M-H8

64K BYTE MEMORY for the HEATHKIT H8 COMPUTER*

Trionyx Electronics
P.O. Box 5131
Santa Ana, CA. 92704

November, 1979

Single Card Memory Uses 16K Dynamic RAM Integrated Circuit

Sophisticated Design Features Transparent Refresh

INTRODUCTION

The M-H8 is a single card, 64K byte, random-access memory designed for the Heathkit H8 computer. This memory is packaged on a 6 inch x 12 inch printed circuit board using standard Heath mounting hardware and plugs directly into the H8 computer. The memory occupies a single card location in the computer.

The memory bytes are 8 bits in length. The total capacity of the memory is 64K x 8 bits. Of this, only 56K can be used, because Heath reserves the first 8K of memory addresses for other uses. This memory may be addressed from 8K to 64K bytes. It will not respond when addressed in the range 0 to 8K bytes.

A single card memory is important not only to lower total memory cost but also because fewer memory cards mean more space in the main computer chassis for peripheral I/O and other types of cards. The future need for a bus expansion chassis will be eliminated in most cases.

The M-H8 may be assembled in any of the following memory sizes: 16K, 32K, 48K and 64K (56K). A small size memory can easily be expanded to a larger size (in 16K byte increments) by simply adding additional parts at some later date. Smaller varsions of the M-H8 are fully compatable with any other memory used with the H8 computer and may be used together in the computer with this memory to increase the memory capacity.

- - - - - -

* Please note that HEATHKIT and H8 are registered trademarks of the Heath Company located in Benton Harbor, Michigan.

A single card memory is made possible by using the very popular 16K x 1 bit dynamic memory integrated circuit chip. This chip became available in 1977 and is now the industry-standard random-access memory device. The 4116, as this part is called, can now be obtained at reasonable prices and should soon be finding its way into many home computer designs. It is available from a large number of manufacturers.

For large memories, dynamic memory chips have almost every advantage over static memory chips. They are less expensive, require much less space and use considerably less power than conventional static memory chips. They do require refreshing and additional control circuitry, but this is more than offset by their many advantages. In a properly designed memory, dynamic memory chips are extremely reliable. Memory chip failures are very rare once the memory chips are properly soldered onto the printed circuit board.

A BRIEF HISTORY OF MEMORY CHIPS

Semiconductor memories began to compete successfully with magnetic-core memories with the introduction of the 4K x 1 bit metal oxide semiconductor (MOS) memory chip. 4K memory chips became available around 1974 and were introduced in a wide variety of styles. There were 16, 18 and 22 pin parts. Some used static storage techniques, while others used dynamic techniques. There was little compatibility between manufacturers of similar parts.

Static memory chips store information in a flip/flop circuit, usually consisting of six transistor elements. Stored information is non-volatile as long as power is applied. Static memory chips require little additional support circuitry and are easy to use. Memory expertise is generally not required to design with static memory chips.

Dynamic memory chips store information as a charge on a capacitor. This capacitor is usually connected to the other memory circuits on the chip with a single transistor switch. A separate capacitor and switch are required for each bit of information stored on the chip. The charge stored on the capacitor gradually leaks away and must be periodically restored. This operation is called refresh. Dynamic memory chips require considerable additional support circuitry. Good printed circuit board layout practices are essential.

Dynamic memory chips have many advantages over static chips. They are less expensive per bit of stored information, store more information per chip and use less power. They provide significant savings in computer chassis space.

One 4K dynamic memory chip (type 4096) was introduced in a 16 pin package which had 6 address inputs. The 12 memory address bits required to address 4K were multiplexed into the memory chip in two steps, using these 6 inputs. This was considered a very radical technique and it was widely predicted that this chip could not compete with more conventional 4K memory chips.

Not only was this chip very successful, it provided the technique required to build the next generation memory device: The 16K x 1 bit dynamic memory chip. All 16K memory chips are alike and chips from all manufacturers may be used interchangeably. This is also a 16 pin part, with 7 multiplexed address inputs. Static 16K memory chips are not available.

At this time, 64K x 1 bit dynamic memory chips (type 4164) have been announced. These parts also use 16 pins and have 8 multiplexed address inputs. These new 64K chips use a single power supply voltage, +5 volts. The older 4K and 16K chips use three power supply voltages: -5 volts, +5 volts and +12 volts.

THE 16K DYNAMIC MEMORY CHIP

The 4116, 16K x 1 bit, memory chip package has 16 pins, as previously noted. Seven of these pins are address inputs, 3 are power supply inputs, 1 is for data input, 1 is for data output and 1 is ground. There are also 3 timing strobes known as RAS (Row Address Select), CAS (Column Address Select) and WRT (Write). All inputs and outputs (address, data and timing strobes) are TTL compatible. They may be connected to standard TTL integrated circuit inputs and outputs.

The three timing strobes are normally held at a high level and are asserted when driven to a low level. The memory cycle starts when RAS is asserted. The seven address inputs must be applied before this time. These addresses are latched into the memory chip when RAS is asserted. A short time after RAS is asserted, seven new addresses are applied to the memory chip address inputs. (A total of 14 address bits is required to address 16K of memory.) CAS is then asserted, latching these additional seven address bits into the memory chip. WRT should remain high for a memory read cycle. Valid memory data will appear at the memory chip data output after a certain time following the assertion of RAS.

This time (memory chip access time) is specified according to the grade of memory chip used. Three grades are available and their respective access times are as follows:

4116-2	150 nanoseconds
4116-3	200 nanoseconds
4116-4	250 nanoseconds

The fastest chip, the 4116-2, is the most expensive. All three grades are made using the same process and are screened into these categories by testing. Memory data will remain valid at the memory chip output terminal until the end of the memory cycle, when RAS and CAS are released.

WRT is asserted for a memory write cycle. Input data is latched into the memory chip upon the assertion of either CAS or WRT, whichever is asserted later. If the input data is changed after this time, it will not be stored in the memory chip.

The preceeding information concerning the 4116 memory chip has been provided to aid in understanding the operation of the memory. It is not sufficient for design purposes or to make design changes. Further information may be obtained from the manufacturer's data specification sheet for this part.

REFRESHING THE MEMORY CHIPS

The entire memory must be refreshed once every two milliseconds. A memory refresh cycle differs from a normal memory cycle in that only RAS is asserted. The memory must be refreshed for all Row addresses. There are seven row address inputs; therefore, the memory must be refreshed 128 times, every two milliseconds. The refresh cycles are normally evenly spaced 16 microseconds apart. This will provide 128 refresh cycles in 2 milliseconds. A refresh cycle typically has a duration of one-half microsecond. The memory will, therefore, be performing refresh cycles less than 5 percent of the time.

Memory refresh cycles potentially interfere with computer memory cycles. Refresh cycles must be synchronized with the computer to prevent interference. The operating cycles of most computers can be extended by the assertion of a "stall" signal, to accommodate memories with excessive access times. Use can be made of this provision when the memory is performing a refresh cycle. Start of a computer memory cycle will be delayed until completion of the refresh cycle and the computer stall signal meanwhile asserted until the memory can accept or provide computer data. The Heathkit H8 bus has a signal line called RDYIN, on connector pin 20, which may be used to extend the operating cycle for the computer for this purpose.

Frequently, memory refresh cycles can be timed in such a manner that the computer does not have to wait for the completion of a memory refresh cycle. In this case, the computer stall signal need not be asserted. Refresh is then said to be transparent. The stall signal may be used occasionally on start up and also when the computer is using the memory at a low repetition rate. The stall signal will not be asserted when the computer is using the memory continuously at high speed.

The M-H8 memory design features transparent refresh. There is sufficient time between computer memory cycles to follow a computer memory cycle with a memory refresh cycle. Refresh cycles will be generated in this case only when required. The memory should not be refreshed excessively, as the power requirements for refresh are significant, especially for larger memories.

MEMORY - COMPUTER BUS INTERFACEE

The Heathkit H8 bus operates the memory with a minimum of signal lines. There are 16 address lines to address up to 64K bytes of memory. There are 8 bi-directional data lines. These lines can either supply data to the memory or receive data from the memory, in 8 bit bytes. There is a memory Read input and a memory Write input. Provision for a memory output to stall the computer (RDYIN) is also available on the bus. There are three power supply voltages (four connections), as well as two ground connections. A total of 33 connections is made to the memory from the computer bus.

The H8 memory cycle has a duration 1500 nanoseconds (see the Memory Interface Timing Diagram). The memory read command pulse width is 750 nanoseconds. The memory write command pulse width is 1150 nanoseconds. A memory cycle will be either a read cycle or a write cycle, depending upon which memory command is used. A memory cycle nominally begins when a memory address is applied to the address lines. This occurs 200 nanoseconds prior to the assertion of either a read or a write command. Both the read and write commands are positive pulses on the bus.

Memory data is expected on the bus 300 nanoseconds after assertion of a memory read command. Similarly, data to be written into the memory is available on the bus 300 nanoseconds after assertion of a memory write command. Address and data information remain valid on the bus until release of the memory read or write commands.

The memory access time requirement, measured from the assertion of the memory addresses on the bus, is 500 nanoseconds. The 8K static memory used by Heath (H8-1) has a specified access time of 450 nanoseconds. However, the static memory cycle starts when the memory addreses are asserted. The dynamic memory cannot start at this time and must wait until a memory read or write command is applied. Accordingly, 200 nanoseconds are lost from the memory access time. This is of some concern and is the primary reason for using the faster "dash 3" (4116-3) memory chips. Using these chips, the memory module access time can be set for less than 500 nanoseconds.

Memory read data is applied to the bus through a latching bus driver. This driver has a high impedance output unless it is applying data to the bus. Read data remains latched on the bus until the release of the memory read command. The memory, itself, is free for refreshing immediately after data is latched on the bus.

Data to be written into the memory may not be available until 300 nanoseconds after assertion of the memory write command. Data must be latched into the memory chips, using the WRT command, after this time.

As previously stated, the memory cycles are started by the assertion of the memory read or write commands. The duration of these commands has no affect on memory operation. The memory cycle finishes as soon as possible, followed by a short pause. The memory is then ready to begin a refresh cycle. There is time for a refresh cycle and another pause before the next memory cycle will be requested. Synchronized in this way, memory refresh cycles do not delay execution of the computer memory cycles.

It may be desirable to examine the memory timing or computer bus signals with an oscilloscope. If the memory is operating satisfactorily, a good way to examine these signals is to do so while running the computer using a short program loop entered in machine code.

Short (minimum program) test loops are very useful, as known operating conditions can be established and related to the oscilloscope signal patterns. Different test loops can be used to provide different bus conditions.

Examples:

(1) Store accumulator in memory (write into memory) and return;
(2) Load accumulator from memory (read from memory) and return.

The front panel monitor routine can be disabled when running these loops, although this is not always necessary. The bus signals can be spread out in intelligible patterns by synchronizing the oscilliscope with an appropriate high order address bit. This synchronizing address change can be written into the beginning of the test routine.

MEMORY MODULE CIRCUIT DESCRIPTION

The memory continuously generates refresh cycles whether the computer is using the memory or not. Much of the memory circuitry can be checked and the timing adjusted simply by working with the refresh signals.

Refresh cycles are timed by oscillator U39 (see Schematic, sheet 2) whose period is approximately 16 microseconds. A 100 nanosecond pulse is generated every 16 microseconds by applying the output of this timer to One Shot U38. A 5 microsecond refresh-enable timing slot is generated when the 100 nanosecond pulse is applied to the other half of U38.

If a memory cycle occurs during the 5 microsecond refresh enable interval, a refresh cycle will be initiated immediately following the memory cycle. In this case, the output at U58-5 will start the refresh cycle. If there is no memory cycle, the termination of the 5 microsecond timing interval will initiate the refresh cycle. In this case, the output at U58-9 will start the refresh cycle.

In either case, the refresh cycle will start when the output of U60-10 is asserted. A delay set by U57-4 will occur before the application of refresh RAS to the memory chips. A delay is required between successive applications of RAS to the memory chips, between the memory and refresh cycles. During this delay, the address multiplexer applies the outputs of the refresh counter to the memory chip address inputs. The refresh counter is advanced with each refresh cycle. When refresh RAS is removed from the memory chips, there is another delay, after which the refresh cycle is terminated. The memory is now ready to begin a computer memory cycle.

A computer memory cycle is initiated when the input at U44-3 goes high. This will cause the output at U44-5 to go high. If a refresh cycle is in progress, the input at U52-5 will be low and a memory cycle will not begin until U52-5 goes high, at the end of the refresh cycle. A memory cycle begins when U54-6 goes low. The duration of the memory cycle is determined by the output of One Shot U45-4.

The WRT input to the memory chips is used for both read and write memory cycles. In the case of a read cycle, the memory chip data output is simply written back into the input. The data content is not changed.

All timing information given on the schematic is for reference purposes only, to aid in following the operation of the schematic circuits. The memory timing should be checked using the Memory Internal Timing Specification sheet. The value of an occasional One-Shot timing resistor may have to be changed to obtain this timing. Timing not shown on the specification sheet is not critical and should be obtained using the component values given on the schematic.

A jumper wire is used to select the memory address range (see Schematic, sheet 2). Common connection A is connected to one of the jumper pads B, C, D or E, depending on the memory capacity installed on the board. The 9 memory signal outputs (8 data outputs and RDYIN) are inhibited when the memory address is outside the range selected by the jumper wire.

Seven address inputs are simultaneously applied to each row of memory chips. These addresses come from the address multiplexing circuits (see Schematic, sheet 3). Fourteen address inputs (A0 through A13) from the computer are separately applied to the memory chips in two steps, using the address multiplexing circuits, to select a memory chip address. In addition, seven refresh addresses are supplied to the memory chips during a refresh cycle. The refresh addresses are generated by the refresh counter, which advances by one count with each refresh cycle. The refresh address counter repeats after every 128 refresh cycles.

The memory address decoder (U37) decodes the memory addresses in 8K blocks throughout the address range (see Schematic, sheet 1). 64K is decoded using 8 decoder outputs. Address bits are asserted on the bus when they are brought low. All address bits will be high for the lowest address count (000 000). All address bits will be low on the bus for the highest address count (377 377). 8K decoding blocks are necessary to disable the memory when addresses are selected in the lowest 8K range. The decoder outputs are then paired using U36 to provide a 1 out of 4 decoded address (four 16K decoding blocks) to select one of the four rows of 16K memory chips. The memory RAS timing signal, MCS-H, is applied to only one row of memory chips at a time. This is the address-selected memory chip row. Conversely, the refresh RAS timing signal, RRS-H, is applied to all rows of memory chips, simultaneously.

CAS and WRT are also applied to all memory chips simultaneously, during computer memory cycles. Only RAS is applied to the memory chips during a memory refresh cycle.

The memory chips are laid out in an array of four rows, with eight chips in each row (see Schematic, sheet 4). Respective address and control inputs are connected together in each row. Each row has separate address and control drivers in order to reduce the load on each driver. The d-c input resistance to the memory chips is very high, whether they are driven high or low. However, the inputs have significant capacitance and this requires good drivers and series damping resistors. The MOS memory chip inputs are TTL compatible, but they are not the same as TTL inputs.

All of the respective memory chip power supply voltages are connected together. This is not shown on the schematic, but is understood. The memory chip data inputs and outputs are connected together by column, rather than by row. There are eight columns of data input and output lines, with four memory chips in each column. The data lines are thus connected to the memory chips in opposition to the address and control lines. In this manner, the memory chip array is formed.

POWER SUPPLY CIRCUITS

The three power supply voltages required by the memory are obtained from the H8 computer. The unregulated H8 voltages, -18, +8 and +18, are converted into -5, +5 and +12 volts, respectively, using standard 3-terminal voltage regulator chips (see Schematic, sheet 3). The power supply current requirements are small: -5 volts at 10 milliamperes, +5 volts at 1.0 amperes and +12 volts at 200 milliamperes. The +5 volts is supplied using four regulators driving separate loads, none of which draw more than 300 milliamperes. The integrated circuits are split into four groups, each connected to a separate +5 volt power supply regulator. This is done to permit using the small Heath mounting bail as a heat sink.

The average (d-c) currents supplied by the +12 volt and -5 volt regulators are quite small. The peak currents these voltages must supply on the memory board however, are large and are many times greater than the d-c currents. These currents are supplied by the filter capacitors in the memory chip array. A large amount of distributed capacitance is required for this purpose. The memory chip array has been very carefully laid out. These currents must be supplied when needed and without inducing spurious signals. During refresh, these currents are drawn by all of the memory chips on the board, simultaneously.

The printed circuit board layout requirements for dynamic memory chips are rather exacting. The memory chip layout on the M-H8 should be appreciated. The memory chips are arranged together in compact group. All three power supply voltages, as well as ground, are respectively cross connected at each memory chip. All three power supply voltages and ground each form a net with a memory chip connected at each interconnection. A pair of filter capacitors is connected to each memory chip.

High quality ceramic capacitors are essential to filter the memory chip power supply voltages. Capacitor quality is directly related to price. Inexpensive capacitors will not work for this purpose. A capacitance of 0.1 microfarads should be used for each capacitor. The CK05BX104K capacitor is suitable for this purpose and was selected for its small size.

The memory +12 volt (H8, +18 volt) power requirement can be reduced by refreshing the memory chips at a lower rate. This should not cause a problem at room temperatures. The refresh rate requirement is very strongly influenced by temperature. The 2 ms refresh rate requirement is specified by the chip manufacturer for operation at 70 degrees centigrade. Increasing the value of R11 or C74 will decrease the refresh rate. The refresh rate can be reduced by several hundred percent, if desired.

ASSEMBLING THE MEMORY MODULE

The memory module uses standard parts as much as possible. The integrated circuits, especially the memory chips may be installed in sockets to facilitate troubleshooting and part replacement. It is important to use high-quality sockets. After a period of time, sockets can be a real source of trouble. Intermittant problems can be very hard to isolate in digital computer systems. Whether or not to use IC sockets is, therefore, an important decision.

Discrete 1/8 watt resistors may used in place of the resistor modules. Eight resistors will be needed in place of each of the 5 resistor modules. The value of this resistor is not critical and may range from 20 to 60 ohms, with 33 ohms being optimum. The memory will probably work OK if bare wire is used in place of a resistor, but this is not recommended.

The memory should also work OK with less than 64 filter capacitors in the memory array. The filter capacitor count can be reduced by one-half for each of the three power supply voltages (see Schematic, sheet 3). The capacitors used in this case should be physically dispersed for each of three voltages. It is still important to use high-quality filter capacitors.

To install less than 64K (56K) of memory capacity, simply omit one or more rows of memory chips and their respective capacitors (and sockets, if used). Alternately, the capacitors and sockets can be installed and the memory expanded simply by adding rows of memory chips at some later date. The memory can be expanded in 16K increments from 16K to 64K, at any time. Each row of eight memory chips constitutes a 16K byte. The memory must be expanded in continuous rows from the top of the board (see Assembly Drawing).

The four +5 volt regulators should be mounted to a standard H8 PC board mounting bail (see Voltage Regulator Installation sheet). The use of sockets for these voltage regulators is discouraged. Wires should be soldered directly to the voltage regulator leads from the printed circuit board.

The bail acts as a heat sink for the voltage regulators. It is a rather poor heat sink and four +5 volt regulators were used to reduce the regulator junction temperatures. The bail gets hot and may reach 150 degrees Fahrenheit. This is normal and should cause no concern. Silicon grease should be used when mounting the regulator chips.

Early versions of the H8 computer used tin-plated bus connectors on both the plug-in modules and the mother board. After a period of time, these connectors can become unreliable to varying degrees. These connectors are manufactured by Molex and are available with either tin or gold plating. The gold-plated connectors are naturally more expensive. Gold does not corrode. A corrosive film usually forms on the tin plating, making tin connectors less reliable. Heath now uses gold connectors for the H8 modules.

Gold connectors should be used on the M-H8 memory module. If the H8 mother board bus connectors are tin-plated they should be replaced with gold-plated connectors. This will make an enormous difference in the reliability of the computer. The best way to do this is to obtain a blank mother board from Heath and install all gold connnectors on it. The computer must then be taken apart to replace the mother board.

Polarity must be observed, of course, when installing the semiconductor components (integrated circuits, diodes and the transistor) and the tantalum capacitors. Any silicon diode may be used for the 4 diodes used on the board. Any silicon switching transistor may be used in place of the 2N4400.

Substitute integrated circuits should not be used. The 74S373, for example, uses considerably more power than the 74LS373. 74S-- series integrated circuits provide more high-level drive and are used to drive the memory chips. In addition, the memory timing components were selected using the integrated circuits specified.

The connector tie bar used on Heath H8 boards serves no function and may be omitted. It is not available from Molex.

Be sure to install a jumper on the printed circuit board between U47, pin 11 and U54, pin 20 to the pad below U46, pin 11 (see Assembly Drawing). This supplies +5 volts to U54.

The memory address range jumper must also be installed (see Memory Address Selection sheet). The memory will not work in the H8 computer unless this jumper is installed.

This is a tightly laid out board and should be soldered with care. Do not use too much solder. Small diameter (.032 in.) solder should be used. Resistor leads are usually dirty and require extra flux and heating time. Make sure all of the integrated circuit (or socket) leads are soldered. It is very easy to miss one of these leads.

A note of caution: The memory chips are MOS integrated circuits and are susceptible to damage by static electricity. They must be handled carefully, observing the precautions normally taken when using MOS devices. The memory chips must be handled one at a time and should be stored on a conducting foam pad, or in their anti-static shipping tubes, when removed from the memory module. If sockets are not used, the memory chips should be soldered on the board one chip at a time, after the TTL integrated circuits are installed. Once the entire board has been soldered, there is no danger to the memory chips.

OPERATION

The M-H8 always starts at the 8K address boundary. The upper address boundary is jumpered to correspond with the memory capacity installed on the M-H8 module. Additional memories (Heath H8 static memories, etc.) may be jumpered to run above the M-H8. It should be noted that the H8 computer can only address memory from 8K to 64K bytes.

After power is turned on, the H8 front panel monitor will continuously address memory in one bit increments starting at the 8K address boundary, until it fails to obtain a reponse. The monitor records this as the upper limit of memory by setting the stack pointer at this location. HDOS opperates in much the same way, printing the located memory capacity on the CRT screen.

It is expected that a lower capacity version of the M-H8 (16K or 32K) will be initially used in most systems together with the existing static memory boards, to extend the memory capacity. Ultimately, however, the M-H8 will most likely be expanded to its full capacity and the static memory boards decarded.

When the memory is operating properly, the front panel (monitor program) can be used. This the first indication that the memory is good. It should be possible to load and store information anywhere in the memory address range using the front panel key pad.

The H8 memory test routine should now be entered and run to insure that the entire memory is functional and that there are no defective memory chips. The Heath H8 memory test routine is listed on pages 61 and 62 of the H8 operating manual. It is also listed and described in detail on pages 9 through 14 of the same operating manual. To test 56K of memory, the data placed in address location 040 105 (split octal notation) of the test routine should be changed from 057 to 377. This raises the upper memory test limit from 12K to 64K. (The memory test always begins at 8K.) To test smaller amounts of memory, the upper memory limit can be varied accordingly. See page 0-58 of the H8 software reference manual for the high byte address boundaries for 4K decimal increments of memory capacity. The high byte addresses are octal numbers and should be decremented by one (except in the case of 377) to set the memory boundary. Example: 200-1=177 sets the upper memory boundary at 32K, for a memory capacity of 24K.

The memory test begins at address 040 160, rather than at the 8K boundary of 040 000. Also, the memory test ends at xxx 260, rather than the upper limit of memory at xxx 377. These small amounts are excluded from the test because they are needed to operate the front panel monitor program, as well as administer the test, itself. If the memory test runs successfully, it can be assumed that these excluded locations are also good. However, these locations can be thoroughly tested by interchanging RAS lines on the memory module. This is easy to do: RAS 0 and RAS 3 are interchanged with RAS 1 and RAS 2, respectivly. This physically changes the memory chip rows responding to given address inputs, in 16K blocks. The RAS line interchange can be effected by simply raising one end of each of the 33 ohm drive resistors (R6, R15, R22, and R27) and cross connecting these ends to the PC board with insulated wire. This is also a good troubleshooting technique to isolate or by-pass a row with a defective memory chip.

It should be noted that the memory test routine is primarily used to locate defective memory chips and certain kinds of problems in the memory chip array wiring (solder bridges, for example). If there are problems in the memory chip timing and control circuits, the front panel monitor program will not run.

Three flip/flop integrated circuit chips (7474) are used on the M-H8 memory module. These must be initialized when the power is turned on. This is done by the initialization circuit using transistor Q1. Board initialization is performed automatically every time the power is turned on. If the board fails to initialize or locks up for any reason during operation, the H8 power switch (located on the rear panel of the computer) should be turned off and then turned back on again after a delay of two or three seconds. The reset function on the H8 front panel is not connected to the M-H8 meemory and cannot be used to reset the memory.

TROUBLESHOOTING THE MEMORY MODULE

In some cases, what appear to be memory problems may actually be problems in the computer hardware or software. It is now well established that the H8 is somewhat unreliable. The sockets on the voltage regulator leads on every H8 module should be removed and these leads soldered directly to the voltage regulator terminals. Some of the other component sockets may also present problems. Socket and connector problems occur unpredictably, at infrequent intervals. Individual sockets can be tested by physical manipulation of the installed parts while running a program. Programs should run without interruption while sockets, or entire printed circuit boards, are flexed or tapped.

Tin-plated bus connectors will cause trouble sooner or later. Both motherboard and PC module mating connectors should be replaced with Molex gold-plated connectors in all cases. The mounting of the H8 bus mother board is not very rigid. Installing or removing one board may affect any other boards on the bus. If any problems occur, all boards should be loosened and reseated on the bus connectors. This is very important. The first step in troubleshooting is to eliminate the bus connectors (and any other connectors or sockets) as the source of trouble.

If an infrequently occuring (intermittent) problem exists there may be a poor connector somewhere. A vibrating tool can be applied all over a suspected printed circuit board or other parts of the system to locate poor connections. (An engraving tool set at the lowest level with a plastic ball attached to its tip is ideal for this purpose.) A program loop should be running while the boards and connectors are vibrated. The program must not stop during this test.

Problems can also exist in the computer software. When running Heath Cassette BASIC, the upper memory limit may have to be set below 40K. The software configuration option can be used to do this.

Be very careful not to short bus connector pin 2 (-18 volt power input) and pin 3 (signal line) together when making measurements. This will destroy a number of integrated circuits on several computer modules. These pins are located at the bottom of the chassis and are very easy to short together.

If the front panel monitor program won't run, there is a problem in the memory timing and control circuits. Check all 6 power supply voltages first (+5 volts from each of the 4 regulators, +12 volts and -5 volts). An oscilloscope suitable for digital troubleshooting should then be used to check circuit operation. An oscilloscope with an external trigger input (preferably dual trace) is required to make differential timing measurements. Frequently, it will be possible to provide signals to the memory interface to check memory circuit operation when the front panel does not operate properly. Repeated pushing of front panel buttons (especially reset) should set this up. All timing adjustments should be checked first wherever possible (see Internal Timing Specification). Circuit operation can then be checked for defective or improperly installed parts. It may be helpful to interchange RAS lines between rows of memory chips, as noted earlier, to locate or bypass rows with bad chips.

If a memory read access time problem is suspected, RDYIN can be asserted for every memory cycle to check this. To accomplish this, U51-1 can be connected to a +5 volt bias resistor or to U51-2. If the memory can be made to work with this modification, but not otherwise, the memory access time may be excessive.

Faster memory chips can be used to reduce the access time. The 4116-2 is 50 nanoseconds faster than the 4116-3. Reducing the respective delays for address multiplexing and the assertion of CAS is necessary to take advantage of the faster chips. This should be done if the 4116-2 memory chip is used. Each of these two delays should be reduced by one-half when using the 4116-2 memory chip.

A memory access time problem can be solved, of course, by simply asserting RDYIN for every memory cycle. Using RDYIN with each memory cycle, however, is undesirable as this will cause the computer to run more slowly. No problem with access time has been experienced using the H8 computer with either the 4116-3 or the 4116-4 memory chips. The 4116-3 memory chip is used to provide a proven margin of reliability.

The refresh cycle must terminate before a computer memory cycle can begin. Access time will be adversely affected if the refresh cycle delays the beginning of a computer memory cycle when refreshing in the transparent mode. The input at U52-5 must go high before the input at U52-4 goes high to begin a computer memory cycle, in this case. This is insured by the 1400 nanosecond maximum specification shown in Memory Internal Timing Specification.

The Heath H8 memory test can be used to test the memory chips once the front panel monitor program is running. This test will locate more than 90 percent of all memory chip problems. No memory test can locate all memory problems. Dynamic memories are susceptable to soft (non-repeatable) errors caused by such things as alpha particle radiation from the memory chip case materials. Fortunately, these errors are very infrequent.

Special test patterns have been devised for semiconductor memories (so called marching, walking and galloping patterns) which are used extensively in commercial practice. They are easy to program, but take extensive periods of time to run. They are all considerably more sophisticated than the Heath H8 memory test. A comprehensive semiconductor memory-test program listing written for the H8 computer will be available soon. This listing may be obtained upon request.