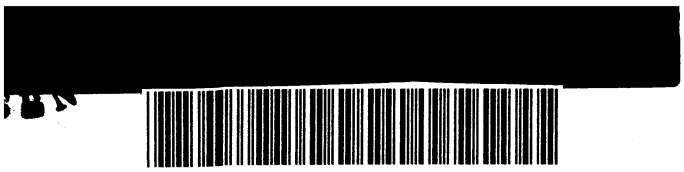

CERN-DD/76/15
September 1976

A.W. Booth*

FIFO Buffer Memories for the Omega PDP 11 High Speed Data Links

Abstract

A high speed packet switching network has been designed and implemented at CERN. It provides remote access in real time to a medium sized computer (CII.10070) for analysis of experimental results from two large experimental facilities. The Omega spectrometer, and the Split Field Magnet. It also provides a general communications system between tasks in any of the computers connected to the network, (the majority of which are DEC. PDP 11s).

The Omega PDP 11 high speed data link interface, which forms the basis of the present network, incorporates a buffered transmission scheme.

This paper is in two parts. Part 'A' Describes an optional 32/63 byte First In - First Out (FIFO) buffer memory, designed to be used in transmission speeds.

Part 'B' describes two large (4096 bytes and 1024 bytes) FIFO buffer memories which have been built as test facilities for the above mentioned interface.

* Student, Leicester Polytechnic, Leicester, England

DD/ib

PART 'A'.

Optional 32/63 Byte FIFO Buffer Memory

Introduction.

Over long data links with fast transmission speeds, the 12 byte storage capacity of the First In - First Out (FIFO) buffer in the PDP 11 data link interface is insufficient. The reason for this is explained in the following text.

1/ General

The PDP 11 data link interface (ref. 1) forms the basis of the present Omega computer network (OMNET). The interface incorporates a First In - First Out buffer. The reason for a buffered transmission scheme and the way in which it is implemented in the interface were descibed in the Cern yellow report 74-9 Chapter D34 (ref. 2), from which a section is reproduced here: -

1.1/ I/O Speed variation.

A safe way to deal with I/O speed variations of the sender and receiver is 'handshaking'. However, over long distances it is slow owing to delays in the cables. To get around this problem a 'buffered transmission' scheme has been implemented (see fig. 1). The transmitting computer (C1) sends data at its maximum rate. At the receiver, data enters a small buffer and is taken out by the receiving computer (C2) at its own specific speed. If the receiver is slower than the transmitter, the buffer tends to fill up and overflow. Before this can happen, a signal (HOLD) is generated and sent back to the transmitter holding up the transfer of data, Because of cable delays this does not happen immediately, and therefore the buffer should still have enough empty places to accept a few more data words. If, after some time, the buffer has been emptied sufficiently, the 'HOLD' signal disappears and data transmission may resume.

(The above paragraph was copied from Cern yellow report 74-9 chapter D34, together with figure 1).

2/ Shortcoming.

Originally, the Omega network was designed to provide high speed transmission of data between computers up to 0.5 kilometres apart. With the expansion of the network, there came long data links and consequently a problem. Over long data links with fast transmission speeds, a number of bytes have already been transmitted down the line while the 'HOLD' signal is being sent back from the receiver to the transmitter. Under these conditions the present 12 byte FIFO buffer in the PDP 11 data link interface overflows. This is because by the time the 'HOLD' signal reaches the transmitter to prevent any more data being transmitted, already there are more bytes in the line than there are empty places remaining

in the buffer. The following table shows the number of bytes to arrive at the buffer after the 'HOLD' signal has been sent.

Distances in kilometres.	Transmission speed.	Number of bytes arriving after 'HOLD' has been sent.
0.5	2.5 mbits/sec	2
1.0	2.5 mbits/sec	3
2.0	2.5 mbits/sec	5
5.0	2.5 mbits/sec	11
10.0	2.5 mbits/sec	21

The above results were obtained by using the following equation: -

$$N = (FL/7V) + 1$$

where 'N' is the number of bytes to arrive after 'HOLD' is sent,
'F' is the clock frequency (transmission speed),
'L' is the length of the link,
and 'V' is the propogation velocity of the cable.

If for example, the 'HOLD' signal is sent after there are six bytes in the buffer, then from the above table it can be seen that the buffer would overflow for data links longer than 2 kilometres, with a transmission speed of 2.5 Mbits/sec. consequently, when using fast transmission speeds over long data links there is a need for a larger FIFO buffer.

NOTE : FOR SIMPLICITY THE SCHEME REPRESENTS THE CASE FOR ONE DIRECTION ONLY

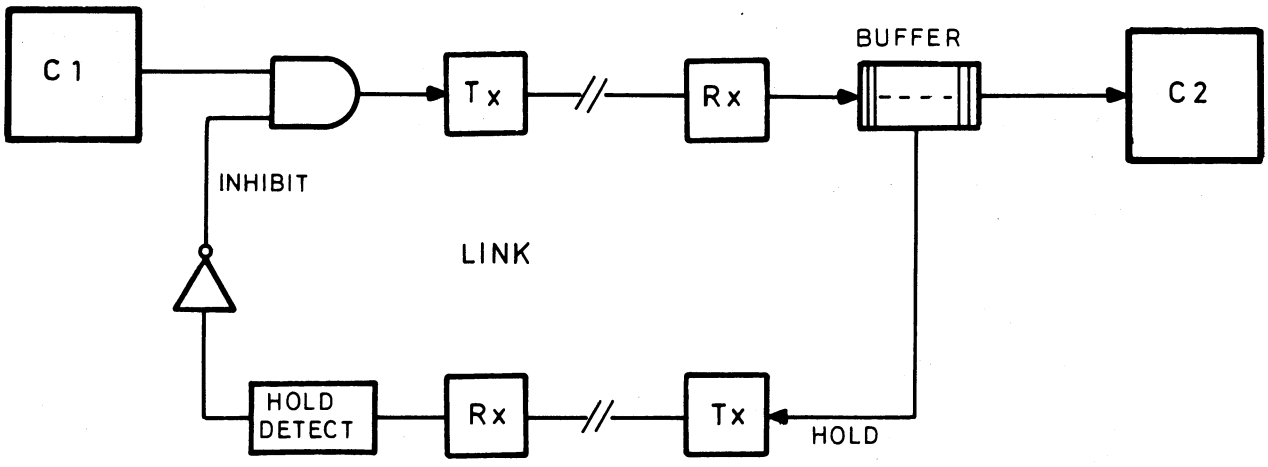

Fig . 1 — BUFFERED TRANSMISSION

3/ Optional 32/63 Byte FIFO Buffer.

Because the length of the data links varies considerably (at the moment between a few metres and 12 km) and also because the versatility of the Omega data link interfaces permits transmission rates as slow as 4.8 Kbits/sec and as fast as 5 Mbits/sec with the same physical link, a compatible optional 32/63 byte FIFO buffer has been designed and built.

In the 32 byte mode, the 'HOLD' signal may be sent after 16 bytes have been written into the buffer (I.E. 1/2 full). In the 63 byte mode, the 'HOLD' signal may be sent after 16 bytes, 32 bytes or 48 bytes have been written into the buffer (I.E. 1/4, 1/2 or 3/4 full). In both modes there exists the facility to kill the 'HOLD' signal (I.E. not to send 'HOLD' at any time) should a mode of operation be encountered where such action would be desirable (e.g. testing the ability of the buffer to store the full 63 bytes). Other features incorporated in the design of the buffer are as follows: -

A/ A DATA OVERRUN signal (indicating that an attempt has been made to write more than 63 bytes into the buffer).
B/ A BUFFER EMPTY signal (inducating that the buffer is empty).
C/ An OUTPUT READY signal (indicating that there is valid data on the outputs of the buffer).
D/ The 'HOLD' signal has been displayed on a light emitting diode (LED) for ease of fault finding.

Changing from the 32 byte mode to the 63 byte mode or vice-versa is accomplished by a small two-way switch mounted on the printed circuit card.

See appendix 'A' for technical description of the optional 32/63 byte FIFO buffer memory.

PART 'B'

TWO LARGE BUFFER MEMORIES.

Introduction.

Two large first in-first out (FIFO) buffer memories have been designed and built as test facilities for the Omega PDP 11 data link interface (ref. 1). The larger of the two is a 4096 by 8 bit buffer and incorporates 8 Semi 4200 integrated circuits, (4096 by 1 bit static MOS Random Access Memories-ref. 3). The smaller buffer is organised as 1024 by 8 bits and incorporates 8 Fairchild 2102 integrated circuits (1024 by 1 bit static MOS Random Access Memories - ref. 4).

The PDP 11 data link interface incorporates a FIFO buffer for hardware flow control, (explained in part 'A' of this paper). This buffer is normally of 12 bytes, (ref. part 'A'), but for test purposes however, a larger buffer is required to allow large amounts of data to be circulated in the 'Local Loop' test, (I.E. connecting the serial line output to the serial line input).

1/ Organisation.

The normal 12 byte FIFO buffer is, unlike the 4K and 1K byte FIFO buffers, a true SILO buffer, i.e. data can be read from one end of the buffer at the same time as data is being written into the other end.

The 4K and 1K buffers do not permit simultaneous read and write operations. The MOS Random Access Memory integrated circuits are organised in such a way that when the memory is 'selected', it must be EITHER to write data into it OR to read data from it. The integrated circuits are not FIFO memories.

The object of designing large buffers was to enable large amounts of data to be circulated in the 'local loop' test. There was no need to make them perform as true SILO buffers. In addition, the possibility of interleaving read and write cycles becomes impractical due to the extra control logic required to overcome the problems caused by having no 'handshake' on input to the buffer, i.e. to ensure that read and write cycles are actually interleaved, even though a request for both might be made simultaneously. Therefore, data is written into the buffer in one block in a write mode, and read from the buffer in one block in a read mode.

'GO IN' (being bit 8 of the link command and status register) when 'HIGH' causes the PDP 11 UNIBUS to be obtained for the first direct memory access transfer on input. It is used as the read/write signal for the Fairchild 2102 memory integrated circuits. The complement of 'GO IN' is used as the read/write signal for the SEMI 4200 memory integrated circuits. When 'GO IN' is high the buffer is in the read mode, and when low, it is in the write mode,

2/ No 'HOLD' signal.

The normal 12 byte buffer tends to fill up when the speed at which data is being read from it, is slower than the speed at which data is being written into it. The hardware flow control system requires that the FIFO buffer be able to warn of its impending overflow. Therefore a signal 'HOLD' is generated and sent to the transmitting computer telling it not to transmit any more data until this signal has been cleared, The 'HOLD' signal is sent after there are 'X' bytes in the buffer, where the value of 'X' depends on transmission speed, length of link, type of line, etc. (ref. part 'A' of this paper). As read and write cycles are mutually exclusive in the 4K and 1K buffers, there is never any need to send 'HOLD'.

3/ Design Features.

Both the SEMI 4200 I.C. and the Fairchild 2102 I.C. are random access memories (RAMs). Because the first byte written into the buffer of the PDP 11 data link interface must be the first byte read from the buffer, the simplest way of filling up and emptying memory locations is sequentially; i.e. when writing into the buffer, the first byte goes into memory location zero, the second byte into location one, etc; when reading from the buffer, the first byte is taken from memory location zero, the second from location one, etc.

The buffers have certain common features included in their design: -

A/ A DATA OVERRUN signal (indicating that an attempt has been made to write more than 4096 bytes into the 4K buffer, or more than 1024 bytes into the 1K buffer.
B/ A BUFFER EMPTY signal (indicating that the buffer is empty).
C/ A DATA READY signal (indicating that there is valid data on the outputs of the buffer).

See appendix 'B' for technical description of 4K and 1K buffer memories.

APPENDIX 'A'.

Technical description of optional 32/63 byte FIFO buffer.

The buffer consists of two 32 by 8 bit FIFO integrated circuits (AM2812-ref. 5) cascaded together (I.E. placed end to end), plus control logic. A clear buffer signal, when inverted and applied to the master reset inputs of the two FIFO units, clears all control register bits and all the outputs of the FIFO'S to 'LOW'.

In the 32 byte mode, the data outputs from the first FIFO unit become the data outputs of the buffer proper. In the 63 byte mode, the data outputs from the second FIFO unit become the data outputs of the buffer proper. The position of the 32/63 mode switch determines which set of data outputs from the FIFO units are selected by the data selector.

Writing into the buffer is accomplished by applying a positive pulse to the 'PL' (parallel load) input of the first FIFO unit.

Reading from the buffer is accomplished by applying a positive pulse to the 'PD' (parallel dump) input, of the first FIFO unit if in the 32 byte mode ore, to the second unit if in the 63 byte mode.

The two FIFO units have been cascaded by driving the 'PD' input of the first unit with the 'IR' (input ready) output of the second unit, and driving 'PL' of the second unit with the 'OR' (output ready) output of the first.

With the jumper pin in position between the pin marked 'HOLD' and the pin marked '1/2', the 'HOLD' signal will be transmitted after 32 bytes have been written into the buffer, and will cease to be transmitted only when there are less than 16 bytes in the buffer. (This peculiarity has occured as a side effect of a technical shortcoming which arises when cascading AM2812 I.C.'S. It is explained at the end of this appendix).

With the jumper pin in position between the pin marked 'HOLD' and the pin marked '3/4' the 'HOLD' signal will be transmitted while there are 48 or more bytes in the buffer.

Because the AM2812 is a complex 'MOS' integrated circuit requiring three supply coltages, ov, +5v and -12v, it was necessary to include a DC/DC convertor in the design of the buffer, as there is only +5v available in the PDP 11 data links, With an input voltage of +5v, the UD5-12D5O gives output voltages of +12v and -12v.

Why 63 and not 64.

Logically, one would expect the buffer to be a 32/64 byte FIFO. The reason for only 63 bytes instead of 64 bytes is due to a technical shortcoming in the design of the 32 8-bit FIFO integrated circuits (AM 2812's) used on this card. A peculiarity occurs when two of the above mentioned I.C.'S are cascaded, in that when the second of the two units is full, its 'IR' (input ready) signal remains 'HIGH'. Clearly then, as this signal drives the 'PD' (parallel dump) input

of the first unit, the data from the first unit will not be dumped until the 'IR' signal of the second as become 'LOW' again. This 'IR' signal of the second unit will not become 'LOW' until the second unit has received a parallel dump command and an empty location has bubbled up from its output to its input.

Therefore, to summarize, when the second FIFO unit is full, the data last written into the second unit remains on the output of the first until an empty location bubbles up from the output of the second and 'IR' of the second becomes 'LOW' again.

Consequently when cascading N of these I.C.'S, the effective storage capacity is 31N+1 bytes.

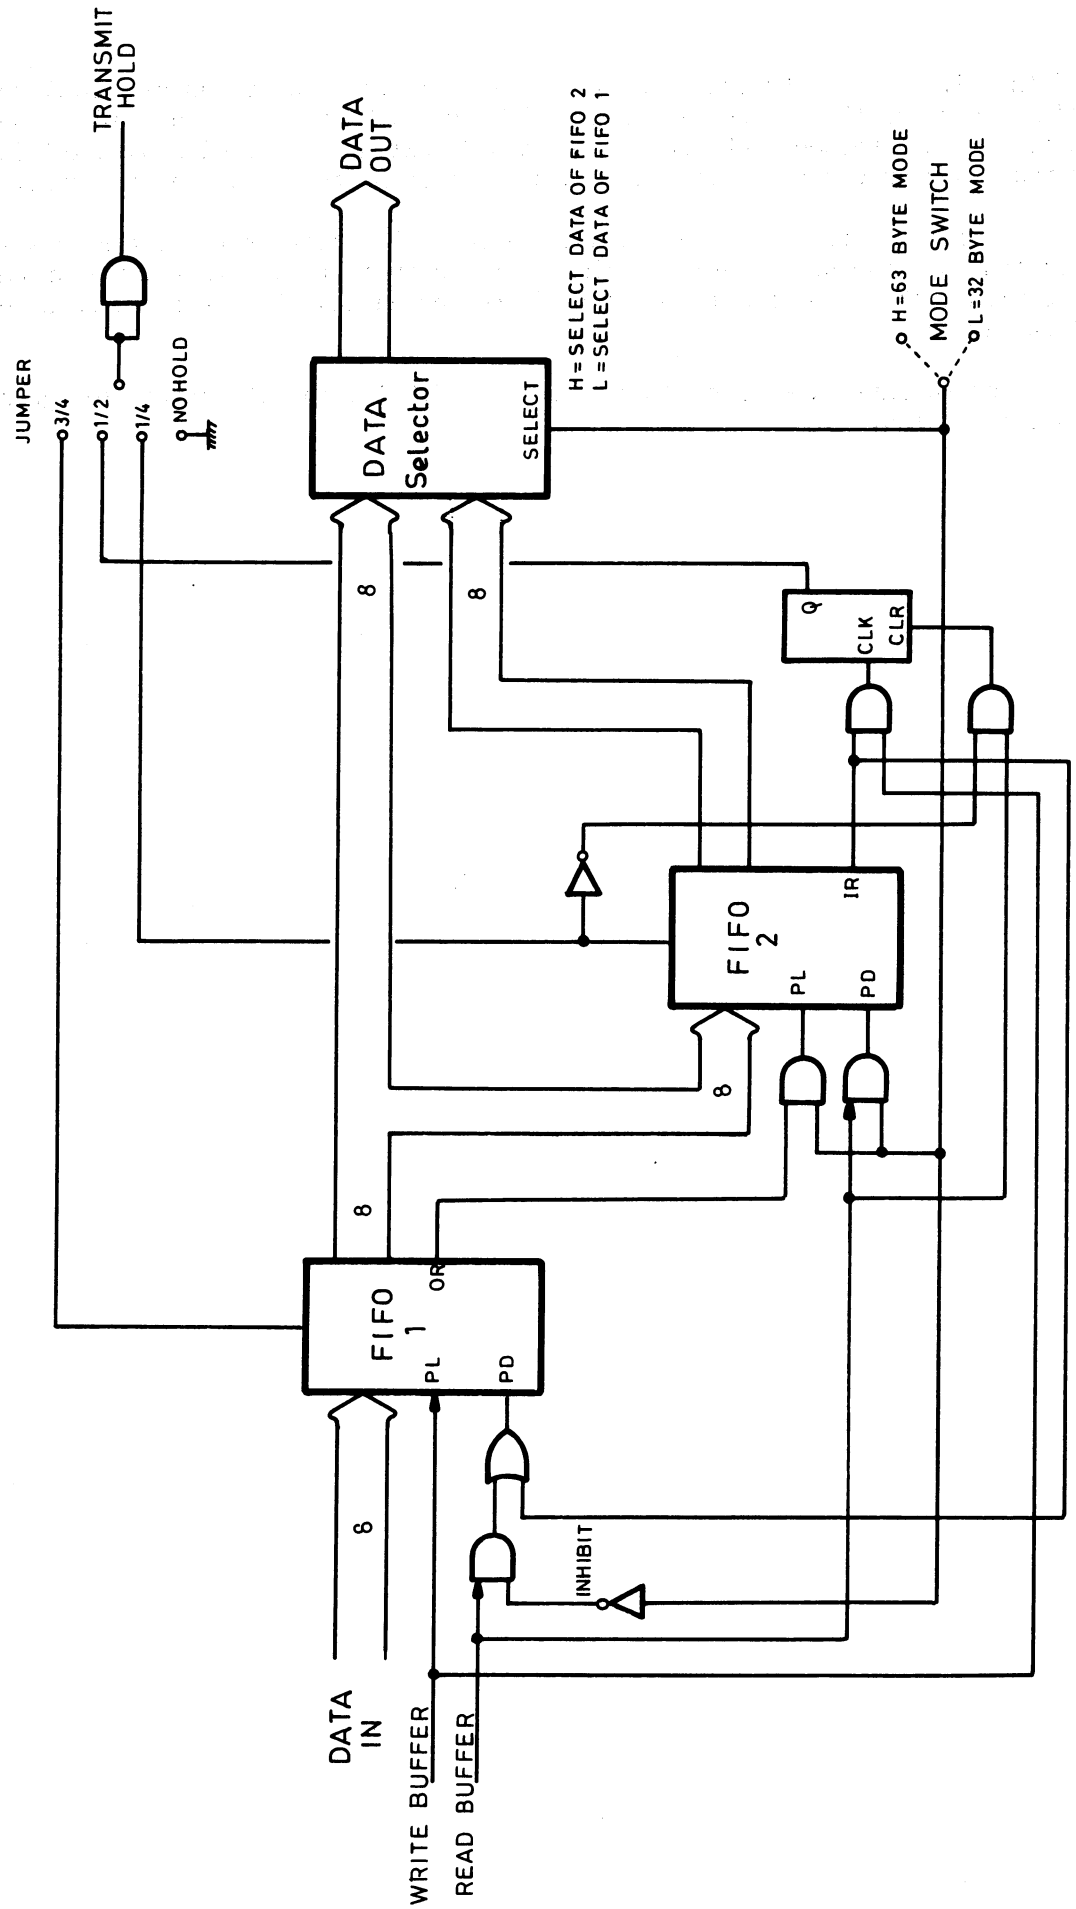

Fig. 2 — SCHEMATIC DIAGRAM FOR OPTIONAL 32/63 BYTE FIFO BUFFER

APPENDIX 'B'.

The term 'One shot (o/s) multivibrator' has been used in place of 'monostable multivibrator', in the following technical descriptions.

1/ Technical description of 4K byte buffer.

A clear buffer (CLB) signal clears the two binary counters to zero and resets all flip-flops. See schematic diagram of 4K buffer.

Writting into the buffer ia accomplished by the write buffer (RDS) signal (from the link serial/parallel convertor) firing a one shot multivibrator whose output is used to provide the chip select (NOT CS) signal for the SEMI 4200 integrated circuit (I.C). The address on the 'A00-A11' inputs of the memory I.C.'S is provided by the address counter (SN74161). And as the write buffer signal implies that stable data is on the inputs of the buffer, writing into the buffer is straight forward. (see fig. 3A).

Reading from the buffer is a little more complicated. The changing of the 'GO IN' signal from a low to a high level causes a pulse to be made which resets the address counter to zero. The same pulse to be 'CS RECOVERY TIME' o/s multivibrator whose outputs are used to make another pulse which performs two functions: -

A/ it sets the 'READ CS' flip-flop (whose output is used to provide the chip select signal in the read mode).
B/ it fires the 'WRITE CS' o/s multivibrator. This produces a pulse of 225ns duration, the trailing edge of which is used to set the "DATA READY" flip-flop.

A 'handshake' exists between 'DATA READY' and the 'NEXT BYTE' signal. Some time after 'DATA READY' has gone high, 'NEXT BYTE' will go high. This causes, 'DATA READY' to go low, which in turn causes 'NEXT BYTE',to go low again. The 'NEXT BYTE' signal also fires the 'CS RECOVERY TIME' o/s. The reason for this one shot is that the data sheet for the SEMI 4200 specifies that a minimum of 150ns should elapse between chip selects. The output of the o/s are used to perform 'A' and 'B' again. So the cycle continues until the buffer is empty.

The SN74193 up/down binary counter is used to keep a count of howmany bytes there are in the buffer. Both counters are clocked by the 'WRITE CS' o/s.

The data sheet for the SEMI 4200 specifies that the maximum time that the chip may remain 'selected' is 1 millisecond. A 'WATCHDOG' has been incorporated in the design to make sure that 1ms is never exceeded.

Consider the read mode, and in the sequence of events, the 'READ CS' and the 'DATA READY' flip-flops have been set. The next occurence should be that the 'NEXT BYTE' signal is sent to reset these flip-flops. If for some reason (e.g. hardware fault), 'NEXT BYTE' is not sent in the first milli-second after 'READ CS' was set, the 'WATCHDOG' one shot multivibrator would inhibit the 'NOT CS' input of the SEMI 4200 I.C.. This can be seen in fig. 5.

2/ Technical desciption of 1K byte buffer.

The control logic for the 1K byte buffer is identical to the 4K buffer, except that there is no 'WATCHDOG' and that the counters are clocked by the logical 'OR' of 'READ CS' and 'WRITE CS'.

The reasons for these differences are: -

1/ There is no limit to the time that the Fairchild 2102 can remain 'SELECTED', consequently no need for a 'WATCHDOG'.

2/ The SEMI 4200 allows the address to be changed any time after 100ns of the chip select pulse have expired. The Fairchild 2102 does not allow the address to be changed until the end of the chip select pulse.

see schematic diagram of 1K buffer.

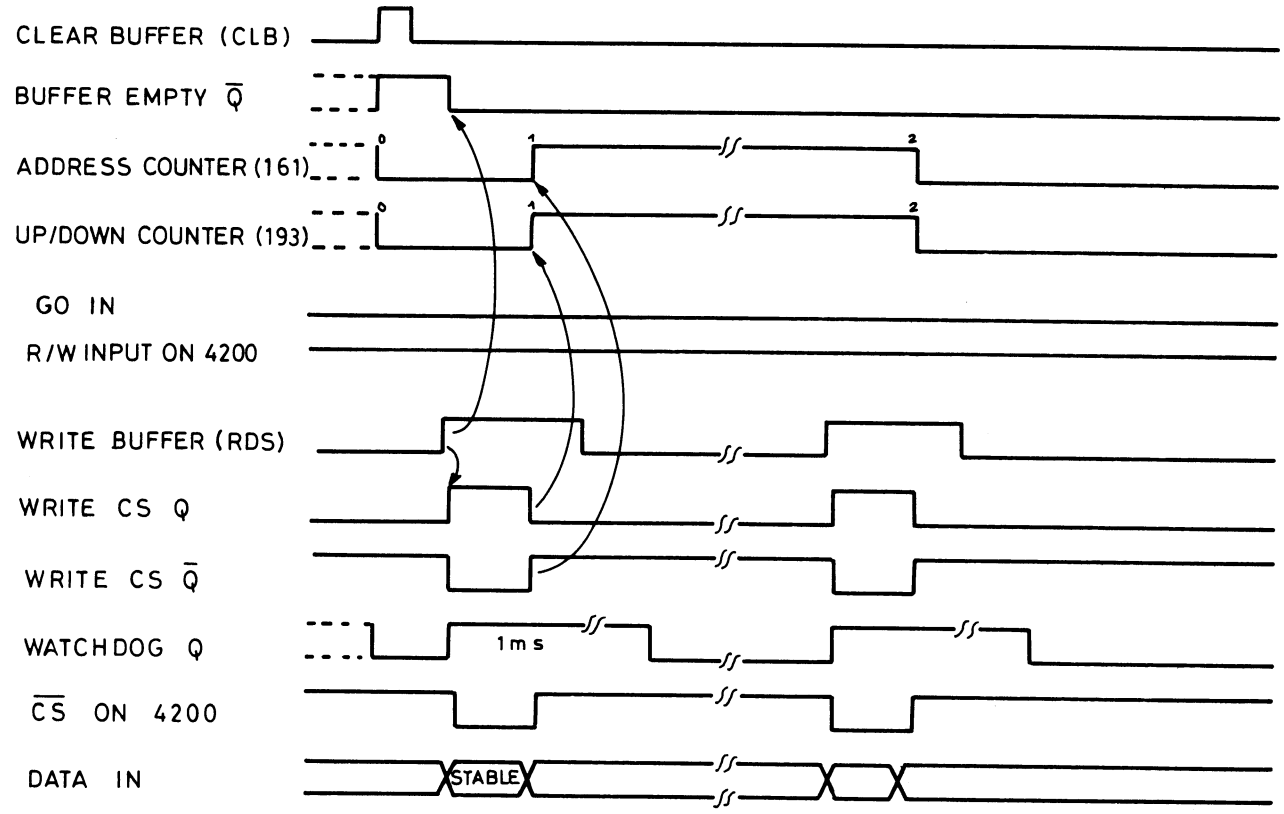

Fig. 3A — WRITING INTO 4K BUFFER

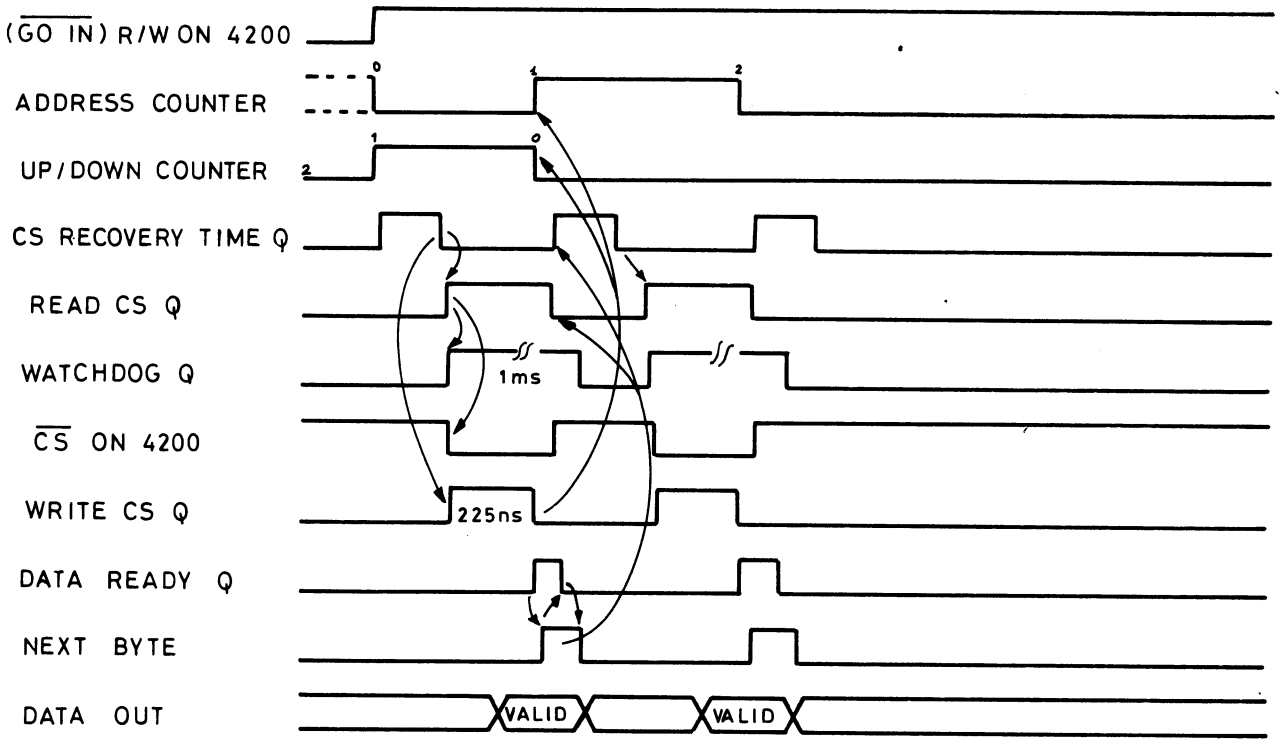

Fig. 3B — READING FROM 4K BUFFER

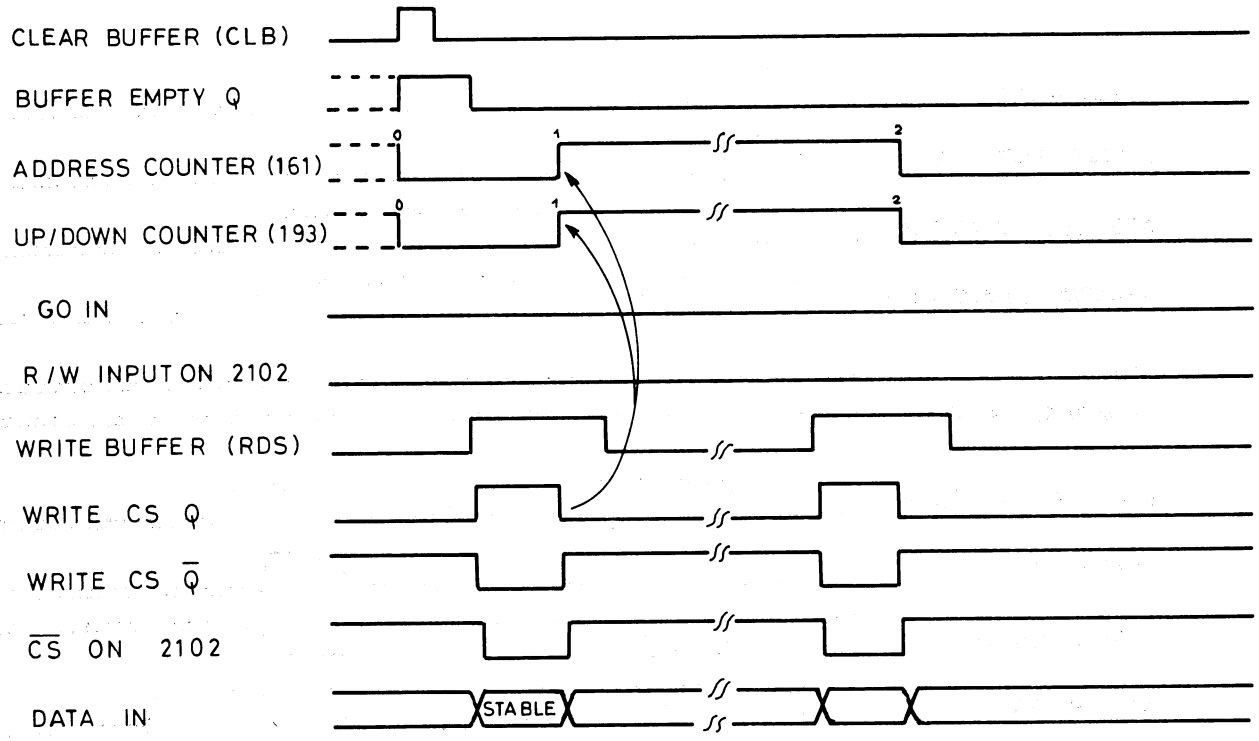

Fig. 4A WRITING INTO 1K BUFFER

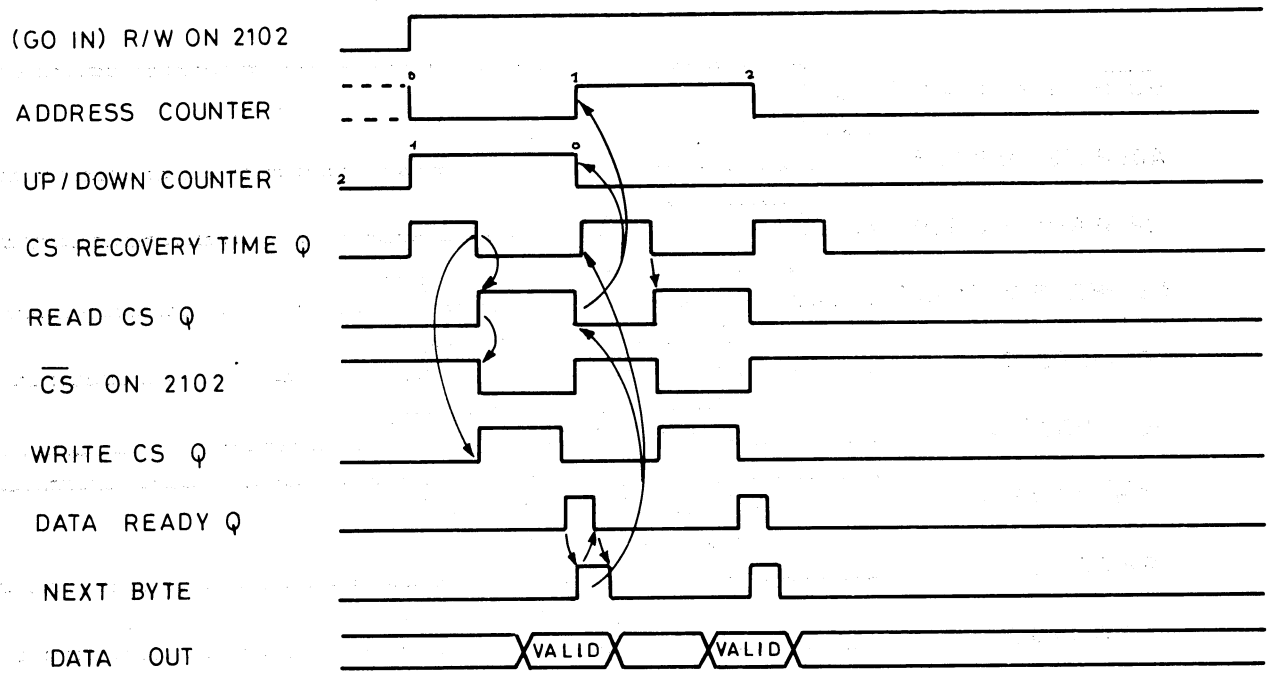

Fig. 4B — READING FROM 1K BUFFER

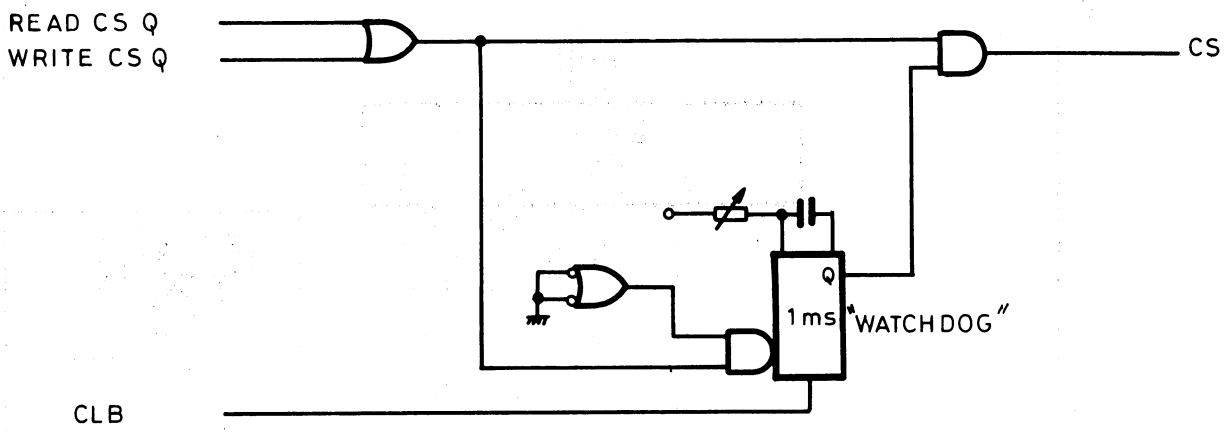

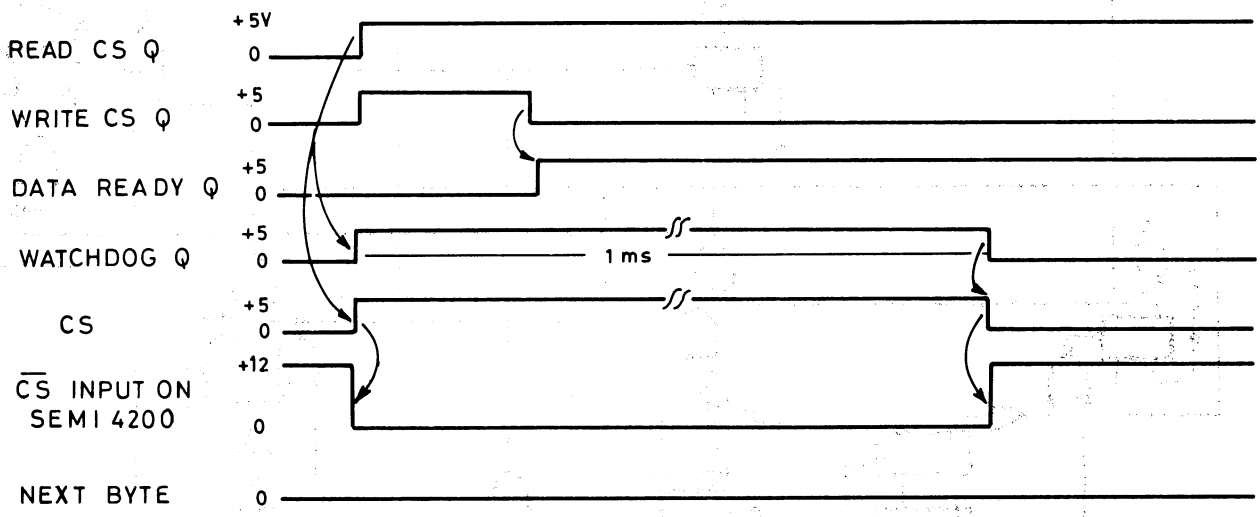

Fig. 5 — "WATCHDOG" IN ACTION ON 4K BUFFER

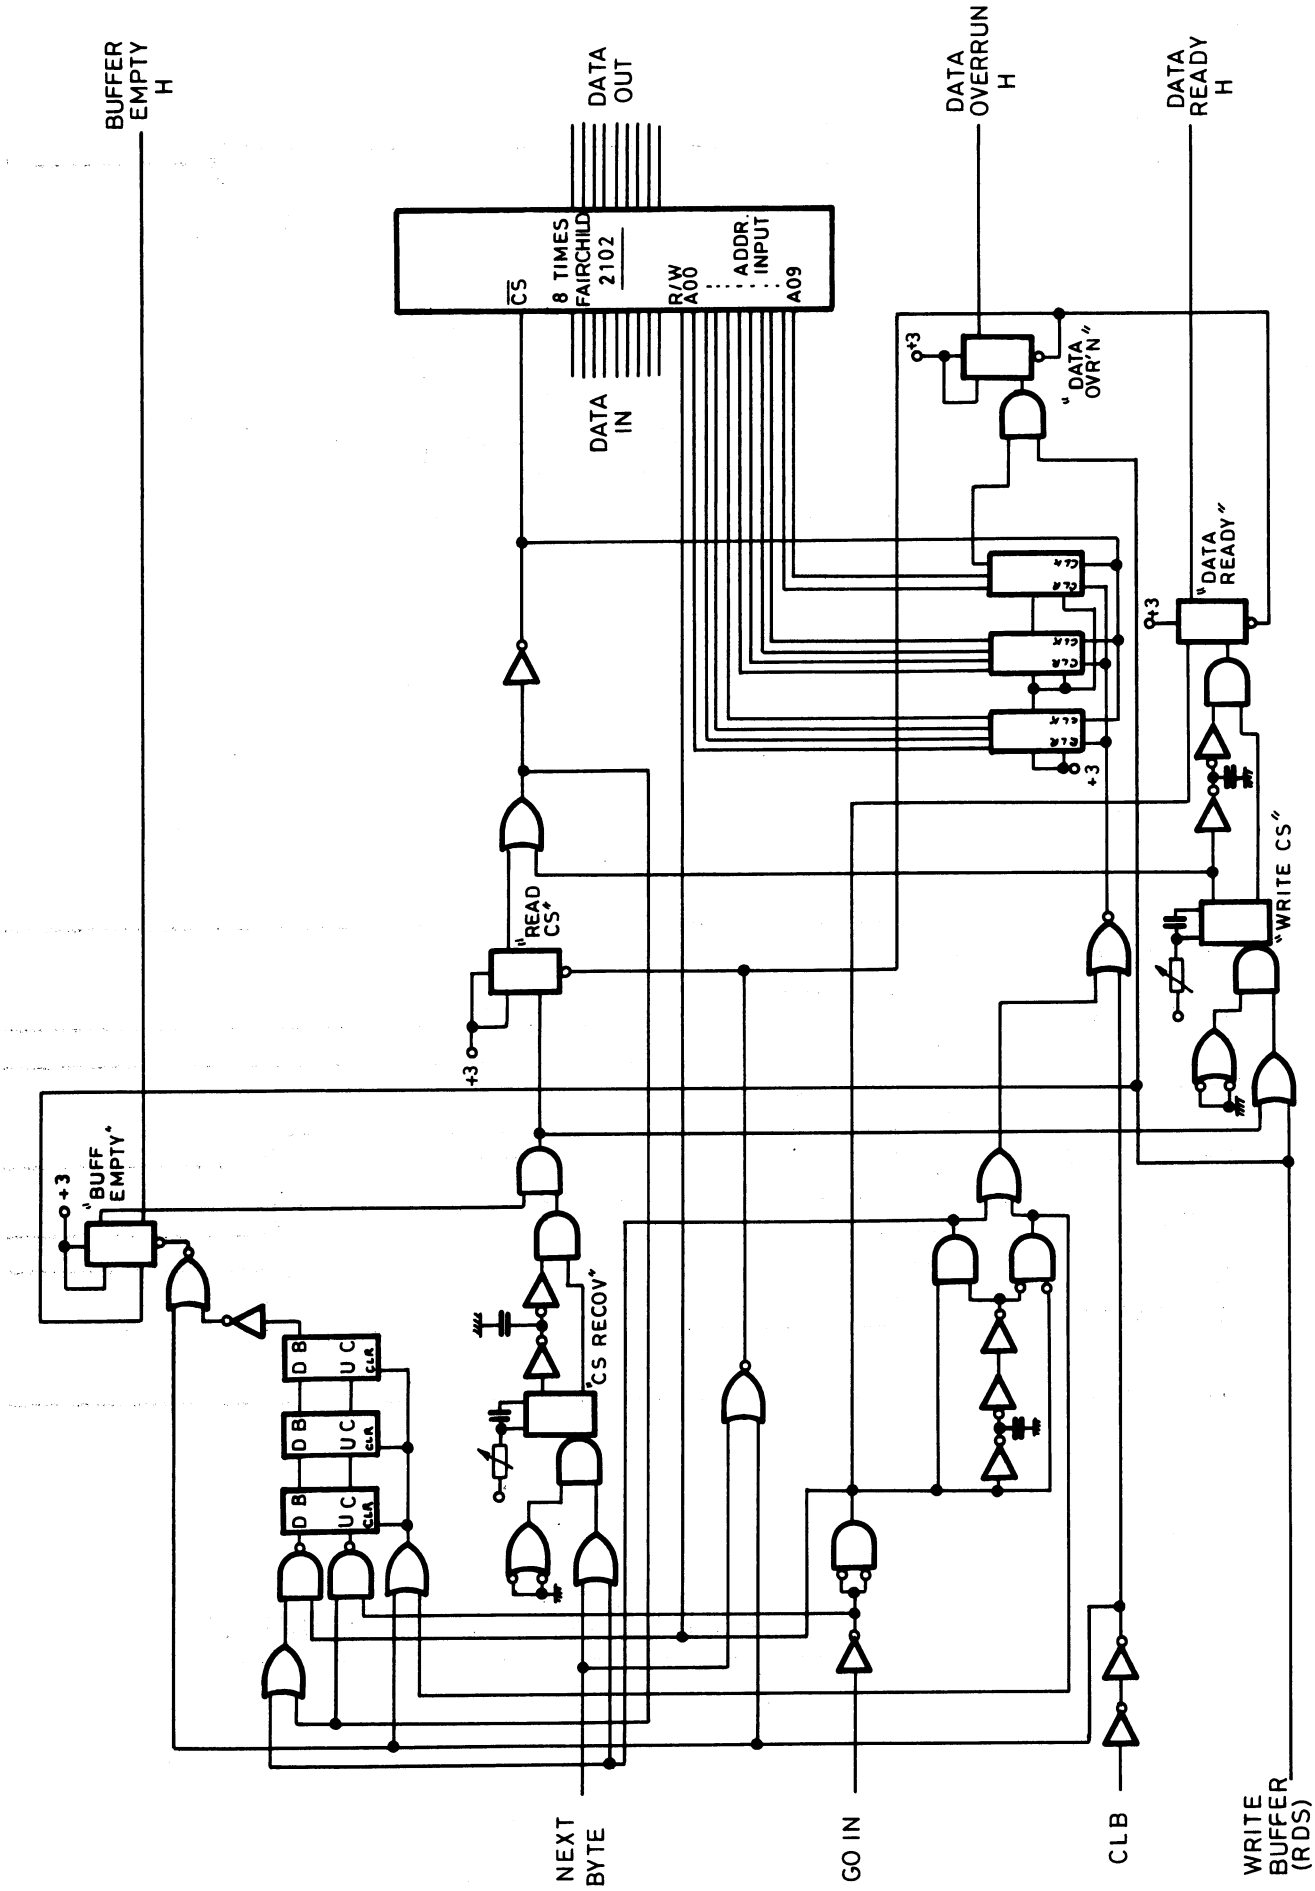

SCHEMATIC DIAGRAM FOR 1K BY 8 BIT FIFO BUFFER

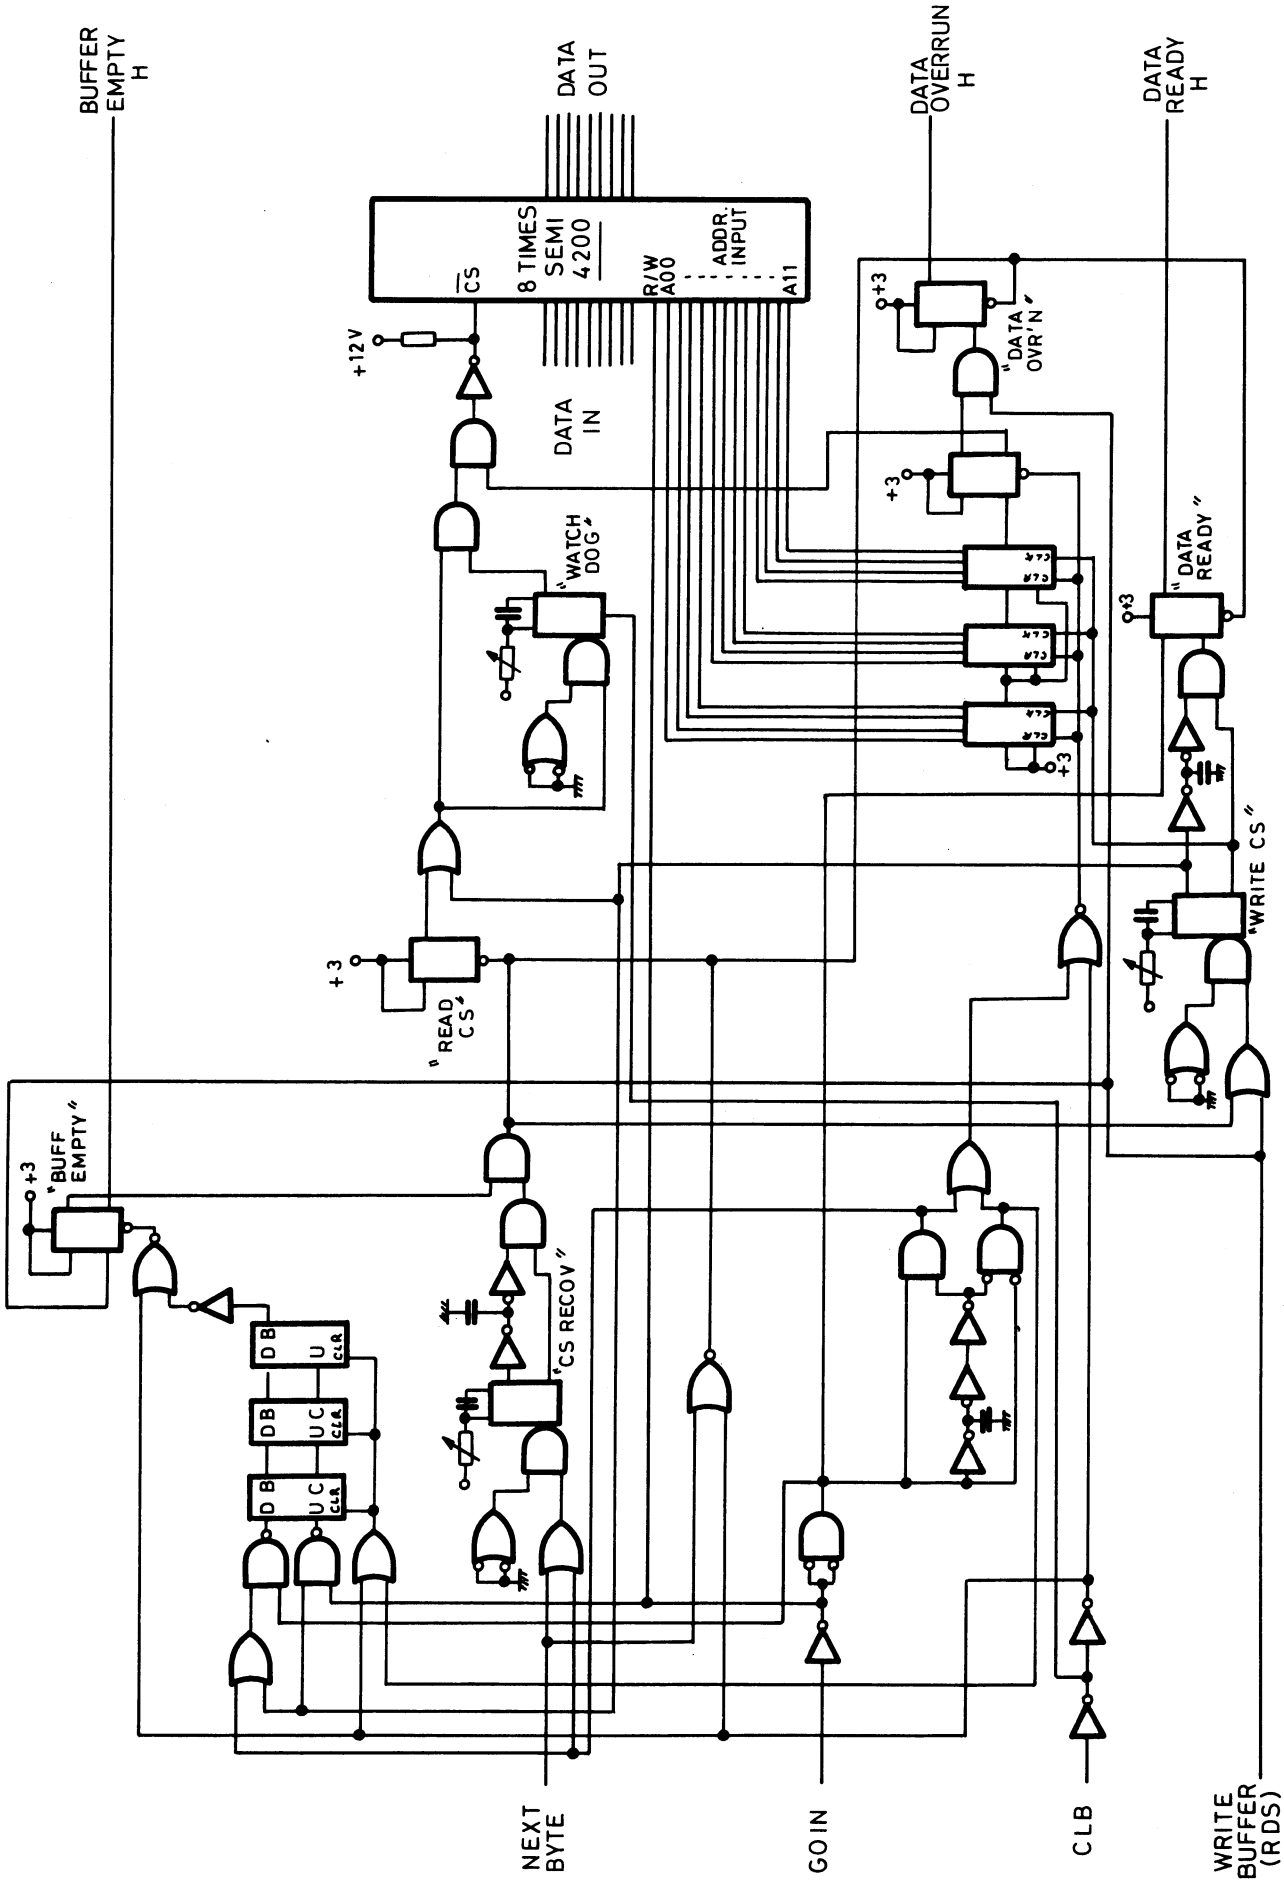

SCHEMATIC DIAGRAM FOR 4K BY 8 BIT FIFO BUFFER

3/ Electrical details and power supplies.

The 'Q' outputs of 'READ CS' and 'WRITE CS' must be inverted and 'pulled up' to +12v before being applied to the 'CHIP SELECT' input of the memory integrated circuits.

The SEMI 4200 requires three power supplies of +12v, +5v and -5v. The Fairchild 2102 requires a single power supply of +5v.

4/ New developments.

There is now available a static MOS RAM organised 4096 by 1 bit which requires a single +5v supply. There is also a static MOS RAM organised 1024 by 4 bits. Consequently a 1K byte memory would only require two 'memory' I.C.'S. Both these memories are made by 'Advanced Micro Devices'. The 4K memory is the 'AM 9140', and the 1k memory is the 'AM 9130'.

Acknowledgements.

I would like to thank H.E. Davies, C.F. Parkman and C. Verkerk for their advice and experience in technical problems, and the writing of this paper.

References:

1. Omega Development Note HW.21 by C.F. Parkman.
2. Cern yellow report 74-9, 'Meeting on technology arising from high Energy Physics'. Vol. 2 (2nd ed.).
3. Data Sheet for Semi 4200 Integrated Circuit.
4. Data Sheet for Fairchild 2102 Integrated Circuit.
5. Data Sheet for Advanced Micro Devices AM2812 Integrated Circuit.