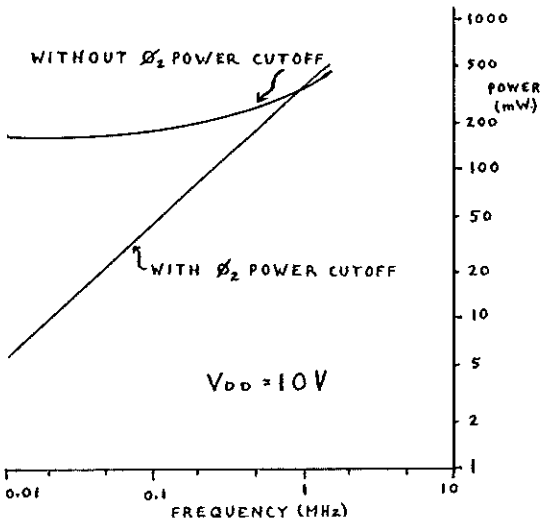

Fig. 5 ϕ_2 Driver Dissipation vs. Frequency With and Without Power Cutoff Circuit.

ive clock loads. The memory is a parallel structure that demonstrates a latency period at least an order of magnitude less than magnetic disc memories. The data throughput is optimized to match the data transfer capability of the processor and DEC Unibus. Total non-optimized system power is less than half that of the equivalent magnetic disc. Expansion of memory capacity and word length is facilitated by the parallel storage structure. Power consumption of this memory could be reduced greatly by a redesign using techniques that were incorporated into the second system design, a serial storage buffer memory. The buffer memory demonstrated the capture of data with a widely varying bandwidth, flexibility of access to any ordered set of memory blocks, and the feasible operation of a bulk memory system with extremely low power. Both of the systems were designed to provide ease of expansion of memory capacity, in size and number of words and blocks. The amount of peripheral circuitry and power consumption need be increased only slightly to handle large increases in memory capacity. It is expected that developments of CCD chips with larger capacities, higher data rates, and lower clock line capacitances will make CCD's increasing attractive in the design of memory systems.

ACKNOWLEDGEMENTS:

The authors wish to thank Archie Fripp, of NASA Langley, who initiated the CCD buffer memory project under contract NAS1-13507 part of program 506-18-21-03