Marinchip Systems M9900 64K RAM

Theory of Operation and Configuration Guide

by Fred Camerer and John Walker

Revised August 1980

Marinchip Systems 16 St. Jude Road Mill Valley, CA 94941 ■ (415) 383-1545

Marinchip 64K RAM Theory and Configuration Guide

1. Introduction

The M9900 64K RAM board is a high speed random access read write memory designed expressly for use with the M9900 CPU. It uses industry-standard 16K dynamic RAM circuits to achieve its goals of high speed, reliability, low cost, and minimal power consumption. The 64K RAM board operates with the M9900 CPU at either 2 Mhz or 3 Mhz, with fully transparent refresh which guarantees reliable data storage without imposing a speed penalty on the CPU.

The M9900 64K RAM takes advantage of the M9900 CPU's 16 bit data transfer mode. This mode permits the CPU to transfer data on a word by word basis rather than the normal 8 bit byte transfers used by other S-100 memories. Use of a 16 bit transfer and the elimination of redundant bus cycles allow memory transfers to be completed in one third the time when the M9900 64K board is being accessed, as compared with conventional S-100 memories. This results in an overall doubling of system performance when the M9900 64K board replaces regular S-100 memories.

The M9900 64K RAM features programmable memory management hardware. Pages of memory on the board may be enabled and disabled in 4K byte increments. The initial configuration of the board is set by switches on the board, and the memory map may be changed at any time by software. When used in conjunction with the Marinchip Multi-User Network Operating System (NOS/MT), this allows system memory configurations as large as 16 megabytes.

2. Configuration and installation

The M9900 64K RAM board is shipped fully assembled and tested. Prior to installation in your system, you need only set up the configuration for the board in the DIP switches.

2.1. Bank select port address

The 8 position DIP switch at the bottom of the board (parallel to the edge connector) labeled "Bank Select Port" sets the I/O address for the port which controls selection of memory blocks on the board. The leftmost switch is the most significant bit (2**7), and the rightmost switch is the least significant bit (2**0). The address is set up in the switches by setting the switch for each zero bit to "on" (or "closed"), and the switch for each one bit to "off" (or "open").

For example, to set the bank select port address to 85 hexadecimal, the switches would be set (from left to right):

off on on on off on off

Marinchip Systems recommends that the bank select ports on memory boards be assigned ascending addresses starting at 80 hex, with 80 being the first 64K board (the one enabled when RESET is pressed, see below), 81 being the next board, etc. If these addreses conflict with I/O devices used in your system, they may be changed with impunity, as long as the memory configuration for NOS/MT is changed to agree with the addresses on the boards. Users of the Disc Executive and Release 1 of the Network Operating System, which do not use the bank select feature of this board, need only make sure that the bank select port address selected does not overlap any other I/O address in the system.

2.2. Initial bank selection

The pair of DIP switches at the right of the board select which banks of memory will be selected when the computer is reset. This selection will stay in effect until the first output is done to the bank select port. The legend at the right of the board indicates the effect of the switches. Each switch controls initial selection of a 4K byte area. The top switch controls the region from 0000 to 0FFF, the next switch 1000 to 1FFF, etc., to the bottom switch which controls F000 to FFFF. If a switch is "off" (or "open"), the corresponding 4K block will be enabled (in the memory map) when the system is reset. If a switch is "on" (or "closed"), the block will be disabled and will not respond to memory read or write signals from the computer.

In a system consisting of the M9900 CPU, the PROM/RAM board, and the 64K board, with no other memories or memory-mapped I/O devices, the initial bank selection switches are normally all turned off, except for the switch controlling the page F000 to FFFF (the switch closest to the edge connector), which is on. This setting enables 60K of memory from 0000 to EFFF, and disables the last 4K in which resides the M9900 CPU I/O area, and the PROM and RAM of the PROM/RAM board.

Systems which have other memory boards (for example, a separate PROM board or static RAM used for DMA transfers), or memory mapped devices

such as video interfaces, should set the initial bank selection switches to disable the areas of memory occupied by those devices.

In systems with more than one M9900 64K board, the normal configuration is all 60K enabled on the first board, and all storage initially disabled (all switches on) on all the other boards. After being loaded from disc, the operating system assumes software control of the memory bank selection and the initial switch settings become irrelevant.

2.3. What does the LED indicate?

The most commonly asked question about the M9900 64K RAM is the significance of the LED (light emitting diode) on the right side of the board. This component is present primarily for its electrical properties rather than as an indicator, but it will be illuminated as long as the bank selection set in the initial bank select switches is in effect. As soon as this initial selection is overridden by output to the bank select port, the LED will be extinguished and will stay out until the computer is reset and the switches become effective again.

3. Bank select programming

This section of the manual explains how to control the bank selection of the M9900 64K board from software. In most applications of the board, the bank selection will be done by the operating system without the user's knowledge. This section is a guide to those developing their own operating systems or programs which access multiple memory banks, and will also aid in understanding how to configure NOS/MT for use with the 64K RAM board.

The bank select port is a 16 bit parallel output port, implemented as an S-100 I/O port. The M9900 CPU is able to perform 16 bit transfers on I/O cycles as well as on memory cycles, and the 64K RAM board takes advantage of this to permit the entire memory map to be set by one output instruction. The output port controlling the bank selection on the board is set by the bank select port DIP switch as described in the configuration section earlier in this manual. This port is written by a program by moving data to the memory mapped address associated with this port. The memory mapped address is computed by the formula:

MMA = 0F200+IO*2

where IO is the I/O address set in the switches. For example, if the bank select port were set to 80 hex, its I/O address would be:

0F200+080*2 = 0F200+0100 = 0F300

Hence, if one wished to store the value in register zero (R0) into the bank select port, one would use the instruction:

```
MOV     R0,0F300
```

The value loaded into the bank select port is a 16 bit value with each bit controlling selection of a 4K byte block of memory. The least significant bit 2**0 (1) controls the bank 0000-0FFF, 2**1 (2) controls 1000-1FFF and so on, up to 2**15 (8000 hex), which controls the page from F000-FFFF. The new bank selection will take effect immediately upon completion of the I/O transfer to the memory board, so the user should be careful to execute the code switching banks from an area of memory not affected by the change in memory map.

The only restriction in programming the bank select port is that at no time may the same address range be enabled in more than one memory board. If this is done, no damage will occur, but the operation of the system will be unpredictable and undesirable.

4. Theory of operation

This chapter of the manual explains the detailed operation of the M9900 64K RAM board. This information is intended to be used in conjunction with the schematics to gain a better understanding of the design and operation of the board, and as an aid to troubleshooting.

4.1. Overview

The M9900 64K RAM is a 64K byte bank switched random access memory organised as 32K sixteen bit words. The timing of the memory board is designed to take full advantage of the M9900 CPU's sixteen bit mode of operation whereby a complete memory cycle requires only two clock cycles instead of the six clock intervals required by two normal S-100 8 bit transfers. The memory may still be intermixed freely with eight bit memory with the only restriction being that two boards may not simultaneously occupy the same address space.

4.2. Bank organisation

The initial bank enable/disable choice for each 4K byte page of the memory is set manually via a row of sixteen independent switches. An additional row of eight switches allows an I/O address to be used to dynamically adjust the bank selection. This I/O port is sent a sixteen bit mask in a single I/O transfer, where each bit in the mask enables a corresponding page in the memory of that board.

4.3. The memory array

The memory itself is a 32 chip array of 4116-type 16K dynamic memory circuits, generally acclaimed to be the most popular LSI component ever. This circuit requires a few well chosen timing signals with which the circuit has been found to be very reliable.

There are three power connections (+12, +5, and -5) plus ground, and three control signals are required. As used in this application, the data in and data out pins are tied together and used as a single signal (in each of the sixteen bit positions, two chips share a single data bus). The remaining seven pins of each chip receive seven common addressing signals as provided by the 3242 address multiplexor/refresh counter circuit.

4.3.1. Timing signals

The three timing signals are RAS-, CAS-, and WRITE-, all low active (the trailing minus sign indicates a low active signal in this manual). Row Address Strobe (RAS-) is used at a negative transition to latch in each chip seven bits called the row address. While the actual circuit inside the chip may vary, one may think of these address bits as selecting one of 128 rows of 128 memory cells each. This process of activation must be repeated for each row with a specified minimum repetition rate. If this "refresh" ritual is not performed devotedly, the memory will lose the data stored therein. The M9900 64K RAM board performs many more refresh cycles than are required by the specifications of the memory circuits; this causes no harm other than a small increase in power consumption. RAS- is

required to be active low for a specified interval. If Column Address Strobe (CAS-) remains high and RAS- returns to the high state, the memory performs a refresh cycle. The data-in and data-out pins are not used, and data stored in the memory remains unchanged.

For every read and write cycle, CAS- goes low for only one row of sixteen chips or the other (depending on the most significant address bit from the processor). The other row of chips will see only RAS- and hence will simply perform a refresh cycle.

After RAS- has captured the row address in each chip, the seven address lines are changed to present the remaining seven bits of the data address, thus specifying one of the 128 bits in the already selected row to be read or written. As CAS- goes low these remaining address bits are latched and the address specification is complete.

4.3.2. Read / write determination

If the WRITE- input to each chip is active (low) before and throughout the duration of CAS- low, then the addressed memory location will be written with the value presented on the data-in pin.

If the WRITE- input remains inactive (high) throughout the CAS- low interval, the data held in the selected memory cell is presented on the data-out pin (the data stored in the memory is unaffected).

At the conclusion of a cycle CAS- must be high a specified time before RAS- goes high, and thereafter RAS- must remain high an interval before the next row address may be strobed.

4.4. Address multiplexing

The 3242 circuit serves both to multiplex the 14 bits of address from the bus to the 7 address pins on the memory circuits, and to generate the refresh addresses and selectively present them to the memory circuits. When the REFEBL signal to the 3242 is inactive, the address from the bus will be sent to the memory chips. Whether the high 7 bits or the low 7 bits will be presented to the chips is controlled by the ROWEBL signal. When REFEBL is active, the internal 7 bit refresh counter within the 3242 will be placed on the output pins. Each negative transition of the COUNT- signal increments the refresh address counter.

4.5. Refresh timing

Refresh timing is simply derived from the main clock signal of the S-100 bus, PH2(24). At every falling edge, RAS- is driven low. The low interval can be extended with the setting of the 74LS121 one-shot circuit as may be required for operation at bus speeds faster than 2 Mhz. At 3 Mhz, 200 nanosecond memory chips require RAS to be 200 to 210 nanoseconds for proper refresh operation. This should be adjusted by observing the RAS- signal on pin 8 of U8 on an oscilloscope, and adjusting the trim pot until RAS- is low for 205 nanoseconds. A board so adjusted will operate at either 2 Mhz or 3 Mhz without problem. If only 2 Mhz operation is required, the one-shot circuit U6 may be removed entirely (simply pulled from its socket), in which case the setting of the trim pot is irrelevant and may be ignored.

4.6. Memory cycle timing

If the RUN flipflop sets when clocked by PSYNC(76) from the bus, it indicates the start of a memory cycle. REFEBL is deactivated at the 3242 so that the row address will be gated from the S-100 address lines A1(80) to A7(83). The CPU is informed that the current cycle is a sixteen bit operation, and it abruptly changes its state from the middle of T1, the SYNC cycle, to the middle of T2, the second last bus cycle for the current memory operation. Bus clock PH1(25) sets CAS1 which disables ROWEBL in the 3242, thus presenting the column address to the memory chips from bus lines A8(84) to A14(86). Note that address line A0(79) (which in eight bit memories selects the odd or even byte) goes completely unused here. Address line A15(32) is used to select which row of sixteen chips is to receive CAS-. When clock PH1(25) falls, CAS2 and the nearly identical CASX are set driving either CASHI- or CASLO- active, thus activating one set of sixteen memory chips.

Meanwhile, signal DBIN is set if the board is involved in a read cycle, and WRITE- is derived from SWO-(97) if the cycle is a write to memory. Signal DBIN enables the data read from the sixteen data cells to be copied to the data in and data out buses of the S-100 connector as required by the M9900 CPU. The data remains on the bus until the subsequent rise of PH1(25) which clears RUN, which in turn clears CAS1 and DBIN if set. CAS2 is also cleared, ending the column address strobe to the memories. The subsequent fall of PH1(25) clears CAX, and the rise of PH2(24) concludes RAS-, at which time the memory is idle, or RUN is again set for a subsequent memory operation.

Note that RUN and CAX conspire to inhibit the signal COUNT-. In this way a count pulse only occurs following a refresh cycle, and the refresh counter thus reliably cycles through all the required refresh addresses.

4.7. Bank selection

The bank selection circuitry makes the decision whether to set RUN at the cycle start. The four most significant address bits are used in a 74150 data selector so that if the selected input is high, and if a signal from the CPU called IO(96) is inactive low, then the signal SEL- is low and RUN will be clocked set by PSYNC(76) as already described.

The flipflop IBS (Initial Bank Select), which is directly set by the Power On Clear bus signal, POC-(99), when the reset button on the mainframe is pushed or the CPU executes an RSET instruction, indicates initial bank select mode. This flipflop when set provides base drive for the transistor. (The LED serves as a voltage drop, and also indicates when glowing that the flipflop is still set). Collector current in the transistor circuit pulls the manually switched 74150 inputs low through isolation diodes, and this disables each bank of 4K bytes corresponding to a switch in the on state.

IBS is reset by an output operation to the bank select port on the interface. U11 and U12 are arranged to give the output U11 pin 6 high whenever the I/O address on the S-100 bus matches that set in an 8 position DIP switch on the board. An I/O operation with the selected

address activates SIXTN-(60), just as RUN does during memory operations, to inform the CPU that the full sixteen bit data path is active. SOUT(45) and PWR-(77) from the CPU clock a sixteen bit selection mask into U16 and U18 and also clear IBS.

With IBS cleared, the initial bank select switches have no effect, and the selected inputs of the 74150 selector are driven by the now-enabled 74LS373 outputs. Additional output operations to the bank select port will continue to adjust the bank select vector, and IBS will remain off.

4.8. DMA inhibit logic

Signal PHLDA(26) when active prevents PSYNC(76) from setting RUN. The result is that the memory does not make any attempt to respond to any Direct Memory Access (DMA) requests. DMA will not be possible to and from this memory for that reason. DMA access is prevented because existing DMA devices require 8 bit data transfers which cannot be accomodated by the M9900 64K RAM board, and also because many DMA devices generate incompatible bus signals which might cause loss of data in the 64K board, even if it were not being referenced. The DMA inhibit logic allows the board to be unaffected by DMA transfers done to and from 8 bit memory installed in the same machine as the 64K RAM board. In order to use a DMA device in a system with 64K RAM boards, a small (4K) 8 bit memory is usually installed, with the corresponding bank disabled in the 64K RAM board. All DMA traffic is done to that memory area, then copied to its ultimate destination in the 64K RAM board. The 64K RAM board was designed to provide the optimum performance with the M9900 CPU, and forcing it to accomodate DMA devices would have required a substantial (50%) sacrifice of performance, and a large increase in parts count, complexity, and hence cost. The tradeoff chosen maximises CPU speed at the cost of somewhat greater overhead when using DMA devices.

4.9. About alpha particles

Currently many people are concerned about the "alpha particle problem" in 16K RAM chips. As is frequently the case when technical matters impinge on the arena of marketing, there is a large body of comment about this topic which tends to reduce, rather than increase, one's understanding of the problem. What follows is our own view of the situation.

Alpha particles are helium nuclei (two protons and two neutrons) which have not yet settled down to a sedate existence in the company of two electrons, but rather are still zipping around with relatively large energies. These particles are generated by radioactive decay of heavy nuclei such as uranium, plutonium, and thorium. Since thorium is fairly common in the earth's crust, atoms of it tend to find their way into the ceramic material used to package ICs. When a thorium atom decays, it emits an alpha particle. If this particle happens to slam into the silicon of an IC, it will create a large number of electron-hole pairs, which will tend to be swept into active regions of the semiconductor. In older parts, the geometry of transistors was large enough and the charge distinguising ones from zeroes was high enough that the impact of an alpha particle did not make enough difference to cause an error. When 16K RAMs started to be applied in

systems, persistent "soft errors" were observed that could not be explained by conventional error mechanisms. These errors were eventually traced to alpha particles impacting in either bit cells, on the sense lines connecting the bit cells to the sense amplifiers, or in the sense amplifiers themselves. The frequency of alpha-induced soft errors depends upon the design of the chip and the alpha flux emitted by the packaging material in which the chip is enclosed. If you take the measured soft error rates of current chips and calculate the probability of error anywhere in a typical memory of 64K to 256K bytes, you find that a soft error will occur, on the average, once every one to three years.

Now, ideally, there should be no soft errors at all, but in general this error rate is acceptable in most applications. Interruptions due to power failure are typically far more common than alpha errors, and almost always far more serious. Typical disc error rates are much higher.

Initially, the alpha problem was thought only to affect dynamic RAMs because of their storage of data as charge on a capacitor, but a recent report from Japan indicates that 4K static RAMs are also vulnerable to alpha errors. Hence, the alpha problem is not a reason to choose static memory over dynamic memory. If one wishes to avoid alpha hits, the only available path is to use memories built from 4K or smaller dynamic chips, or 1K static chips (although current evidence is that 4K statics are pretty safe, as well), or to design single bit error correction into the memory. Neither of these options leads to an inexpensive bulk memory.

This discussion of alpha induced soft errors is intended to shed some light on the problem, not convince you that there is no problem. We feel that the problem is not worth worrying about since its maximum incidence is so much less than that of numerous more serious problems such as power outages. At this time, all major semiconductor manufacturers are investigating the problem and testing greatly improved chip designs and packaging materials. Already, new ceramics have been developed which reduce alpha flux close to the lowest level that can be measured by existing instruments. Consequently, this not very serious problem will soon be getting even less severe.

NOTE: ALL UNUSED S/R INPUTS TIED TO VCC
ALL LS241-EACH SHEET-ALWAYS ENABLED

REVISIONS			
LTR	DESCRIPTION	DATE	APPROVED

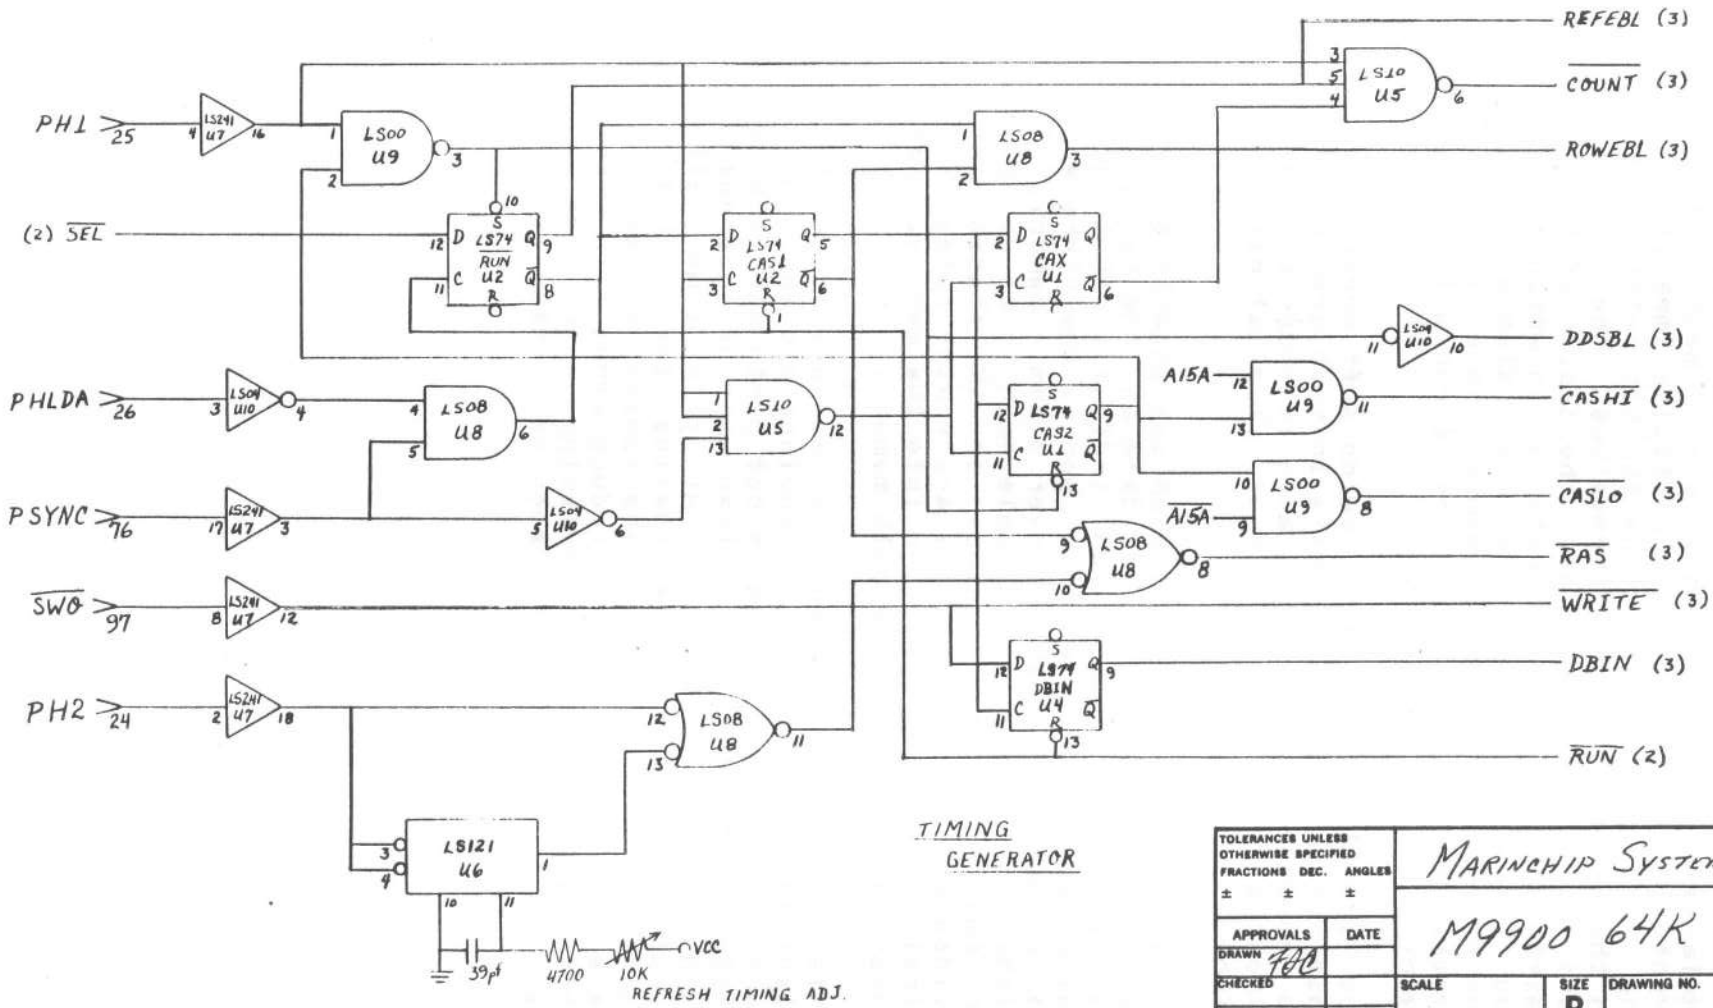

TIMING GENERATOR

TOLERANCES UNLESS OTHERWISE SPECIFIED		
FRACTIONS	DEC.	ANGLES
±	±	±

APPROVALS	DATE
DRAWN TAC	
CHECKED	

MARINCHIP SYSTEMS

M9900 64K RAM

SCALE	SIZE	DRAWING NO.
	B	

DO NOT SCALE DRAWING — SHEET 1 OF 3

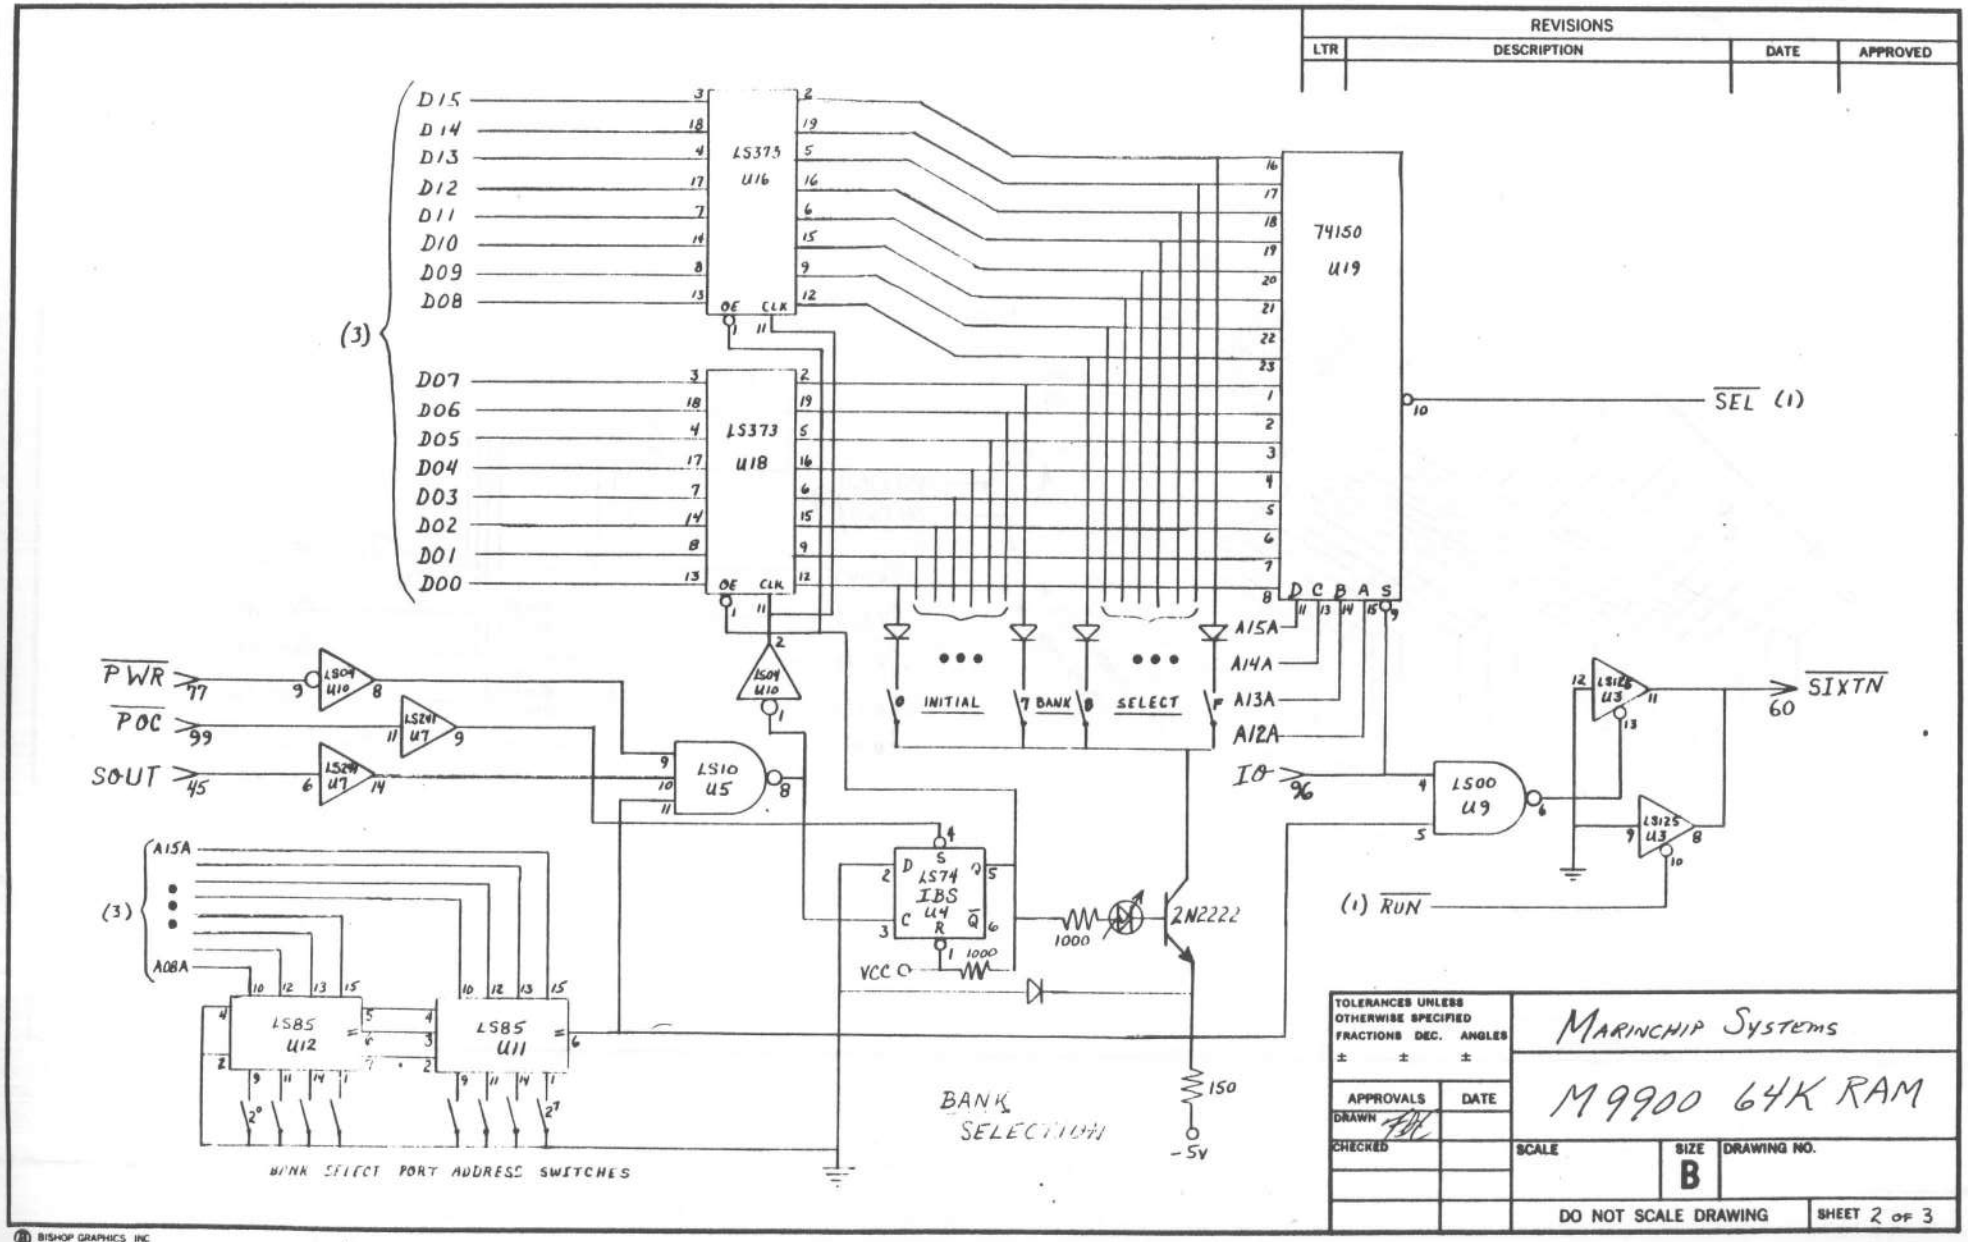
REVISIONS
LTR
DESCRIPTION
DATE
APPROVED
D15
D14
D13
D12
D11
D10
D09
D08
(3)
LS373
U16
D07
D06
D05
D04
D03
D02
D01
D00
LS373
U18
OE
CLK
74150
U19
SEL (1)
A15A
A14A
A13A
A12A
INITIAL
BANK
SELECT
PWR
POC
SOUT
LS04
U10
LS241
U7
LS10
U5
IO
LS00
U9
LS125
U3
SIXTN
60
RUN (1)
LS74
IBS
U4
2N2222
1000
VCC
150
-5V
A08A
LS85
U12
LS85
U11
BANK SELECT PORT ADDRESS SWITCHES
BANK SELECTION
TOLERANCES UNLESS OTHERWISE SPECIFIED
FRACTIONS DEC. ANGLES
APPROVALS
DATE
DRAWN
CHECKED
MARINCHIP SYSTEMS
M9900 64K RAM
SCALE
SIZE
B
DRAWING NO.
DO NOT SCALE DRAWING
SHEET 2 OF 3
BISHOP GRAPHICS INC
REORDER NO. 20521

REVISIONS			
LTR	DESCRIPTION	DATE	APPROVED

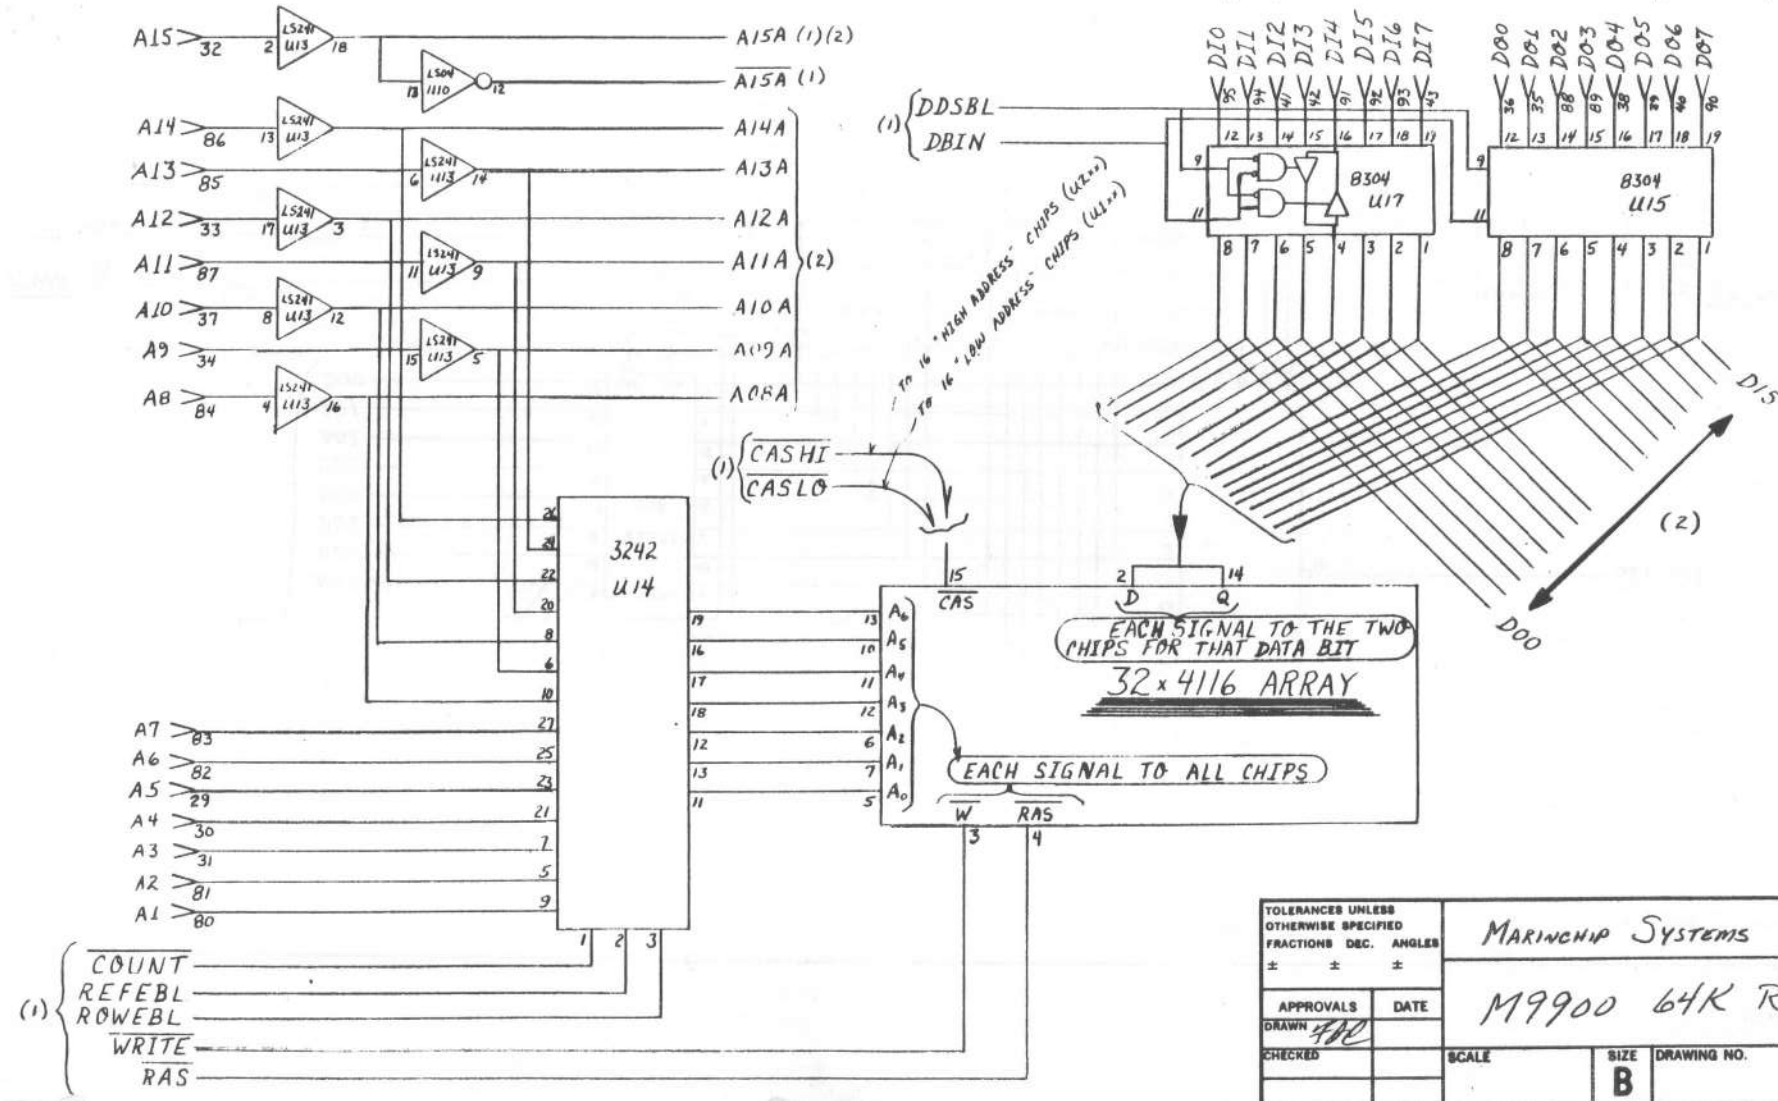

TOLERANCES UNLESS OTHERWISE SPECIFIED		
FRACTIONS ±	DEC. ±	ANGLES ±

APPROVALS	DATE
DRAWN TOR	
CHECKED	

MARINCHIP SYSTEMS

M9900 64K RAM

SCALE	SIZE	DRAWING NO.
	B	

DO NOT SCALE DRAWING — SHEET 3 OF 3