
```
CASE$TASK: PROCEDURE REENTRANT;
  DO CASE RUNNING$TASK;
    CALL TASK00;
    CALL TASK01;
    CALL TASK02;
    CALL TASK03;
    CALL TASK04;
    CALL TASK05;
    CALL TASK06;
    CALL TASK07;
    CALL TASK08;
    CALL TASK09;
  END;
TASK(RUNNING$TASK).STATUS=TASK(RUNNING$TASK).STATUS AND
  NOT (BUSY$BIT OR READY$BIT);
TASK(RUNNING$TASK).PNTR=0;
IF RUNNING$TASK=NEW$TASK THEN DO;
  IF NEW$DELAY<>0 THEN DO;
    CALL ACTIVATE$DELAY;
  END;
  ELSE /* SINCE NEW$DELAY=0 */ DO;
    CALL ACTIVATE$TASK;
  END;
END;
RUNNING$TASK=0;
RETURN;
END CASE$TASK;

PREEMPT:PROCEDURE REENTRANT; /* ASSUMES INTERRUPTS DISABLED */
IF PREEMPTED$TASK=0 THEN DO;
  IF (HIGH$PRIORITY$HEAD<>0) AND (RUNNING$TASK>=
    FIRST$LOW$PRIORITY$TASK) THEN DO;
    PREEMPTED$TASK=RUNNING$TASK;
    RUNNING$TASK=0;
    DO WHILE PREEMPTED$TASK<>0;
      CALL DISPATCH;
    END;
  END;
END;
RETURN;
END PREEMPT;
```

```
DISPATCH:PROCEDURE REENTRANT; /* ASSUMES RUNNING$TASK=0 */
DISABLE;
IF HIGH$PRIORITY$HEAD<>0 THEN DO;
  RUNNING$TASK=HIGH$PRIORITY$HEAD;
  HIGH$PRIORITY$HEAD=TASK(RUNNING$TASK).PNTR;
  IF HIGH$PRIORITY$HEAD = 0 THEN HIGH$PRIORITY$TAIL = 0;
  CALL CASE$TASK;
END;
ELSE IF PREEMPTED$TASK<>0 THEN DO;
  RUNNING$TASK=PREEMPTED$TASK;
  PREEMPTED$TASK=0;
END;
ELSE IF LOW$PRIORITY$HEAD<>0 THEN DO;
  RUNNING$TASK=LOW$PRIORITY$HEAD;
  LOW$PRIORITY$HEAD=TASK(RUNNING$TASK).PNTR;
  IF LOW$PRIORITY$HEAD = 0 THEN LOW$PRIORITY$TAIL = 0;
  CALL CASE$TASK;
END;
ELSE RETURN;
RETURN;
END DISPATCH;
```

```
TASK00: PROCEDURE REENTRANT; /*ERROR CODE*/RETURN;END TASK00;

TASK01: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK01 CODE*/
  DISABLE;
  RETURN;
  END TASK01;

TASK02: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK02 CODE*/
  DISABLE;
  RETURN;
  END TASK02;

TASK03: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK03 CODE*/
  DISABLE;
  RETURN;
  END TASK03;

TASK04: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK04 CODE*/
  DISABLE;
  RETURN;
  END TASK04;

TASK05: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK05 CODE*/
  DISABLE;
  RETURN;
  END TASK05;

TASK06: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK06 CODE*/
  DISABLE;
  RETURN;
  END TASK06;

TASK07: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK07 CODE*/
  DISABLE;
  RETURN;
  END TASK07;
```

```
TASK08: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK08 CODE*/
  DISABLE;
  RETURN;
  END TASK08;

TASK09: PROCEDURE REENTRANT;
  ENABLE;
      /*TASK09 CODE*/
  DISABLE;
  RETURN;
  END TASK09;

      /*INITIALIZE*/

DISABLE;
DO STATUS=0 TO 9;
  TASK(STATUS).PNTR=0;
  TASK(STATUS).STATUS=0;
  TASK(STATUS).DELAY=0;
  NEW$TASK,NEW$DELAY=0;
  HIGH$PRIORITY$HEAD,HIGH$PRIORITY$TAIL=0;
  LOW$PRIORITY$HEAD,LOW$PRIORITY$TAIL=0;
  RUNNING$TASK,PREEMPTED$TASK=0;
  END;

      /* MAIN LOOP */

DO WHILE TRUE<>FALSE;
  CALL DISPATCH;
  ENABLE;
  STATUS=STATUS;
  END;

END TM86;
```

REFERENCES

1. Hansen, Brinch, *Operating System Principles*, Prentice-Hall, Englewood, N.J., 1973.
2. Knuth, D. E., *The Art of Computer Programming*, Addison-Wesley, Reading, Mass., 1969.
3. Wirth, Nicklaus, *Algorithms + Data Structures = Programs*, Prentice-Hall, Englewood, N.J., 1976.
4. "PL/M-86 Programming Manual," Intel Corporation, 1978, manual order number 9800466A.
5. "RMX/80 User's Guide," Intel Corporation, 1977, manual order number 9800522B.



September 1979

**Debugging Strategies and
Considerations for 8089
Systems**

Debugging Strategies and Considerations for 8089 Systems

Contents

INTRODUCTION

STATIC (OR FUNCTIONAL) DEBUGGING

- Hardware Testing
- External Processor Interface
- Software Testing

REAL-TIME TESTING

- Logic Analyzer Techniques

A REVIEW OF IOP OPERATION

- Task Execution
- Going from Instruction Execution into DMA
- DMA Termination
- Priorities/Dual Channel Operation

SUMMARY

Appendix I. CHECKLIST OF POSSIBLE PROBLEMS

Appendix II. BREAKPOINT ROUTINE AND CONTROL PROGRAM

Our thanks to John Atwood and Dave Ferguson, the authors of this note. Both John and Dave are members of Intel's 8089 design engineering group. Please direct any questions you may have to your local Intel FAE (field application engineer) or to MPO Marketing at Intel, Santa Clara.

INTRODUCTION

The Intel 8089 is the first integrated I/O processor available. This I/O processor (IOP) makes available the power of I/O channels, as used in mainframes and mini-computers, in a microcomputer form. Designed as part of the MCS-86™ family, the IOP can be interfaced with the MCS-80™ and MCS-85™ families as well.

An I/O channel is basically a processor remote from the main CPU, which independently runs I/O operations upon command of the CPU. To relate the 8089 to existing LSI components, it is similar to a microprocessor that is time-multiplexed with a DMA controller, but with two channels available. However, since the 8089 processor is optimized for I/O and multiprocessor operations, and the DMA has been made much more flexible than existing DMA controllers, a truly general purpose and powerful I/O control system is available on one chip.

Due to the uniqueness of the 8089, this application note was written to review debugging strategies and point out possible pitfalls when developing an IOP system. Debugging an IOP system is very similar to debugging microprocessor/DMA controller systems, and many of the techniques described here are standard microprocessor techniques. However, several factors are present which can complicate the debugging process:

1. Multiprocessor Operation

Although usable by itself, the IOP is designed to be used with other processors. All factors normally encountered with multiprocessor operation, including bus arbitration, processor communication, critical code sections, etc., must be addressed in the design and debug of an IOP system.

2. DMA Tie-in to IOP Program Execution

The relationship between IOP program execution and DMA transfers and termination is different from earlier DMA controllers and should be fully understood to properly run the system.

3. Dependency of Programs on Real-Time I/O Operations

Requirements by I/O devices for maximum data rates and minimum latency times force the software programmer to be aware of hardware timing constraints and can complicate program debugging.

4. Dual Channel Operation

Related to multiprocessor operation and real-time dependencies, the two independent channels available on the 8089 may have to be coordinated with each other to make the whole system function. Dependence of one channel on the other can also complicate debugging.

Due to the complexities of running in a real-time environment, as many steps as possible should be taken to facilitate debugging. A major help here is to make sure as much of the hardware and software as possible is working before running real-time tasks. This is a good practice anyway, but it should be reemphasized that a complex multichannel system can quickly get out of hand if more than a few things are not right.

An aid to debugging any system is a clean, well organized system design. The 8089 lends itself to structured, modular software interfaces to the host CPU, via the linked-list initialization structure, and parameter communication through the parameter block (PB) area. Some of the aspects of structured programming that aid debugging are:

- *Top Down Programming* — The functions done by low-level routines are well understood, and the number of program fixes, which can cause more errors, is minimized.
- *Program Modularity* — Small, easy to manage sub-programs can be debugged independently, increasing the chance that the entire system will work the first time.
- *Modular Remoteness* — By having all program modules communicate only through a well-defined interface, one module's knowledge of the "inner workings" of another is minimized. System software complexity is reduced. Updates to program modules are more reliable, too.

Two major areas of debugging will be outlined here — static (or functional) debugging in which the hardware and software are not tested in a real-time environment, and real-time debugging. Applying a logic analyzer to IOP debugging will also be explained, and a review of IOP operation and potential problems will be done.

STATIC (OR FUNCTIONAL) DEBUGGING

The predominant errors in a system, when first tried out, are either errors in implementation (i.e., wrong hookups or coding errors), or an incorrect implementation (a wrong assumption somewhere). Most of these bugs can be found through static debugging techniques that are usually easier to work with than real-time testing.

Hardware Testing

Static hardware testing is done mainly to see if all individual parts of the system work, so the whole system will "play" when run. The level of testing can run from checking for continuity and shorts (which finds only hookup errors) to trying to move data around and running I/O devices from a monitor or special test programs (which can also find incorrect circuit design). In all but the simplest systems, the latter approach is recommended since it is a step towards software debugging.

Several approaches to hardware testing will be covered. Running diagnostic programs (such as a monitor) out of the IOP's host system, in both the LOCAL and REMOTE modes, will be covered. The case where the host system cannot support diagnostic software and must have an external processor to exercise the IOP and its peripherals will also be explained.

The case where the host system can run diagnostics or test programs that have interactive user I/O, such as a CRT terminal or teletype, provides the most straightforward way to test the IOP. Naturally, before these programs can be run, the basic hardware must be correct enough to run programs. When this point is reached, a monitor program can be used to exercise memory and I/O controllers on the system bus.

It should be mentioned that aids, other than just testing with software, are helpful for hardware debugging. While a necessity for real-time debugging, a logic analyzer is also a definite help for static hardware debugging. Its main use in hardware debugging is showing timing relationships between address or data paths and other signals. It is especially useful for functional software debugging, to be described shortly. The last debugging section outlines the use of an analyzer with the IOP. Of course, an oscilloscope, logic probes and pulsers, etc., can be used to trace out specific logic or timing problems.

LOCAL Mode

When the IOP is running in the LOCAL Mode, all I/O controllers and memory are accessible by the host or controlling CPU. Thus a standard monitor, such as the one supplied with the SDK-86 or available for the ISBC-86/12™ development kit, can exercise all hardware on the bus.* The breakpoint routines, however, will not work due to the different instruction set. The 8086 or 8088 is best suited for running the IOP in the LOCAL mode due to identical status lines and bus timing, as well as the Request/Grant line, which eliminates bus arbitration hardware. Figure 1 shows the general LOCAL mode configuration.

*The SDK-86 serial monitor is a good basis for a general 8086 monitor. The IOP cannot be used directly with the SDK-86, since the 8086 is running in the minimum mode. The SDK-86 can be converted to run in the maximum mode, if desired.

REMOTE Mode

From a system design standpoint, running the IOP in the REMOTE Mode is advantageous in that it removes the I/O bus cycles from the system bus. Normally, the remote I/O is not accessible to the host CPU. Until the IOP is able to run its own test programs to transfer data from the REMOTE bus to the system bus, I/O controllers and memory on the REMOTE bus will be invisible to the host. To get around this problem during prototyping, either an external processor interface can be used (see next section), or a temporary bypass can be made to access the REMOTE bus from the system bus.

Bypassing the normal REMOTE/SYSTEM interface is a handy technique for doing preliminary debugging on the REMOTE bus. This can be done by memory-mapping the IOP's I/O space into an unused portion of the host CPU's system memory space. When accessing this space, the IOP access to its own I/O space is disabled, and a separate set of address buffers, transceivers and bus control signal buffers are enabled. Reads and writes can then be done to the formerly inaccessible REMOTE bus by the host CPU.

A simple system (Figure 2) implements this bypassing scheme. It was designed for just forcing or examining devices on the REMOTE bus and may not read or write correctly if the IOP is simultaneously trying to do bus cycles. A more sophisticated arbitration system would permit reliable run-time checking also.

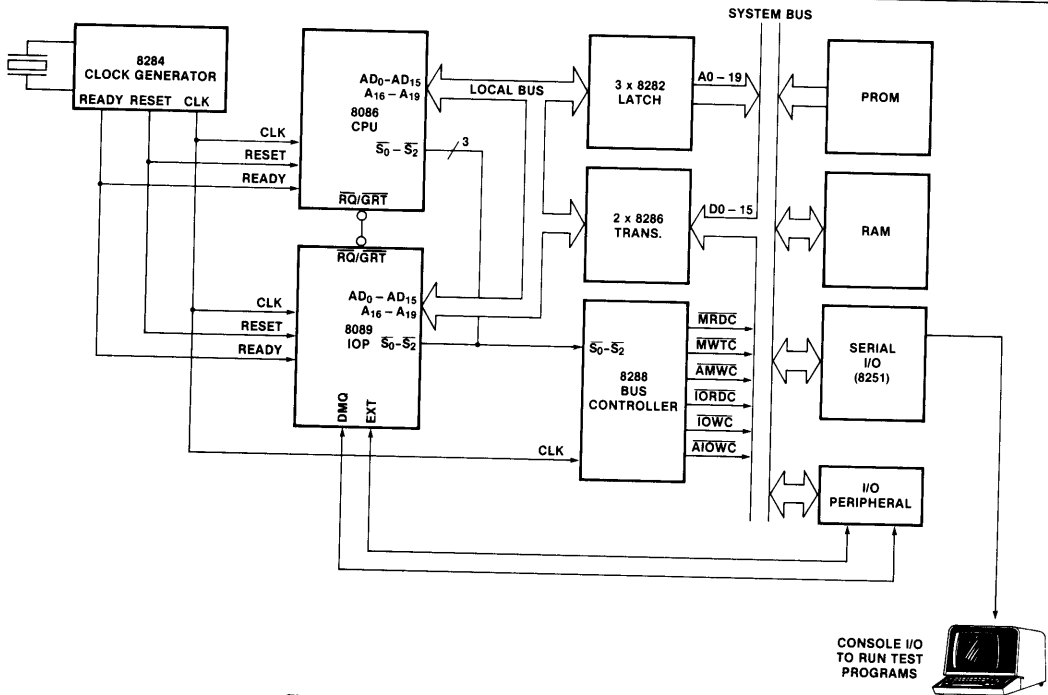


Figure 1. Generalized LOCAL Configuration—8086 in Max Mode

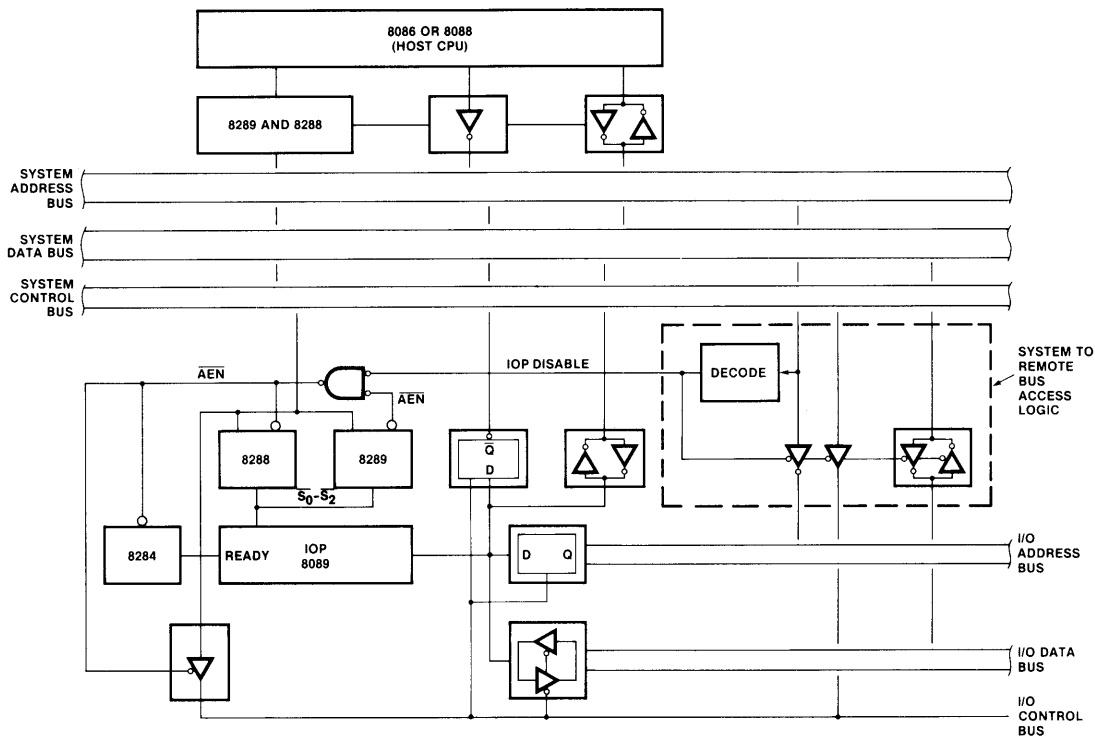


Figure 2. Remote Mode Bypass for Debugging

Running the IOP in the REMOTE mode, particularly if the MULTIBUS™ protocol is adhered to, has the advantage that the IOP can be exercised with any MULTIBUS-compatible processor. If the main processor is not amenable to being used as a debugging tool, another processor could be used to debug the hardware interface. If the microprocessor is of the same type as the intended host processor, software debugging can be done as well. A generalized REMOTE mode configuration using the MULTIBUS is shown in Figure 3.

External Processor Interface

A technique that can be used if the host processor cannot run any debugging or monitor routines is to have an external processor tie into the host processor's bus. This is useful if the main system CPU cannot run an interactive monitor or other debugging programs. If a MULTIBUS interface is being used, an 8289 bus arbiter and a set of address/data/control buffers can be used. A somewhat simpler system, similar to the remote bus access system mentioned above, could be used for static debugging of non-MULTIBUS systems. Again, if true bus arbitration is added (which brings us nearly to a MULTIBUS interface), it could also be used for run-time testing. Intel processors that have the MULTIBUS interface include the iSBC-80/20™, iSBC-86/12™, iSBC-

80/10™, iSBC-80/05™, the Intellec® development systems, among others.

In the previously described systems, the external processor would disable the host CPU's access to the bus, either by some form of bus request or by a "brute force" disabling of the CPU's buffers. In the latter case, the external processor could only control the bus during a time that the CPU is halted, without destroying the program flow. Mapping the processor's memory space into the external processor memory space is the simplest method, but can impact programs being run on the external processor. If the processor under test utilizes the MULTIBUS interface (with bus arbitration), then a processor like the iSBC-80/30™ or iSBC-86/12™ could be used as the debug vehicle with no special hardware. A more flexible interface that would have less impact on the system memory space would have the addresses for the system under test generated from latches loaded by the I/O instructions from the external processor. This case must have software routines to interface to the I/O ports and handle the desired debugging routines (see Figure 4).

Software Testing

It is desirable to check as much of the IOP program as possible statically, since various tools and techniques are available which may not be usable during real-time

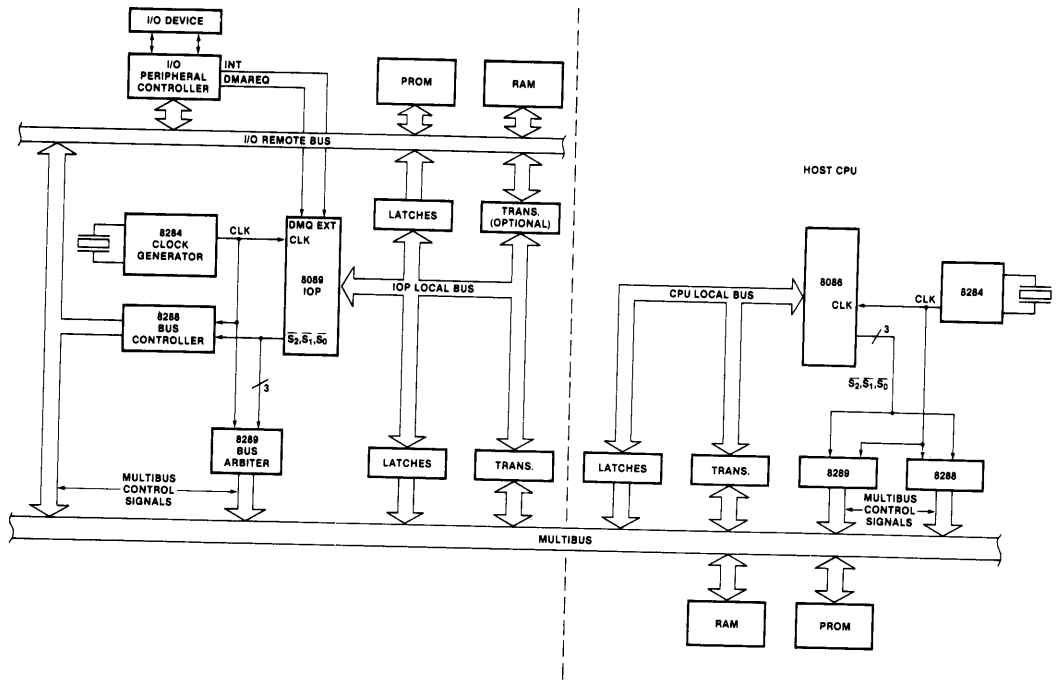


Figure 3. Generalized Remote Bus Using MULTIBUS Interface

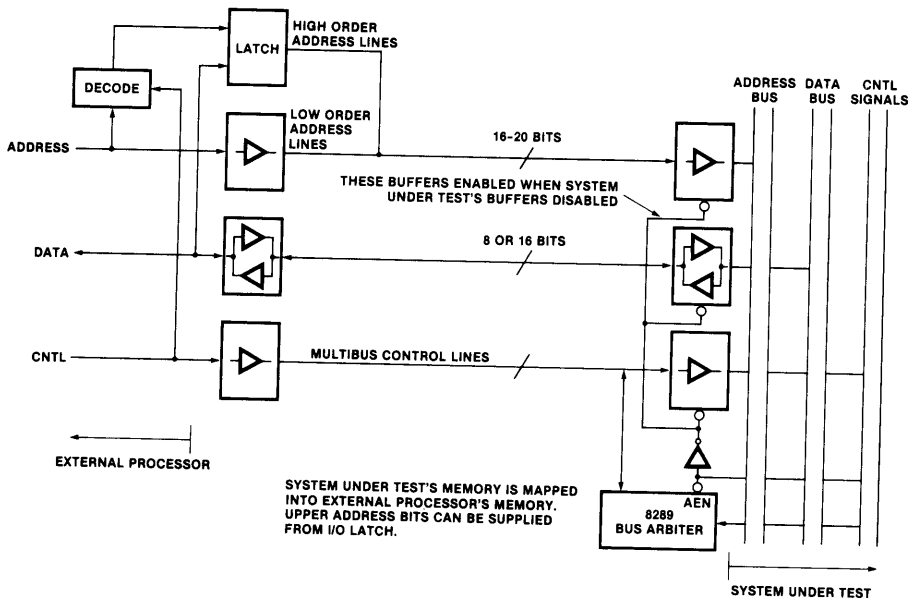


Figure 4. External Processor Interface

testing. This "static" software testing is not applicable to heavily I/O-dependent or DMA-dependent routines, but is best suited to longer computational or data handling routines. The idea is to test the correctness of algorithms, rather than seeing if the whole system runs.

There are two main approaches to functional software testing. One is to essentially run the program in real time and monitor program flow on a logic analyzer. The difference between this and real-time testing is that program subsections can be tested separately by using different TP (Task Pointer) starting addresses. If it is necessary to set up certain registers or parameters in memory, a small "setup" program can be run after initialization, which can load up registers or memory, then jump to the program section desired.

Another technique is to run the programs with breakpoint routines so that one can step through code segments and follow program execution. Software breakpoints are usually implemented by inserting a jump or restart to a monitor routine at the breakpoint location. This jump or restart is machine language dependent so, unfortunately, the existing breakpoint routines within monitors for the 8080 or 8086 are not applicable.

New routines tailored to the 8089 can be used, and, if done properly, can even be used to examine programs running on a REMOTE bus. Using breakpoints is somewhat complicated on the 8089 because the minimum instruction length is two bytes. There is no absolute CALL instruction, only a relative one (which would have to have its displacement recalculated each time it was used). But, with a several-byte absolute jump inserted at each place a breakpoint is desired, full breakpoint capabilities can be obtained.

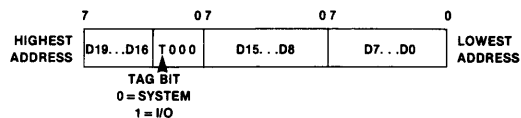
There are many ways the breakpoints can be implemented. When a breakpoint is reached, the 8089 itself could output the machine state to a console through its own routines. Better suited to debugging, though, is a system that has the 8089 place its machine state in memory, alert the host processor, and then halt. The host then picks up the 8089's state and can treat it in the same way it runs its own breakpoint routines. Since the host processor is more likely to be running a monitor or some other kind of debugging routine (and most likely has at least temporary console I/O), it is the logical system to initiate and examine 8089 breakpoints. If the IOP is running in the REMOTE mode, and the host processor has access to the I/O bus via the scheme mentioned in the hardware debugging section, then IOP programs running on the REMOTE bus can be examined.

The breakpoint itself can consist of an escape sequence that is used to save the TP value and jump to the save routine, or just a jump to the save routine. This routine saves all register contents for the channel the breakpoint is in, signals the host processor, and stops the IOP. All user programmable registers (GA, GB, GC, IX, MC, BC, TP), as well as the pointer tags, are accessible. The PP (Parameter Pointer) and PSW are not normally accessible, but if the generation of the CA is such that the IOP can send itself a CA, then by sending a CA HALT, the PSW will appear at PP + 3. Remember that

since the IOP doesn't have arithmetic or logical condition codes, the PSW is not as important as in other machines.

The most straightforward way to pass data from the IOP to the host processor is through the PB (Parameter Block) area since the PP will normally remain relatively fixed throughout the IOP program. In order not to infringe on the PB areas used by the programs, an area 18 bytes long should be allocated at the end of the PB block to hold the register contents. Using other areas to store the register data requires saving and reloading a pointer register as part of the breakpoint escape sequence.

The data returned from the breakpoint save routine will appear to the host processor as a sequential block of data in the PB area. Sixteen-bit data can easily be extracted, but 20-bit pointer data will have to be reconstructed from the move pointer (MOVP) format:



Several means are available to signal the host processor that a breakpoint has been reached. A bit could be set in memory or an interrupt sent to the CPU. The best way, though, is to use the BUSY flag (at CP + 1 or CP + 9). After starting the IOP, the BUSY flag is set to FF. When a breakpoint is reached, the IOP performs its save routine and does either a software or CA HALT. These result in clearing the BUSY flag, which then signals the CPU to obtain valid breakpoint data. The CPU can then restart the IOP by either a CA START or CA CONTINUE.

The breakpoint routine outlined above will work for a "one-shot" test. However, to be more useful as a general purpose debugging tool, some refinements must be added. To keep from destroying the program whenever a breakpoint is placed, the supervisory program running from the host processor must save the IOP code that is occupied by the escape sequence. When the breakpoint is completed and IOP execution is to resume, the host program restores the IOP code, sets the TP in the CB area back to where the breakpoint was placed, and sends a CA START. Since the length of each instruction can be easily found from bits 1-4 of the opcode, a single stepping function can also be done.* By the time this is implemented, the host program is becoming a full-fledged debugging routine. Appendix 3 describes a debugging program that makes use of the ideas presented here.

Breakpoint routines can be quite useful, but some restrictions and limitations should be mentioned. The processor examining the breakpoints must have access to the IOP program memory, either directly, or through IOP programs that simulate direct access. The program memory must be in RAM. The breakpoint must be

*The formula for length of instructions is: length (in bytes) = 2 + 1 (if bits 1,0 = 01) + 1 (if bits 3,2 = 01) + 2 (if bit 3 = 1) + 2 (if LPDI).

placed on an instruction boundary, and multiple breakpoints must not be placed so that they overlap. There may be some impact on the PB area. CA generation may have to be different than usual. But, despite these limitations, the breakpoints offer a useful and more conventional software debugging tool than analyzers.

REAL-TIME TESTING

Running an IOP program in its final environment with real I/O devices is the true test of dynamic operation. The program is no longer in a static, isolated environment. The demands of DMA and multiprocessing may reveal unplanned timing dependencies or critical section problems. There may also be sections of hardware or software, which couldn't be tested statically, that may have bugs. The whole purpose of static or functional testing is to dig these problems out while convenient debugging tools can be used. Since there are no simple techniques for real-time debugging, the use of a logic state analyzer and techniques to fully understand the IOP's real-time operation will be emphasized.

Multiprocessing operations and real-time asynchronous I/O requests can cause the timing complexity of the system as a whole to rise beyond the point of complete comprehension by an individual. It is then essential that techniques to ensure correctness are used. These include good design methods, especially a clean, well-structured design, as well as good testing. A thorough test requires the attitude that the system should be tested for failures, rather than tested for correctness. In other words, one should try to make the system fail, tests should be chosen that will put the worst stress on critical timing areas.

The best way to do this is to write a diagnostic program that puts the CPU, IOP, and I/O devices through the worst conceivable timing and program combinations. Ideally, the program should be self-checking so that it can be run without supervision, printing any data or program errors that occur, much like a memory test.

The two main real-time problem areas are insufficient data rates or latency, and critical section problems. To

test for data rate problems, run the system clock at its lowest expected frequency and use memory and I/O with maximum expected wait states. Identify the tightest program timings and try to have these sections coincide with worst case DMA or other heavy bus utilization (see dual channel operation later). Critical section problems can occur when two independent processors communicate with each other with improper "handshaking." This can result in one processor missing another's message, or even having both processors hang up, waiting for each other to go ahead. The 8089 provides aids to these problems, including the TSL instruction (to implement semaphores) and the BUSY flag. However, any interprocessor communication (including one channel of the IOP to the other) should be checked. Beware of cases when one processor is running considerably slower than the other (due to DMA overhead or chained instruction sequences).

The techniques for real-time debugging evolve from functional testing using a logic analyzer. For all but the simplest systems, an analyzer is essential, since it can graphically show program execution and timing relationships during real-time execution. Another aid is a delayed oscilloscope. Triggering the scope from the logic analyzer, the delay can be adjusted so that any signal in the system can be monitored.

To facilitate the use of the logic analyzer, especially if its memory is not very deep or when using it to trigger an oscilloscope, a repetitive system can be used to continually update the display. Using a repetitive reset helps to debug the software-hardware interface, since oscilloscope or logic analyzer probes can be readily moved around the circuit to observe new signals without manually retriggering the display. At its simplest, the reset to the host processor can be strobed, say every 10 ms. The processor will then provide the two channel attentions (CAs) that are needed to initialize the IOP. Where this isn't feasible, the CAs can be externally forced by either a string of one-shots or a simple processor with timing loops (such as a SDK-85 or SDK-86). See Figure 5 for initialization timing.

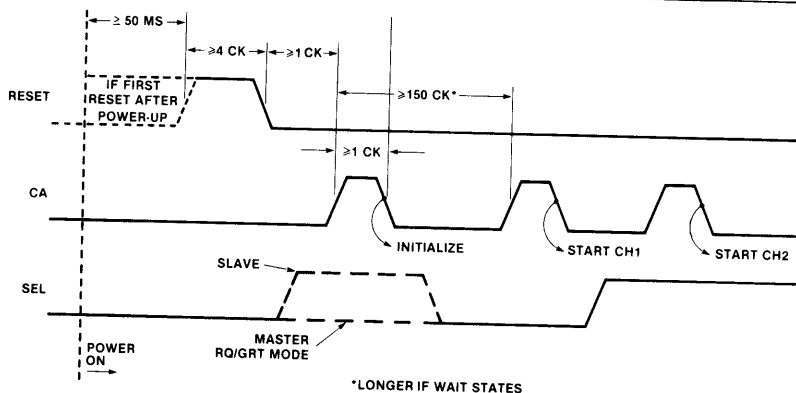


Figure 5. Initialization Input Sequence

Memory protection of the IOP and system programs is helpful when debugging DMA operation. It is quite easy for runaway DMA to wipe out memory. Another precaution to avoid this problem is to set an upper limit on the number of bytes transferred by always specifying a byte count termination.

Logic Analyzer Techniques

In the absence of other powerful debugging systems, the logic analyzer has shown to be an extremely useful tool. Because of its importance in debugging an IOP system, some basic techniques and observations that relate to monitoring IOP operation will be reviewed here. The particular brand or type of analyzer used is not too important, but would be desirable to have the following features:

- At least a 24-bit data width
- Flexible triggering and qualification control
- Display after triggering on a sequence of states
- Capability for hexadecimal data display

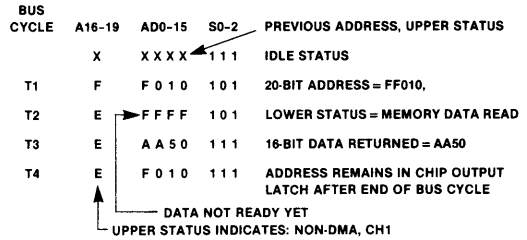
It is best to hook up to the address/data lines at the IOP, as opposed to looking at the separate address and data lines, since 39 lines would be required just to look at address, data and status lines. The three lower status lines should be monitored to show the type of bus cycle being run. Other lines can be connected where needed, at places like the DRQ lines, the EXT lines or other lines related to the system.

For general purpose debugging, triggering the analyzer on the rising edge of the IOP clock shows the most useful data concerning bus cycles. Of course, using the falling edge may be necessary to check certain signals, particularly ones that are active only while the clock is low. The following discussion is based on sampling data on the clock's rising edge.

One should be careful when setting up the triggering for the analyzer that the desired event is what is displayed and not a later event with the same trigger word. This can happen when the logic analyzer is in the repetitive trigger mode. It may retrigger before the system actually resets. A sequence restart feature is helpful.

The basis of following program execution and DMA on a logic analyzer is to follow an 8089 bus cycle, which is identical to a 8086 and 8088 bus cycle. The following diagram shows a typical 8089 bus cycle.

For general purpose debugging, displaying every clock is useful, but for quickly finding one's way around a program, the analyzer can be qualified so that only instruction fetches (status = 100 or 000), with ALE active, are trapped. A much more compact display of execution flow results.



As mentioned earlier, on a 16-bit bus, most instructions starting on odd addresses won't show the first fetch, since the internal queue is in use. It is a good idea in that case to use only even instruction boundaries as trigger words. When following dual channel operation, one should keep an eye on the upper status bits (S3-S6), since S3 indicates which channel is running (0 = CH1, 1 = CH2), and S4 indicates DMA/non-DMA transfer (0 = DMA, 1 = non-DMA).

A REVIEW OF IOP OPERATION (With things to look out for)

When trying to get an unfamiliar system going for the first time, it is too easy to stumble on apparent problems that are really just unexpected operation modes or peculiarities of the machine. For this reason the basic principles of IOP operation will be reviewed here with special emphasis on possible problem areas or pitfalls that a user might encounter when debugging a 8089 system. The topics are covered generally in the order encountered when bringing up a system. For complete details of operation and some design examples, see the 8086 Family User's Manual.

RESET

RESET must be active (HIGH) for at least four clocks in order to fully initialize all internal circuitry. On power up, RESET should be held high for at least 50 microseconds. The chip is only ready to accept a Channel Attention (CA) one clock after RESET goes inactive.

Note that the SEL pin is sampled on the falling edge of the first CA after RESET to tell the 8089 whether it is a master (0) or a slave (1) for its request/grant circuitry. If a master, it will assume it has the bus from the beginning. If a slave, it will strobe the $\overline{RQ}/\overline{GT}$ Line to request the bus back and will not start any bus transfers until it has been granted the bus. If the $\overline{RQ}/\overline{GT}$ line is not being used, make sure the IOP comes up in the master mode.

Initialization

Upon the first CA after reset, a sequence of instructions is executed from an internal ROM. These instructions pick up parameters and load data from the linked list sequence (Figure 6). The instruction sequence is essentially:

```

MOVb SYSBUS from FFFF6
LPD System Configuration Block (SCB) from FFFF8
MOVb SOC from (SCB)
LPD Control Pointer (CP) from (SCB) + 2
MOVbI "00" to CP + 1 (clears BUSY flag)
    
```

Remember that four bytes must be fetched during an LPD. If on a 16-bit bus, with even addressed boundaries, only two fetches are needed. Otherwise (8-bit bus or odd boundaries), four fetches are needed.

Even though no bus cycles are run to fetch these instructions, the CH1 Task Pointer (TP) appears on the address latches during the short internal fetch periods. On power up, this value is meaningless, but if a repetitive RESET is used, the TP remains unchanged from the end of the last program run. See Figure 6 for the start of a typical initialization sequence as viewed on a logic analyzer.

Bit 0 in the SYSBUS field sets the actual (or physical) system bus width that the IOP expects. In the 8-bit mode, only byte accesses are made, and all 8-bit data should appear on the lower eight data lines. In the 16-bit mode, word accesses can be made (if the address is even), all data on even addresses appears on the lower eight data lines, and all data at odd addresses appears on the upper eight.

Bit 0 in the SOC field sets the physical width for the I/O bus. The same rules for the system bus apply here. Note that these bits should reflect the actual hardware implementation and are not to be confused with the DMA logical widths set by the WID instruction.

The R bit (bit 1) in the SOC field is used to change the mode of the RQ/GT circuitry. When the IOP is on the same bus as an 8086, it is required to have the R bit be 0, with the 8086 as the master and the 8089 as the slave.

The master (8086 or 8088) can never take the bus away from the slave (8089); only the slave can give back the bus. In other words, during DMA transfers, the 8089 would not have the bus taken away. This is the only mode compatible with the 8086 or 8088.

When two IOPs are being used on the same bus, the RQ/GT circuitry can be put into an equal priority mode by setting the R bit to one. A slave can only be granted the bus if the master is doing unchained instructions or running idle cycles. The master can request the bus back from the slave at any time. The slave grants it if doing unchained instructions or if it is idling. The master and slave are put on essentially the same priority.

At the end of initialization, the "BUSY" flag of CH1 is cleared. For systems where the 8086 is waiting for the initialization sequence to end before giving another CA, it can set the BUSY flag high prior to initialization. The BUSY flag going low is a sign that the IOP is ready for another CA. It is important to remember that the IOP will not respond to, nor latch, a CA during an initialization sequence.

Channel Attention

The main system processor initiates communications with the IOP through the Channel Attention (CA) line. As mentioned earlier, the first CA after system RESET initializes the IOP. All subsequent CAs cause the IOP to do a two-step process. It first fetches the Channel Control Word (CCW) from the appropriate channel at (PP) for channel 1 or (PP + 8) for channel 2. (SEL at the time of CA falling determines the channel for all following actions.) The lower three bits of the CCW Command Field (CF) are examined and then cause the IOP to execute the desired function.

Command Field (CF)

Control of task block programs is accomplished through the command field. The various CF functions are:

CF

- 000 — Examine other field only and set BUSY flag
- 001 — Start task program in I/O space
- 011 — Start task program in system memory

The start command causes the following instructions to be executed out of the internal ROM:

- LDP CP from (CP) + 2 (CH1) or + 10 (CH2)
- LDP TP from (PP) (for TP in system) or
- MOVB TBP from (PP) (for TBP in I/O)
- MOVBI "FF" to (CP) + 1 or + 9 (set BUSY flag)
- 111 — HALT channel. BUSY flag cleared to "00"
- 110 — HALT channel. Save state of machine and clear BUSY flag by executing:
 - MOVP TP to (PP)
 - MOVB PSW to (PP) + 3
 - MOVBI "00" to (PP) + 1 or + 9

CA	A ₁₉ -A ₀	S ₃ -S ₀	T	COMMENTS
1	FFF F	111		Trigger CLK ↓
1	FFF F	111		
	FFF F	111		
	FFF F	111		
	E00 0	111		Bus un-tristated
	E00 0	111		
	FFC 6	D 111		
10				TP to latch
CK	FFC 6	D 111		
	FFF F	6 101	T1	Address loaded to latch
	FFF F	F 101	T2	Data not ready yet (nothing on bus)
	FFF 0	1 111	T3	SYSBUS loaded into chip (01)
	FFF F	F 111	T4	Nothing on bus
	FFF F	6 111		After bus cycle, address remains in latch
	FFF F	6 111		TP is loaded to latch, even though fetches are from internal ROM
	FFC 6	D 111		
14				Address to latch
	FFC 6	D 111		
	FFF F	8 101	T1	
	FFF F	F 101	T2	
	FFF 0	1 111	T3	1st 2 bytes of LPD data fetched (FFF0)
	FFF F	F 111	T4	
	FFF F	8 111		
	FFF F	8 111		
	FFF F	8 111		
	FFF F	8 111		
	FFF F	8 111		
	FFF A	101		
	FFF F	A 101		
	FFF A	111		2nd 2 bytes of LPD data fetched (FFFA)
6				
	FFF A	111		
6	FFF A	111		
CK	FFC 6	D 111		

Figure 6. Start of Initialization Sequence On a 16-Bit Bus

The channel will HALT and the machine will continue execution on the other channel or go to idle if the other channel is idle.

101 — Continue channel. The channel is revived after a HALT by executing:

```
MOVP TP from (PP)
MOVB PSW from (PP) + 3
MOVBI "FF" to (CP) + 1 or + 9
(set BUSY flag)
```

Do not do a CONTINUE after initialization without doing a CA START first since the (PP) register in CH1 is used as a temporary register (to hold SCB) and is only correctly loaded by a CA START.

The upper 5 bits in the CCW will have affect if CF = 000 or upon a CA START. Some things to note about these upper fields are:

- **Priority Bit** — If both channels are doing tasks of the same overall priority, the tasks with the higher priority bit will run. If the priority bits are the same, execution will alternate between the two channels.
- **BLL Bit (Bus Load Limit)** — Keeps nonchained instructions from occurring more often than once every 128 clocks. However, channel attention or termination cycles, even on the other channel, may disrupt the exact time interval to the next instruction.

It should be noted that the setting or clearing of the BUSY flag occurs after the loading or storing of registers, so that in a system where the main CPU uses the BUSY flag as a form of semaphore to tell when the IOP is truly finished, there is no danger that the SCB, CP, PP or TP could be changed before the IOP loads them.

Also since DMA termination cycles and chained instruction execution have a higher priority than CA, it is possible for CA to be "shut-out" by these higher priorities running on the other channel. However, since CA is always latched (except during initialization), it won't be forgotten.

How Can a Channel be Halted?

Sometimes a channel may stop its operation unexpectedly. To see what could cause this, and to show the impact of halting a channel, the various ways of stopping a channel are explained:

HALTED CHANNEL — If the channel has never started after initialization, if it has received a CA HALT command or a software HALT, channel operation is suspended. If the other channel can run, it will, otherwise idle cycles will run. Only a CA START or CONTINUE can resume operation.

WAITING FOR A DMA REQUEST — If the channel is in a source or destination synchronized DMA transfer mode, it will wait until DRQ is active before running its synchronized transfer. To minimize the impact on the overall throughput of the chip, the other channel can run during these DRQ wait periods.

WAITING TO GET THE BUS BY $\overline{RQ}/\overline{GT}$ — If the IOP has given the bus away via $\overline{RQ}/\overline{GT}$, it won't initiate any bus transfers until it has the bus back. The machine will run up to just before T1 of a bus clock cycle and will three-state its address/data and status pins until it has been granted the bus.

WAITING FOR READY — When running bus transfers, READY is sampled at T3 of a busy cycle. If inactive, the whole chip will wait until READY goes active.

The last two cases of waiting (or "wait" states) stop the whole chip and do not permit the other channel to run. However, with READY inactive or with the bus not acquired, there is not much that can be done on the other channel anyway. These two cases only stop the chip when running bus cycles. Any internal operations can proceed without having the bus or with the system not READY.

Note the difference between when the chip is HALTED when using $\overline{RQ}/\overline{GT}$ and an external arbiter (8289) for bus arbitration. Not having the bus due to $\overline{RQ}/\overline{GT}$ will inhibit the bus cycle from even starting. Since the 8289 stops the chip by forcing AEN inactive, which goes through the 8284 clock generator to force READY inactive to the IOP (or 8086/8088), a bus cycle has already been started, with ALE asserted, and the address on the address/data lines. When the bus is obtained, operation proceeds at T3 of the bus cycle.

As will be mentioned later, many invalid opcodes will cause the machine to hang up. In these cases the address/data lines will point to where the bad opcode was fetched.

Task Execution

Although optimized for fast and flexible DMA operation, the IOP is also a full-fledged microprocessor. The 8086 Family User's Manual deals with programming strategies and other details. Some of the things to be noted during debugging will be mentioned here.

Instruction Fetching

Unlike the 8085 (but like the 8086), the 8089 labels all fetches from the instruction stream, whether OPCODE, offset, displacement, or literal data, as an instruction fetch on the status lines. In some cases, such as MOV R,I and ADD R,I, the instruction fetch time greatly exceeds execution time because literals are treated as instruction fetches. When following programs on a logic analyzer, triggering on status = 100 or 000 (instruction fetch) and a known program address is the handiest way to trace the flow of the program.

When running programs on a 16-bit bus, a 1-byte queue register comes into play, saving the upper byte fetched from the last instruction fetch, if not used by the previous instruction. This reduces fetch time and bus utilization since the odd byte doesn't need to be fetched again. An internal four-clock cycle fetches data from the queue. Like the internal ROM fetches, the task pointer is put out on the address/data lines, but no bus cycle is run.

The queue can have some possible unexpected affects that have to be taken into account during debugging. These apply only to 16-bit systems and are:

1. Instructions that start on odd boundaries will not likely have bus cycles run to fetch the odd byte unless jumped to, unless preceded by LPDI (which clears the queue), or an instruction that modifies the task pointer is executed. The latter causes the queue to be cleared so that part of an old instruction won't become part of the new one.
2. There is a queue register for each channel so loading or clearing the queue on one channel has no affect on the other channel's queue.
3. The second word of immediate data fetched by a LPDI is done during a pseudo-instruction fetch cycle that cannot make use of the queue or already fetched data. Thus, if on an odd boundary, fetching an LPDI will be byte, word, byte, byte, byte, and the queue will not be loaded.

When Can the Other Channel Interrupt Instruction Execution?

This will be explained more in the "dual channel" operation section, but a few points will be mentioned here. All instructions are made up of internal cycles, with each cycle composed of two to eight clocks. Each bus cycle is one internal cycle, but there can be internal cycles with no communications to outside the chip. Internal cycles will be extended by the number of wait states in each bus cycle. Between any of these cycles, DMA from the other channel can intervene if the priorities permit it. Instruction fetching and execution can only interrupt instructions on the other channel when the instruction has been completed, not between internal cycles.

Registers

All the registers have some special purpose use in the Instruction Execution or DMA, but all except the CC register can be used as general purpose registers during instruction sequences. A few are loaded specially:

- *CP* — Is only loaded during an initialization sequence. There is one CP register that handles both channels. (All others are duplicated, one set for each channel.)
- *PP* — Is only properly loaded during a CA START command. It holds the SCB value after the initialization sequence.
- *TP* — This is included as part of the registers in the RRR field, but cannot be operated on unless you plan on having your program execution jump around. Everytime this is operated on, the queue is cleared. The TP is loaded from two words (address and displacement) on a CA START, LPD, or LPDI, and loaded from 3-byte MOVDP format (see illustration on page 5) on a CA CONTINUE, and can be operated on using any register oriented instructions.

The following registers are loaded during program execution, but can have special effects:

- *CC* — The only thing that affects instructions in the CC register is the chaining bit. If chaining doesn't matter (if only one channel is being used without channel attentions, for example), then the CC register can be general purpose. However, for portability of programs, it is strongly suggested not to use the CC register except for altering DMA parameters and chaining.
- *MC* — Is a general purpose 16-bit register, but is also used to do a masked comparison either for DMA search/match termination or for the JMCE and JMCNE instructions.
- *BC, IX* — Both general purpose 16-bit registers. In instructions that reference memory using the AA field, if AA = 11, the IX register is incremented by the number of bytes fetched or stored.
- *Pointer Registers (GA, GB, GC and TP)* — Are 20-bit registers, but can also be used as 16-bit registers. Adds will carry into the upper 4 bits, but other operations (COMP, OR, AND) are done only on the lower 16 bits. Note that when used as pointers to system memory, it is possible to add a large 16-bit number to the pointer and to put the pointer into another 64K block of memory.

Sign Extension

All program data brought into the chip, either literals or displacements in opcodes, or program data fetched from memory, is sign-extended. Offsets used for calculating addresses are not sign extended. Any 8-bit data brought in has bit 7 sign-extended up to bit 19. Sixteen-bit data is sign-extended from bit 15 to bit 19. It is important to note this, because it can affect logical operations. For example, if one wanted to OR 0084H with 1234H in register GC, you couldn't do ORBI GC, 84H, because bit 7 would sign-extend into the upper byte. Instead, you should code ORI, 0084H to do this properly (note that this has a word for the immediate data). The non-ADD operations will cause the upper four bits of the pointer registers to be invalid since the upper four bits of the ALU come only from the adder.

Tags

It should be noted that the way the IOP knows which bus to access (system or I/O) is via the Tag bit associated with the pointer register used. The TAG can only be set in these ways: loading as a 16-bit register (MOV R,M, MOV R,I) sets TAG to I/O space, loading as a pointer (LPD, LPDI) sets TAG to a system space), or bringing the TAG in from memory by a MOVDP instruction.

Effects of Invalid Opcodes

The upper 6 bits of the 2-byte opcode actually determine which opcode will be executed. If these bits are a valid opcode, but lower bits are invalid, the chances are good that the bad bits will be ignored. But if the upper six bits are invalid, there is a very good chance that the chip will hang up and stop execution in that channel. The only way to get out of this mode is to reset the chip. If this hang-up occurs, it can usually be traced because the last address of the instruction fetch will still be on the

address/data lines, showing where the program went astray.

Going from Instruction Execution into DMA

The XFER instruction places the current channel into the DMA mode after the next instruction. This permits one last instruction to start up an I/O device (start CRT display on an 8275, for example). However, in order for the IOP to get setup for DMA, the GA, GB, and CC registers should not be altered during this last instruction. Failure to observe this will probably result in an improper first DMA fetch. The WID instruction can be placed after XFER.

DMA Transfers

Incrementing/Non-Incrementing pointers

A memory or I/O pointer can be made to increment for each byte transferred during DMA or it can remain fixed. Incrementing is used primarily for memory block transfers, and non-incrementing is used to access I/O ports.

B/W Mode

Each DMA transfer is composed of separate fetch and store cycles so that 8/16-bit data can be assembled and disassembled, and translation and termination may also be easily handled. There are four possible transfers or B/W modes. They are:

- B – B — 1 byte fetched, 1 byte stored
- B/B – W — 2 bytes fetched, 1 word stored
- W – B/B — 1 word fetched, 2 bytes stored
- W – W — 1 word fetched, 1 word stored

The B/W mode used depends on the logical bus width (selected by the WID instruction), address boundary, and incrementing mode.

All systems with 8-bit physical buses will run in the B/B mode. On 16-bit physical buses the other modes are possible, depending on the logical widths selected. Note that the logical bus width can be different than the physical bus width since there are cases where an 8-bit peripheral may be used on a 16-bit bus. The selection of the logical width, and not the physical width, is what determines the B/W mode. Thus it is the responsibility of the programmer not to program an invalid combination (i.e., don't specify a 16-bit logical width on an 8-bit physical bus).

Any transfer on an odd boundary will be B/B but if the pointer is incrementing and on a 16-bit logical bus, after the first transfer, the pointer will be on an even boundary. The IOP will then try to maintain word transfers in order to transfer data as efficiently as possible. See the user's manual for details. The change in B/W mode occurs only after the first transfer or, as explained in the termination section, upon certain byte count terminations.

Synchronization

In the unsynchronized mode, transfers occur as fast as priorities will allow. This is the IOP's "block-move" mode. Most I/O peripherals only want a DMA transfer on demand; the DRQ lines, along with synchronization specified, will handle this need. Source synchronization

is used for I/O reads and destination synchronization is used for I/O writes.

If the IOP is waiting for a DMA request, it will run programs or DMA on the other channel, or execute idle cycles if nothing is pending. If running idle cycles when the DRQ comes, the transfer starts five clocks after DRQ is recognized. If running DMA or instructions on the other channel, the DRQ cannot be serviced until the current internal cycle is done, and may require a maximum of 12 clocks (without bus arbitration or wait states).

Consecutive DRQ-synchronized DMA transfers on the same channel are separated by four idle clocks (assuming no other delays) by an internal sampling mechanism. This happens between the 2-byte fetches on source-synchronized B/B-W cycles, and between the two stores on destination-synchronized W-B/B cycles. This delay between consecutive DMA cycles allows adequate time for proper acknowledgement of the current DMA request before the next request is processed. On destination-synchronized DMA, this isn't a problem, but on source-synchronized DMA, there will be four extra clocks per transfer. Unless one is running right at the speed limit, this won't be a problem. Near the maximum data rate, unsynchronized transfers can be used, with synchronization done by manipulating the READY line.

Translate Mode

When the translate bit is set, the data fetched during DMA will be added to the GC register. This new pointer will in turn be used to fetch, via a seven clock extra fetch cycle, new data, which will then be stored. Translate is only defined for byte transfers. The bytes are added to GC as a positive offset, so a lookup table for translating data can be a maximum of 256 bytes long. Even if the data to be translated falls within a smaller range (such as ASCII code), a full 256-byte lookup table is recommended so that erroneous data can be flagged and controlled.

Translate can be run on any of the B/B transfer modes, so it is useful for doing block translation within program execution as well as translation directly to or from an I/O port.

DMA Termination

One of the powerful features of the IOP is its varied DMA termination conditions and their close tie-in with resuming Instruction Block programs. However, because of the multitude of DMA modes, care must be taken in predicting the exact termination parameters. Various things to be careful about will be outlined here.

Byte Count (BC) Termination

The BC register is decremented for every byte transferred whether or not BC termination is set. If BC termination is set, the last transfer done is the one that results in BC being zero. To avoid the problem of missing BC = 0 on word transfers, if BC is odd between every transfer, the IOP detects when BC is 1, and forces the last transfer to be in the B/B mode. Since both the fetch and store cycles are complete, the source and destination pointers point exactly to the next byte or word that would have been fetched.

Masked Compare (MC) Termination

An MC termination occurs when a pattern matches (or doesn't match, depending on mode selected) the lower half of the MC register (the match pattern) with only the bits that are enabled by the upper half of MC (the mask pattern) contributing to a match. Thus the masked bits can be "don't cares" in both the data byte and the match byte.

The masked comparison is only done on store (deposit) cycles. Any bytes transferred (in B/B or W-B/B mode) will be compared. But, since the MC comparison is done on only one byte, any words stored (W-W or B-B/W) have only their lower byte compared. This may be fine, but if not, make the destination logical width 8 bits.

Just like BC termination, the pointers will point to the next data to be transferred. The BC will also be decremented correctly, except if the termination occurs on the first byte of a W-B/B transfer. In this case the BC will be decremented as if the entire transfer (both bytes) had taken place.

The store cycle that causes an MC termination will be lengthened by two extra clocks (or by one extra clock if there are wait states), to allow time to set up the termination cycle.

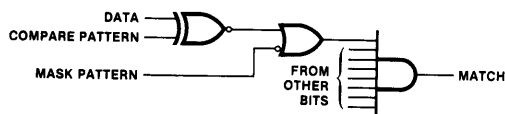


Figure 7. Masked Compare Logic for 1-Bit

External (EXT) Termination

External termination allows the I/O device or controller to use its own conditions to generate a termination. Basically, the IOP will halt DMA as soon as it recognizes an EXT terminate, even if a transfer is only partially complete. There might be concern that multibyte cycles (W-B/B or B/B-W) might have data lost if an EXT terminate stopped the store cycle. In unsynchronized DMA this would happen, but this mode is typically not used with I/O controllers that could generate external terminations. In synchronized DMA modes, it is assumed that the I/O controller will only do a DRQ for valid data transferred, and that it won't give an EXT terminate with its DRQ active. In destination synchronization, the possible problem occurs in the W-B/B mode, where EXT terminate comes after the first store but before the second. This is fine, since even though data was overfetched, the proper amount was actually transferred. In source synchronization, the B/B-W mode raises problems since if an EXT terminate came after the first byte fetched and before the second byte fetched, normally no store cycles would be done at all, thus losing the first byte fetched. In this case (i.e., source synced, DRQ inactive, and 1 byte already fetched), a single byte store cycle is run before the termination cycle, ensuring data integrity.

In order to prevent an invalid signal level from becoming trapped from the asynchronous EXT term lines, two clocks of delay and signal conditioning are done on these lines. In addition, a termination cycle can only be started at certain times during DMA (or TB on the other channel — see dual channel operation section). The EXT terminate lines should be valid eight clocks before the start of the DMA cycle to be stopped.

EXT is sampled even when the IOP is running something on the other channel. Remember though, that despite the high priority of termination, the current instruction on the other channel has to finish before the termination cycle is run. Simultaneous EXTs on both channels result in CH1 termination being done first.

In order to have enough time to process a byte count termination, the BC register is always decremented during DMA fetch cycles. Because of this, external or MC terminations that occur during W-B/B cycles will result in the byte count always being decremented by two, even if only one byte is stored. This also occurs in the block-to-block or block-to-port B/B-W modes. To find the exact number of bytes transferred, the source pointer address can be checked in the block-to-port and block-to-block modes during B/B-W cycles and in the block-to-port W-B/B mode. The destination pointer address can be used to find the number of bytes transferred in the port-to-block and block-to-block modes during W-B/B cycles.

Termination Cycles and Multiple Terminations

Upon termination, the user can run different task block programs, depending on which type of termination has occurred, by specifying an appropriate termination offset. That is, instruction fetching will begin after a termination cycle starting at either the TP value before the DMA started, TP + 4 or TP + 8. These offsets permit long or short jumps to termination routines.

The termination cycle is an add immediate instruction that runs from the internal ROM and adds the proper offset to the TP. It is 15 clocks long for TP + 4 and TP + 8 termination and 12 clocks long for TP + 0 termination.

As mentioned earlier, EXT terminate must come a certain time before the end of a transfer to ensure that the next transfer doesn't start. If it comes in time and MC termination also occurs on the current transfer, then the termination cycle with the largest offset is run. A simultaneous BC terminate cycle will have priority over MC and will result in the running the BC termination program.

Priorities/Dual Channel Operation

The IOP can share its internal and external hardware between two separate channels. The user sees two identical IOP channels with all registers, machine flags, etc., independent of the other channel. The only register in common is the CP register, loaded by the initialization sequence. The mechanism for achieving dual channel operation is time multiplexing between the two channels.

Since interleaving two channels affects their response time to external events and since interfacing to these events is the prime purpose of the IOP, several means of adjusting the priorities of the channels are provided.

Before going into the priority algorithms in detail the four types of cycles that are affected by the priorities will be outlined:

1. *DMA Cycles* — Any type of DMA transfer cycle, including single transfers and translate cycles. DMA can be interrupted after any bus transfer by the other channel.
2. *Instruction Cycles* — Any instructions that have been fetched out of I/O or system memory. Instruction cycles are made up of internal cycles, each two to eight clocks long (assuming no wait states). Some cycles may not run bus transfers. Instructions can be interrupted by DMA after any one of the internal cycles, but can only be interrupted by instructions on the other channel (normal ones or ones from internal ROM) after the current instruction is completed.
3. *Termination Cycle* — Performed when DMA transfers end and instructions resume (except on single transfers).
4. *Channel Attention Cycles* — Performed when channel attention is given, performs actions specified in the CCW field. Both termination and CA cycles can be interrupted by DMA after any internal cycle, but can only be interrupted by instruction cycles after the complete sequence of internal cycles is done.

Termination and channel attention cycles as well as the initialization cycle (which never runs concurrently with other operations) are sequences of instructions fetched from an internal ROM.

Recognizing the higher importance in doing DMA, termination and (to a lesser extent) CA cycles, the following priority scheme is built into the IOP. Any channel that has a higher-priority operation will run continuously until done. If both channels are running the same priority, execution will alternate between them.

Highest Priority

1. DMA transfers, termination, chained instructions
2. Channel attention cycles
3. Instruction cycles
4. Idle cycles

Lowest Priority

Two ways exist to alter the priority scheme. One way is to utilize the priority bits for each channel. If one is greater than the other, that channel will run at the expense of the other if both channels are otherwise running at the same priority. Thus the P bit only has effect on channels running at the same priority level.

If one wants to run instructions along with or in place of DMA on the other channel, the other technique is to set the chaining bit (in the CC register) which brings the instruction priority up to the level of DMA. Care should be taken with this since now CAs are at a lower priority than instructions and will not be serviced unless that

channel goes idle. Chaining will also lock out normal instructions on the other channel. Chaining should thus be used with care.

In order to reduce the possibility of shutting out channel attentions, an exception is made to the above priority scheme. After every DMA transfer, whether synchronized or unsynchronized, the IOP will service any pending CA. However, chained task block execution will still shut out CAs on the other channel.

What is the importance of priorities? Well, as an example, let's say that we are running long periods of non-time-critical block moves (via DMA) on one channel and running short bursts of DMA that must be serviced promptly on the other channel. With the default priorities, the short DMA channel bursts would be interleaved with the longer DMA, reducing the maximum transfer rate for both channels. If, however, the priority bit was one on the burst mode DMA and zero on the other, the bursts would be serviced continuously at the fastest possible data rate.

An even more critical case would be the same low priority, long DMA transfers on one channel with DMA on the other channel that must terminate, run a short instruction sequence, and resume DMA again within a short, fixed time. (This might be the case in running a CRT display with linked list processing between lines.) Normally, the low priority, long DMA could indefinitely block the short TB sequence. By setting the high-priority channel's priority bit to one and putting it into the chained instruction mode, the low priority channel would stop its DMA entirely so that the termination/instruction sequence could run.

When establishing the priorities to be run, care should be taken that both channels will run successfully under a worst case combination. This can be tricky when the channels are running asynchronously with fast data rates and/or short latencies, but must be taken into account. Of course, running only one channel on the IOP is an easy solution, but if more than one IOP is being used in the system, the priorities and delays of the bus arbitration used (either $\overline{RQ}/\overline{GT}$ or an 8289 bus arbiter) must be taken into account. It may be found that the on-chip arbitration between the two channels is faster and more powerful than external arbitration.

SUMMARY

It is hoped that the material presented here will aid those who are putting together and debugging an 8089 IOP system, and help them in understanding the operation of the IOP. Many of the debugging techniques should be familiar to those who have worked with micro- and minicomputer systems before. Other debugging techniques not mentioned here, which work well with microprocessor systems, could be just as applicable to the 8089. The unique nature of the IOP among LSI devices warrants special consideration for its I/O functions and multiprocessor capabilities.

Appendix I

CHECKLIST OF POSSIBLE PROBLEMS

HARDWARE PROBLEMS

- Is RESET at least four clocks long?
- Are both V_{SS} lines connected to ground?
- Does the first CA falling edge come at least two clocks after RESET goes away?
- Does the second CA come at least 150 clocks (16-bit system, no wait states) after the first CA?
- Is READY correctly synchronized and gated by local/system bus lines?
- Is SEL correct for first CA so that IOP comes up correctly as master or slave?
- If two IOPs are local to each other, is a 2.7K pull-up resistor used on RQ/GT?

SOFTWARE PROBLEMS

- Are the initialization parameters in the initialization linked-list correct?
- Is BUSY flag being properly tested by host CPU software before modifying PB or providing a new command?
- Has the chaining, translate, or lock bit in the CC register been erroneously set?
- Have DMA termination conditions been met? The IOP could be trying to do endless DMA.

Appendix II

BREAKPOINT ROUTINE AND CONTROL PROGRAM

The debugging program described here is an example of the kind of software development tool that can be developed for the 8089 IOP. It was written to try out various breakpoint schemes, and has been used to debug an engineering application test system. The program is not meant to be the ultimate debugging tool, but is an example of what can be put together to utilize the breakpoint routine described earlier in the application note.

The debugging program was tested on a 8086-based system that emulates the SDK-86 I/O structure, and uses the SDK-86 serial monitor. This enables it to use the SDK-86 Serial Downloader to interface to an Intellec® development system on which the software was created. The 8086 system is interfaced via a MULTIBUS™ interface to an IOP running in the REMOTE mode. The remote bus access technique, mentioned earlier in this note, is implemented on this system, but was not used in the software debugging program.

The breakpoint routine uses a simple jump to a save routine. The PL/M-86 supervisory or control program handles the placement of the jump within the users program. Since it can not normally access the remote bus, all IOP programs to be tested must run out of system memory.

When the control program starts, it assumes the IOP has just been reset. It then prompts the user for the CP and PP values. After this, it sends the first (initialization) channel attention. It then asks the user for the channel to be run, and the starting and stopping addresses. After the stopping address has been entered, a Channel Attention Start is given. If the breakpoint is reached, a HALT is executed, and the control program prints the register contents. If the breakpoint hasn't been reached, the user can type any character, and a Channel Attention Halt will be sent to the IOP. If the IOP responds within 50 ms, the TP where it was halted is printed. Otherwise, the control program issues an error message. If, at any time, the user wants to get out of the program, typing an ESC will pass control back to the SDK-86 monitor. Figure 9 shows the flow of the control program.

Note that, unlike a single CPU debugging routine, having the 8086 supervise the 8089 enables a clean exit from crashed IOP programs. The program code where jumps had been placed are always restored. The control program is a good example of how the power of dual processors can be put to good advantage.

Comments within the control program indicate parameters that need to be changed to run on different systems. It should be noted that channel attentions are invoked by the recommended method of using an I/O write to a port to generate CA and using A0 for SEL.

Source and object files of this program are available through Intel's INSITE™ User's Program Library as program 8089 Break. 89 (number AD6).

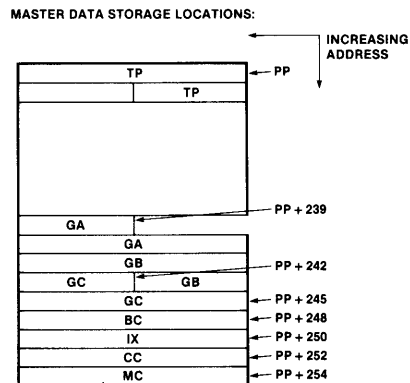


Figure 8. Breakpoint Routine to Run 8089 Program out of System Memory

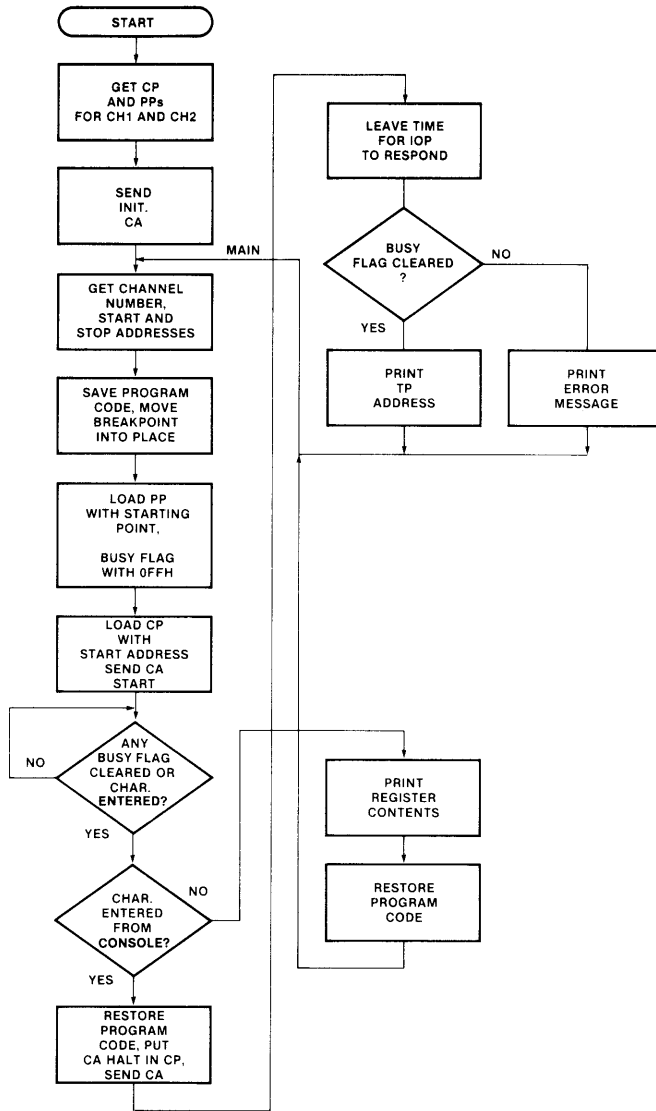


Figure 9. Breakpoint Routine to Run 8089 Program out of System Memory

PL/M-86 COMPILER BOB9 BREAKPOINT ROUTINE

PAGE 1

ISIS-II PL/M-86 X103 COMPILATION OF MODULE BREAKPOINT
 OBJECT MODULE PLACED IN BREAK OBJ
 COMPILER INVOKED BY: F1:PLM86 BREAK.SRC PAGEWIDTH (100)

```

$TITLE ('BOB9 BREAKPOINT ROUTINE')
/* .....
BOB9 BREAK POINT PROCEDURE
WRITTEN BY DAVE FERGUSON 2/2/79   REV 2   8/14/79
INTEL CORPORATION
..... */

1      BREAK$POINT:
2      DD;
3      1      DECLARE I BYTE;
4      1      DECLARE SAVECODE (4) WORD; /*BUFFER FOR STORAGE*/
5      1      DECLARE ONEPP POINTER; /* CHAN ONE PP */
6      1      DECLARE TWOPP POINTER; /* CHAN TWO PP */
7      1      DECLARE STARTBYTES (4) BYTE; /* BUFFER FOR START ADDRESS */
8      1      DECLARE STARTPOINTER POINTER; /* POINTER FOR START ADDR. */
9      1      DECLARE ENDPONTER POINTER; /* POINTER FOR END ADDR. */
10     1      DECLARE PRESENT POINTER AT (@INPNTR); /* POINTER BUFFER */
10     1      DECLARE TRUE LITERALLY 'OFFH',FALSE LITERALLY 'OOOH';

/* YOU MUST CONFIGURE YOUR I/O STRUCTURE AND
SYSTEM TO MATCH THE PROGRAM OR VISA VERSA */
11     1      DECLARE CRSTATUS LITERALLY 'OFFF2H', /* 8251 STATUS PORT */
CRDATA LITERALLY 'OFFFOH', /* 8251 DATA PORTS */
CHANATTEN LITERALLY 'OFAH', /* CHANNEL ONE CHANNEL ATTENTION PORT */
/* CHANNEL TWO CHANNEL ATTENTION PORT = CHANATTEN + 1 */
CHANNELONE LITERALLY 'OOH',
CHANNELTWO LITERALLY 'OIH',

/* ASCII IS A STRING OF HEX CHARACTERS IN ASCII FORM */
ASCII (*) BYTE DATA ('0123456789ABCDEF'),
TITLE$STRING (*) BYTE DATA (OAH,ODH,'BOB9 BREAKPOINT VER 1.0',
OAH,ODH,'TYPE ESCAPE TO RETURN TO MONITOR.',
OAH,ODH,0),
CHANGIVEN (*) BYTE DATA ('CHANNEL ATTENTION GIVEN TYPE ANY KEY TO ABORT.',
OAH,ODH,0),
BKREACHED (*) BYTE DATA (OAH,ODH,'BREAKPOINT REACHED',OAH,ODH,0),
GETCP (*) BYTE DATA ('INPUT CP IN HEX',OAH,ODH,00),
GET$PP (*) BYTE DATA ('INPUT PP IN HEX FOR ',OOH),
GETSTART (*) BYTE DATA (OAH,ODH,'INPUT STARTING ADDRESS IN HEX',OAH,ODH,OOH),
STOPADDR (*) BYTE DATA ('INPUT END ADDRESS IN HEX',OAH,ODH,OOH),
CHANNELNUMBER (*) BYTE DATA (OAH,ODH,'CHANNEL ONE OR TWO? ',OOH),
ABORT (*) BYTE DATA (' FATAL ERROR - IOP DOES NOT RESPOND TO CHANNEL',
' ATTENTION. RE-INITIALIZE SYSTEM ',0),
ABORTAT (*) BYTE DATA (' TP WAS ',0),
ONE (*) BYTE DATA (' CHANNEL ONE',OAH,ODH,OOH),
TWO (*) BYTE DATA (' CHANNEL TWO',OAH,ODH,OOH),
GASTRING (*) BYTE DATA ('GA = ',OOH),

```

PL/M-86 COMPILER 8089 BREAKPOINT ROUTINE

PAGE 2

```

GBSTRING (*) BYTE DATA ('GB = ',00H),
GCSTRING (*) BYTE DATA ('GC = ',0H),
BCSTRING (*) BYTE DATA ('0AH,0DH, 'BC = ',00H),
IXSTRING (*) BYTE DATA ('0AH,0DH, 'IX = ',00H),
CCSTRING (*) BYTE DATA ('0AH,0DH, 'CC = ',00H),
MCSTRING (*) BYTE DATA ('0AH,0DH, 'MC = ',00H)

12 1  DECLARE CHAR BYTE;
13 1  DECLARE ONETWO BYTE;

/* SDKMON IS A PLM TECHNIQUE USED TO FORCE THE CPU INTO AN
   INTERRUPT LEVEL 3. IN ORDER TO USE THIS THE PROGRAM MUST
   BE COMPILED (LARGE). */
14 1  SDKMON;
      PROCEDURE;
15 2  DECLARE HERE (*) BYTE DATA (0CCH),
      /* THIS IS AN INT. 3 */
      WHERE WORD DATA(.HERE);
16 2  CALL WHERE;
17 2  END;

/* CO SENDS A CHAR TO THE CONSOLE WHEN READY */
/* THIS ROUTINE IS WRITTEN TO RUN VIA THE SERIAL
   PORT OF AN SDKB6 */
18 1  CO:
      PROCEDURE (C);
19 2  DECLARE C BYTE;
20 2  DO WHILE (INPUT(CRTSTATUS) AND 01H) = 0; END;
22 2  OUTPUT (CRTDATA) = C;
23 2  END;

/* CI GETS A CHARACTER FROM THE USER VIA THE SERIAL PORT */
/* CI AUTOMATICALLY ECHOS THE CHARACTER TO THE USER CONSOLE */
/* DECLARE ESCAPE LITERALLY '1BH';
24 1  CI:
      PROCEDURE BYTE;
25 1  DO WHILE (INPUT(CRT*STATUS) AND 02H) = 0; END;
26 2  CHAR = INPUT (CRTDATA) AND 07FH;
28 2  CALL CO(CHAR);
29 2  IF CHAR = ESCAPE THEN CALL SDKMON; /* GO TO SDK MONITOR */
30 2  RETURN CHAR;
32 2  END;
33 2  END;

/* VALIDHEX CHECKS THE VALIDITY OF A BYTE AS A HEX CHARACTER*/
/* THE PROCEDURE RETURNS TRUE IF VALID FALSE IF NOT */
34 1  VALIDHEX:
      PROCEDURE (H) BYTE;
35 2  DECLARE H BYTE;
36 2  DO I=0 TO LAST(ASCII);
37 3  IF H=ASCII(I) THEN RETURN TRUE;
39 3  END;
40 2  RETURN FALSE;
41 2  END;

```

PL/M-86 COMPILER 8089 BREAKPOINT ROUTINE

PAGE 3

```

/* HEXCONV CONVERTS A HEX CHARACTER TO BINARY FOR MACHINE USE.
   IF THE CHARACTER IS NOT A VALID HEX CHAR, THE PROCEDURE RETURNS
   THE VALUE OFFH */
42  1  HEXCONV:
      PROCEDURE (DAT) BYTE;
43  2      DECLARE DAT BYTE;
44  2      IF VALIDHEX(DAT) <> OFFH THEN RETURN TRUE;
45  2      DO I=0 TO LAST(ASCII);
46  2          IF DAT = ASCII(I) THEN RETURN I;
47  3      END;
48  2      END;

/* HEXOUT WILL CONVERT A VALUE OF TYPE BYTE TO AN ASCII STRING
   AND SEND IT TO THE CONSOLE */
51  1  HEXOUT:
      PROCEDURE(C);
52  2      DECLARE C BYTE;
53  2      CALL CO(ASCII(SHR(C,4) AND OFH));
54  2      CALL CO(ASCII(C AND OFH));
55  2      END;

/* WORDOUT CONVERTS A VALUE OF TYPE WORD TO AN ASCII STRING
   AND SENDS IT TO THE CONSOLE */
56  1  WORDOUT:
      PROCEDURE (W);
57  2      DECLARE W WORD;
58  2      CALL HEXOUT(HIGH(W));
59  2      CALL HEXOUT(LOW(W));
60  2      END;

/* GETADDRESS IS A PROCEDURE TO GET AN ADDRESS FROM THE CONSOLE.
   THIS PROCEDURE WILL ONLY CONSIDER THE LAST 5 CHARACTERS ENTERED
   */
61  1  DECLARE INPNTR (4) BYTE;

62  1  GET*ADDRESS:
      PROCEDURE POINTER;
63  2      DECLARE BUFF BYTE;
      /*CLEAR ALL VALUES TO ZERO */
64  2      INPNTR(0) = 0;
65  2      INPNTR(1) = 0;
66  2      INPNTR(2) = 0;
67  2      INPNTR(3) = 0;

68  2      BUFF = 0;
69  2      DO WHILE BUFF <> TRUE;
/* THIS SEQUENCE OF SHIFTS ALLOW THE USER TO TYPE IN FIVE
   OR MORE CHARACTERS TO BECOME THE ACTUAL POINTER FOR 8089
   OR 8086. THIS PROCEDURE RETURNS THE LAST FIVE IN PROPER
   SEQUENCE STORED IN INPNTR(0-3). THE STORAGE
   IS AS FOLLOWS:
   1. THE LAST CHARACTER INPUT GOES INTO
      THE LOW FOUR BITS OF INPNTR(0).
   2. THE NEXT TO LAST CHARACTER GOES INTO
      THE LOW FOUR BITS OF INPNTR(2).

```

PL/M-86 COMPILER 8089 BREAKPOINT ROUTINE

PAGE 4

```

3. THE THIRD CHARACTER INPUT GOES INTO
   THE HIGH FOUR BITS OF INPNTR(2)
4. THE SECOND CHARACTER INPUT GOES INTO
   THE LOW FOUR BITS OF INPNTR(3)
5. THE FIRST CHARACTER INPUT GOES INTO
   THE UPPER FOUR BITS OF INPNTR(3).
THE 86 SHIFTS INPNTR (2,AND3) LEFT FOUR BITS AND ADDS THIS TO
INPNTR(0) RESULTING IN THE ADDRESS THE USER TYPED IN. */
70 3      INPNTR(3) = (SHL(INPNTR(3),4) OR (SHR( INPNTR(2),4) AND OFH));
71 3      INPNTR(2) = (SHL(INPNTR(2),4) OR (INPNTR(0) AND OFH));
72 3      INPNTR(0) = BUFF;
73 3      BUFF = CI;
74 3      BUFF = HEXCONV(BUFF);
75 3      END;
76 2      CALL CO(OAH); /*LINE FEED TO CRT*/
77 2      CALL CO(ODH); /*CARRIAGE RET TO CRT*/
78 2      RETURN PRESENT; /* PRESENT IS A POINTER TO THE ARRAY INPNTR. */
79 2      END;

/* STRINGOUT IS A PROCEDURE TO SEND THE CONSOLE AN ASCII STRING
   ENDING IN THE VALUE 00. STRINGOUT NEEDS A VALUE OF TYPE POINTER
   */

80 1      STRINGOUT:
          PROCEDURE(PTR);
81 2          DECLARE PTR POINTER, STR BASED PTR (1) BYTE;
82 2          I = 0;
83 2          DO WHILE STR(I) <> 0;
84 3              CALL CO(STR(I));
85 3              I = I + 1;
86 3          END;
87 2          END;

88 1      DECLARE TAGIS (*) BYTE DATA (' OPERATING IN ',0),
          TAGISONE (*) BYTE DATA ('IO SPACE',OAH,ODH,0),
          TAGISZERO (*) BYTE DATA ('SYSTEM SPACE',OAH,ODH,0);
/* TAGTEST TESTS THE TAG BIT AND SENDS A MESSAGE TO THE CONSOLE
   THE TAG IS LOCATED IN BIT THREE. A TAG BIT OF ONE MEANS THE
   POINTER IS TO I/D SPACE, AND A TAG BIT OF ZERO MEANS THE
   POINTER IS TO SYSTEM SPACE */
/* THE CALLER MUST DECIDE WHICH BYTE HAS THE TAG AND PASS IT TO TAGTEST */

89 1      TAGTEST:
          PROCEDURE(TEST);
90 2          DECLARE TEST BYTE;
91 2          CALL STRINGOUT(@TAGIS);
92 2          IF (TEST AND 01000B) <> 0
          THEN
93 2              DO;
94 3                  CALL STRINGOUT(@TAGISONE);
95 3              END;
          ELSE
96 2              DO;
97 3                  CALL STRINGOUT(@TAGISZERO);
98 3              END;

```

PL/M-86 COMPILER 8089 BREAKPOINT ROUTINE

PAGE 5

```

99 2      END.
100 1      DECLARE SAVE$ADDR LITERALLY '2000H',
          SAVE$SEG LITERALLY '00COH';

101 1      DECLARE BREAK89 (4) WORD DATA (98B1H,0891H,SAVE$ADDR,SAVE$SEG);
          /* BREAK89 IS AN 4 WORD ESCAPE SEQUENCE TO ADDRESS 2000H
           CONSISTING OF AN LPDI TP,SAVE$ADDR WITH SEGMENT.
           LOCATED AT 00COH. */

          /* BRKRTN IS 33 BYTES OF CODE THAT STORES ALL REGISTERS
           AS FOLLOWS:
           GA STORED AT PP + 239
           GB STORED AT PP + 242
           GC STORED AT PP + 245
           BC STORED AT PP + 248
           IX STORED AT PP + 250
           CC STORED AT PP + 252
           MC STORED AT PP + 254
           */

102 1      DECLARE BRKRTN (33) BYTE AT (02C00H)
          /* 02C00H IS ACTUALLY (SAVE$ADDR + (SHL(SAVE$SEG),4)), AND SHOULD
           MATCH ADDRESS AND SEGMENT WHERE BREAK ROUTINE IS WANTED */
          INITIAL
          (03H,09BH,0EFH,023H,09BH,0F2H,043H,09BH,0F5H,063H,087H,0FBH,0A3H,087H,
          0FAH,0C3H,087H,0FCH,0E3H,087H,0FEH,020H,048H) ;

103 1      DECLARE PP POINTER;
104 1      DECLARE PPP BASED PP (1) BYTE;

105 1      START$PRGM
          PROCEDURE(ONE$TWO,PPP),
106 2      DECLARE ONE$TWO BYTE,PPP POINTER,
          WHERE BASED PPP (1) BYTE;

107 2      WHERE(0) = START$BYTES(0);
108 2      WHERE(1) = 0;
109 2      WHERE(2) = START$BYTES(2);
110 2      WHERE(3) = START$BYTES(3);
111 2      CPDAT((ONE$TWO) * 8) = 3;
          /* IF ONETWO = 1 THEN OUTPUT TO PORT 0FBH, IF ONETWO
           IS 0 THEN OUTPUT TO PORT 0FAH */
112 2      OUTPUT(CHANATTEN + (ONETWO)) = 0,
113 2      CALL STRINGOUT(@CHANGIVEN);
114 2      END;

          /* THIS PART OF THE PROGRAM ALLOWS THE USER TO DEFINE THE
           CP,PP OF EACH CHANNEL */
115 1      DECLARE BREAKOUT BASED ENDPOINTER (1) WORD;

116 1      DECLARE CP POINTER;
117 1      DECLARE CPDAT BASED CP (1) BYTE;

118 1      DECLARE ONEPPDAT BASED ONEPP (1) BYTE;
119 1      DECLARE TWOPPDAT BASED TWOPP (1) BYTE;

120 1      CALL STRINGOUT (@TITLESTRING);

```

```

121 1      CALL STRINGOUT(@GETCP);
122 1      CP = GETADDRESS;
123 1      CALL STRINGOUT(@GETPP);
124 1      CALL STRINGOUT(@ONE);
125 1      ONEPP = GETADDRESS;
126 1      CALL STRINGOUT(@GETPP);
127 1      CALL STRINGOUT(@TWO);
128 1      TWOPP = GETADDRESS;
129 1      OUTPUT (CHANATTEN) = 0; /* INITIALIZATION CA */

130 1      MAIN:
131 1          CALL STRINGOUT(@CHANNUMBER);
132 1          CHAR = CI; /* GET CHANNEL NUMBER */
133 1          IF (CHAR AND 01H) <> 0 /* CHECK BIT ZERO TO DEFINE
134 1              CHANNEL NUMBER */
135 1              THEN DO;
136 1                  CALL STRINGOUT(@ONE);
137 1                  ONETWO = CHANNEL*ONE;
138 1              ELSE
139 1                  DO;
140 1                      CALL STRINGOUT(@TWO);
141 1                      ONETWO = CHANNEL*TWO;
142 1              END;
143 1          CALL STRINGOUT(@GET*START); /* GET STARTING ADDRESS
144 1              FROM USER */
145 1          STARTPOINTER = GETADDRESS;
146 1          DO I = 0 TO 3; /* MOVE STARTING ADDRESS INTO CP AREA */
147 1              STARTBYTES(I) = INPNTR(I);
148 1          END;
149 1          CALL STRINGOUT(@STOPADDR); /* GET STOP ADDRESS
150 1              FROM USER */
151 1          ENDPINTER = GETADDRESS;
152 1          DO I = 0 TO 3; /* MOVE CODE TO SAFE AREA */
153 1              SAVECODE(I) = BREAKOUT(I);
154 1          END;
155 1          DO I = 0 TO 3;
156 1              BREAKOUT(I) = BREAKB9(I); /* MOVE ESCAPE SEQUENCE INTO PLACE */
157 1          END;
158 1          CPDAT(1) = OFFH; /* SET CHANNEL ONE BUSY FLAG */
159 1          CPDAT(9) = OFFH; /* SET CHANNEL TWO BUSY FLAG */
160 1          DO CASE ONETWO;
161 1              PP = ONEPP;
162 1              PP = TWOPP;
163 1          END;
164 1          CALL START*PRGM(ONE*TWO,PP);
165 1          /* WAIT FOR ONE OF THE FOLLOWING
166 1              1. CPDAT(1) = 0 CH1 NOT BUSY
167 1              2. CPDAT(9) = 0 CH2 NOT BUSY
168 1              3. THE 8251 REC. BUFFER IS FULL BECAUSE USER HAS DEPRESSED A KEY
169 1          */
170 1          DO WHILE ( (CPDAT(1) AND CPDAT(9)) AND (NOT (INPUT(CRT*STATUS) AND 02H))) = OFFH;

```

PL/M-86 COMPILER 8089 BREAKPOINT ROUTINE

PAGE 7

```

162 2      END;
163 1      IF (INPUT(CRT$STATUS) AND 02H) <> 0
          THEN
164 1          DO;
165 2              CHAR = CI;
166 2              DO I = 0 TO 3;
167 3                  BREAKOUT(I) = SAVECODE(I);
168 3              END;
          /* IF ONETWO = 0 THEN PUT CHA HLT IN CPDAT(0)
          IF ONETWO = 1 THEN PUT CHA HLT IN CPDAT(8)
          */
169 2      CPDAT(ONETWO * 8) = 06H;
          /* IF ONETWO = 0 THEN OUTPUT TO PORT OFAH, IF ONETWO
          IS 1 THEN OUTPUT TO PORT OFBH.
          */
170 2      OUTPUT(CHANATTEN + ONETWO) = 0;
171 2      DO I = 0 TO 5;
172 3          CALL TIME(100);
173 3      END;

          /* IF BUSY FLAG HAS BEEN CLEARED, THEN A CA HALT&SAVE
          WAS EXECUTED. IF SO, PRINT SAVED TP; IF NOT, ABORT */

174 2      IF CPDAT(SHL(ONETWO,3) + 1) <> 0 /* CHECK BUSY FLAG */
          THEN
175 2          DO;
176 3              CALL STRINGOUT(@ABORT);
177 3          END;
          ELSE
178 2          DO;
179 3              CALL STRINGOUT(@ABORTAT);
180 3              CALL CO(ASCII(SHR(PPP(2),4))); /* UPPER NIBBLE OF ADDR
          STORED BY HALT */

181 3          CALL HEXOUT(PPP(1)); /* MIDDLE BYTE OF ADDR
          STORED BY HALT */

182 3          CALL HEXOUT(PPP(0)); /* LEAST SIG BYTE OF ADDR
          STORED BY HALT */

183 3      END;
184 2      CPDAT(ONETWO * 8) = 3H; /* CA START IN CPDAT(0) OR CPDAT(8) */
185 2      GO TO MAIN;
186 2      END;
187 1      DO;

188 2      CALL STRINGOUT(@BKREACHED);

189 2      CALL STRINGOUT(@GASTRING);
190 2      CALL CO(ASCII(SHR(PPP(241),4)));
191 2      CALL HEXOUT(PPP(240));
192 2      CALL HEXOUT(PPP(239));
193 2      CALL TAGTEST(PPP(241));

194 2      CALL STRINGOUT(@GBSTRING);
195 2      CALL CO(ASCII(SHR(PPP(244),4)));
196 2      CALL HEXOUT(PPP(243));

```

PL/M-86 COMPILER 8089 BREAKPOINT ROUTINE

PAGE 8

```

197 2      CALL HEXOUT(PPP(242));
198 2      CALL TAGTEST(PPP(244));

199 2      CALL STRINGOUT(@QCSTRING);
200 2      CALL CD(ASCII(SHR(PPP(247),4)));
201 2      CALL HEXOUT(PPP(246));
202 2      CALL HEXOUT(PPP(245));
203 2      CALL TAGTEST(PPP(247));

204 2      CALL STRINGOUT(@BCSTRING);
205 2      CALL HEXOUT(PPP(249));
206 2      CALL HEXOUT(PPP(248));

207 2      CALL STRINGOUT(@IXSTRING);
208 2      CALL HEXOUT(PPP(251));
209 2      CALL HEXOUT(PPP(250));

210 2      CALL STRINGOUT(@CCSTRING);
211 2      CALL HEXOUT(PPP(253));
212 2      CALL HEXOUT(PPP(252));

213 2      CALL STRINGOUT(@MCSTRING);
214 2      CALL HEXOUT(PPP(255));
215 2      CALL HEXOUT(PPP(254));

216 2      END;
          /* RESTORE CODE TO ORIGINAL LOCATION */
217 1      DO I = 0 TO 3;
218 2          BREAKOUT(I) = SAVECODE(I);
219 2      END;

220 1      GO TO MAIN;

221 1      END.

```

MODULE INFORMATION

```

CODE AREA SIZE.      = 0619H   1561D
CONSTANT AREA SIZE = 01EFH   495D
VARIABLE AREA SIZE = 0020H   32D
MAXIMUM STACK SIZE = 0014H   20D
427 LINES READ
0 PROGRAM ERROR(S)

```

END OF PL/M-86 COMPILATION

8089 ASSEMBLER

ISIS-II 8089 ASSEMBLER X004 ASSEMBLY OF MODULE AP50_BREAKPOINT_ROUTINE
 OBJECT MODULE PLACED IN :FO:BRKASM.OBJ
 ASSEMBLER INVOKED BY ASMB9.4 BRKASM.SRC

```

0000          1 NAME      AP50_BREAKPOINT_ROUTINE
                2 BRKPNT  SEGMENT
                3 ;*****
                4 ; BASIC 8089 BREAKPOINT ROUTINE
                5 ;   BY JOHN ATWOOD  REV 3  B/13/79
                6 ;   INTEL CORPORATION
                7 ;*****
                8
                9 ; THE FOLLOWING CODE IS CONTAINED IN THE PL/M-86
                10 ; CONTROL PROGRAM(BREAK.89) AND IS ASSEMBLED HERE
                11 ; TO ILLUSTRATE HOW THE ESCAPE SEQUENCE AND SAVE
                12 ; ROUTINE CODE WAS GENERATED. TO USE THE 8089 BREAK-
                13 ; POINT PROGRAM, THIS ASMB9 PROGRAM WOULD NOT BE
                14 ; NEEDED. SAVE_ADDR IS THE SAME AS SAVE$ADDR IN THE
                15 ; BREAK.89 PROGRAM.
                16
                17 SAVE_ADDR      EQU      2000H      ;SAVE ROUTINE ADDRESS
                18
                19          LPDI  TP,SAVE_ADDR      ;JUMP TO SAVE ROUTINE
                20
                21 ;*****
                22
                23 ; REGISTER SAVE LOCATIONS WITHIN PB:
                24
                25 REGS      STRUC
                26 PBLOCK:  DS      239      ;PARAMETER BLOCK
                27 GASAV:   DS      3        ;GA AREA
                28 GBSAV:   DS      3        ;GB AREA
                29 GCSAV:   DS      3        ;GC AREA
                30 BCSAV:   DS      2        ;BC AREA
                31 IXSAV:   DS      2        ;IX AREA
                32 CCSAV:   DS      2        ;CC AREA
                33 MCSAV:   DS      2        ;MC AREA
                34 REGS      ENDS
                35
                36 ; REGISTER SAVE ROUTINE:
                37
                38 ORG      SAVE_ADDR
                39
                40          MOVP [PP],GASAV,GA      ;SAVE GA
                41          MOVP [PP],GBSAV,GB      ;SAVE GB
                42          MOVP [PP],GCSAV,GC      ;SAVE GC
                43          MOV  [PP],BCSAV,BC      ;SAVE BC
                44          MOV  [PP],IXSAV,IX      ;SAVE IX
                45          MOV  [PP],CCSAV,CC      ;SAVE CC
                46          MOV  [PP],MCSAV,MC      ;SAVE MC
                47
                48          HLT                      ;STOP THIS CHANNEL,
                49                      ;CLEAR BUSY FLAG.
                50 ;*****
                51 BRKPNT  ENDS
                52
                53 END

```

March 1979

**Designing 8086, 8088, 8089
Multiprocessing Systems
with the 8289 Bus Arbiter**

Designing 8086, 8088, 8089 Multiprocessor Systems with the 8289 Bus Arbiter

Contents

INTRODUCTION

BUS ARBITER OPERATING CHARACTERISTICS

**MULTI-MASTER SYSTEM BUS SURRENDER
AND REQUEST**

**8289 BUS ARBITER INTERFACING TO THE
8288 BUS CONTROLLER**

8289 BUS ARBITER INTERNAL ARCHITECTURE

**8086 FAMILY PROCESSOR TYPES AND
SYSTEM CONFIGURATIONS**

8289 SINGLE BUS INTERFACE

IOB INTERFACE

RESB INTERFACE

INTERFACE TO TWO MULTI-MASTER BUSES

WHEN TO USE THE DIFFERENT MODES

Single Bus Multi-master Interface

IOB Mode

Resident Bus Mode

CONCLUSION

Our thanks to Jim Nadir, the author of this application note. Jim is a design engineer in the microprocessors and peripherals operation division. Please direct any technical questions you may have to your local Intel FAE (Field Application Engineer).

INTRODUCTION

Over the past several years, microprocessors have been increasing in popularity. The performance improvements and cost reductions afforded by LSI technology have spurred on the design motivation of using multiple processors to meet system real-time performance requirements. The desire for improved system real-time response, system reliability and modularity has made multiprocessing techniques an increasingly attractive alternative to the system design engineer; techniques that are characterized as having more than one microprocessor share common resources, such as memory and I/O, over a common multiple processor bus.

This type of design concept allows the system designer to partition overall system functions into tasks that each of several processors can handle individually to increase system performance and throughput. But, how should a designer proceed to implement a multiprocessing system? Should he design his own? If so, how are the microprocessors synchronized to avoid contention problems? The designer could put them all in phase using one clock for all the microprocessors. This may work, until the physical dimensions of the system become large. When this occurs, the designer is faced with many problems, like clock skew (resulting in bus spec violations) and duty cycle variations.

A better approach to implementing a multiprocessor system is not to have a common processor clock, but allow each processor to work asynchronously with respect to each other. The microprocessor requests to use the multiple processor bus could then be synchronized to a high frequency external clock which will permit duty cycle and phase shift variations. This type of approach has the benefit of allowing modularity of hardware. When new system functions are desired, more processing power can be added without impacting existing processor task partitioning.

One approach to implement this asynchronous processing structure would be to have all the bus requests enter a priority encoder which samples its inputs as a function of the higher frequency "bus clock". The inputs would arrive asynchronously to the priority encoder and would be resolved by the priority encoder structure as to which microprocessor would be granted the bus. Another approach, that used by Intel, is rather than allowing the requests to arrive asynchronously with respect to one another at the priority encoder, the bus requests are synchronized first to an external high frequency bus clock and then sent to the priority encoder to be resolved. In this way, the resolving circuitry common to all microprocessors is kept at a minimum. Overall system reliability is improved in the sense that should a circuit which serves to synchronize the processor's request (which is now located on the same card as the microprocessor itself) fail, it is only necessary to remove that card from the system and the rest of the system will continue to function. Whereas in the other approach, should the synchronizing mechanism fail, the whole

system goes down, as the synchronizing mechanism is located at the shared resource. In addition to the improved system reliability, moving the synchronization mechanism to the processor permits processor control over that mechanism, thereby permitting system flexibility (as will be shown) which could not be reasonably obtained by any other approach.

This synchronizing or arbitrating function was integrated into the 8289, a custom arbitration unit for the 8086, 8088, and 8089 processors. This note basically describes the 8289 arbitration unit, illustrates its different modes of operation and hardware connect in a multiprocessor system. Related and useful documents are: 8086 user's manual, 8289 data sheet, Article Reprint -55: Design Motivations for Multiple Processor Microcomputer Systems (which discusses implementing a semaphore with the MULTIBUS™) and Application Note 28A, Intel MULTIBUS™ interfacing.

BUS ARBITER OPERATING CHARACTERISTICS

The 8289 Bus Arbiter operates in conjunction with the 8288 Bus Controller to interface an 8086, 8088, or 8089 processor to a multi-master system bus (the 8289 is used as a general bus arbitration unit). The processor is unaware of the arbiter's existence and issues commands as though it has exclusive use of the system bus. If the processor does not have the use of the multi-master system bus, the bus arbiter prevents the bus controller, the data transceivers and the address latches from accessing the system bus (i.e., all bus driver outputs are forced into the high impedance state). Since the command was not issued, a transfer acknowledge (XACK) will not be returned and the processor will enter into wait states. Transfer acknowledges are signals returned from the addressed resource to indicate to the processor that the transfer is complete. This signal is typically used to control the ready inputs of the clock generator. The processor will remain in wait until the bus arbiter acquires the use of the multi-master system bus, whereupon the bus arbiter will allow the bus controller, the data transceivers and the address latches to access the system bus. Once the command has been issued and a data transfer has taken place, a transfer acknowledge (XACK) is returned to the processor. The processor then completes its transfer cycle. Thus, the arbiter serves to multiplex a processor (or bus master) onto a multi-master system bus and avoid contention problems between bus masters.

Since there can be many bus masters on a multi-master system bus, some means of resolving priority between bus masters simultaneously requesting the bus must be provided. The 8289 Bus Arbiter provides for several resolving techniques. All the techniques are based on a priority concept that at a given time one bus master will have priority above all the rest. These techniques include the parallel priority resolving techniques, serial priority resolving and rotating priority techniques.

A parallel priority resolving technique has a separate bus request ($\overline{\text{BREQ}}$) line for each arbiter on the multi-master bus (see Figure 1). Each $\overline{\text{BREQ}}$ line enters into a priority encoder which generates the binary address of the highest priority $\overline{\text{BREQ}}$ line which is active at the inputs. The output binary address is decoded by a decoder to select the corresponding $\overline{\text{BPRN}}$ (bus priority in) line to be returned to the highest priority requesting arbiter. The arbiter receiving priority ($\overline{\text{BPRN}}$ active low) then allows its associated bus master onto the multi-master system bus as soon as it becomes available (i.e., it is no longer busy). When one bus arbiter gains priority over another arbiter, it cannot immediately seize the bus, it must wait until the present bus occupant com-

pletes its transfer cycle. Upon completing its transfer cycle, the present bus occupant recognizes that it no longer has priority and surrenders the bus, releasing $\overline{\text{BUSY}}$. $\overline{\text{BUSY}}$ is an active low OR-tied signal line which goes to every bus arbiter on the system bus. When $\overline{\text{BUSY}}$ goes high, the arbiter which presently has bus priority ($\overline{\text{BPRN}}$ active low) then seizes the bus and pulls $\overline{\text{BUSY}}$ low to keep other arbiters off the bus. (See waveform timing diagram, Figure 2.) Note that all multi-master system bus transactions are synchronized to the bus clock ($\overline{\text{BCLK}}$). This allows for the parallel priority resolving circuitry or, any other priority resolving scheme employed, time to settle and make a correct decision.

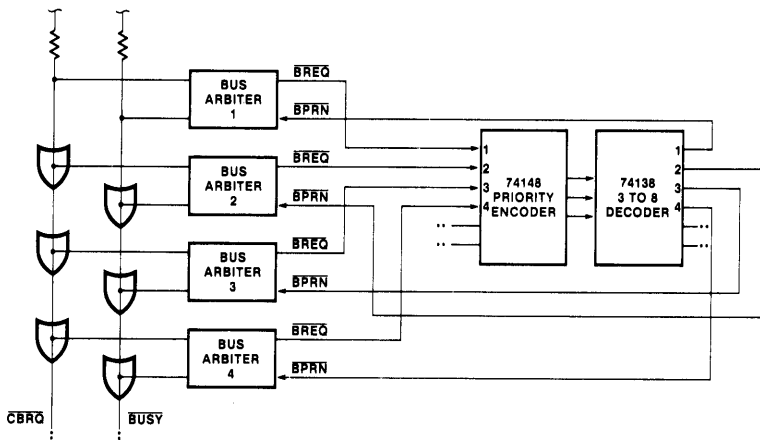


Figure 1. Parallel Priority Resolving Technique

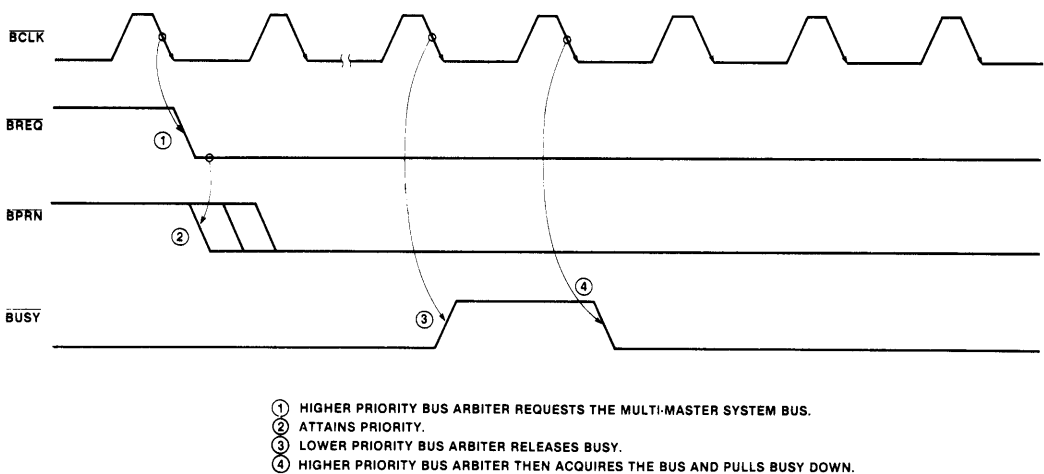
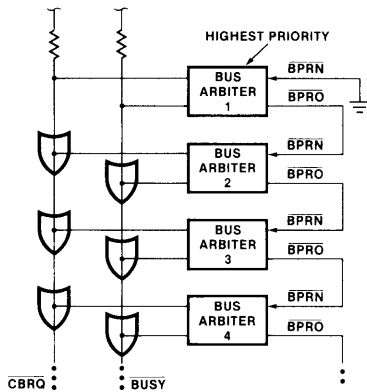


Figure 2. Higher Priority Arbiter Obtaining The Bus From A Lower Priority Arbiter

A serial priority resolving technique eliminates the need for the priority encoder-decoder arrangement by daisy-chaining the bus arbiters together. This is accomplished by connecting the higher priority bus arbiter's $\overline{\text{BPRO}}$ (bus priority out) output to the $\overline{\text{BPRN}}$ of the next lower priority (see Figure 3). The highest priority bus arbiter would have its $\overline{\text{BPRN}}$ line grounded, signifying to the arbiter that it always has highest priority when requesting the bus.



THE NUMBER OF ARBITERS THAT MAY BE DAISY-CHAINED TOGETHER IN THE SERIAL PRIORITY RESOLVING TECHNIQUE IS A FUNCTION OF $\overline{\text{BCLK}}$ AND THE PROPAGATION DELAY FROM ARBITER TO ARBITER. NORMALLY, AT 10 MHz ONLY 3 ARBITERS MAY BE DAISY-CHAINED. SEE TEXT.

Figure 3. Serial Priority Resolving

A rotating priority resolving technique arrangement is similar to that of the parallel priority resolving technique except that priority is dynamically reassigned. The priority encoder is replaced by a more complex circuit which rotates priority between requesting arbiters, thus guaranteeing each arbiter equal time on the multi-master system bus.

There are advantages and disadvantages for each of the techniques described above. The rotating priority resolving technique requires an extensive amount of logic to implement, while the serial technique can accommodate only a limited number of bus arbiters before the daisy-chain propagation delay exceeds the multi-master system bus clock ($\overline{\text{BCLK}}$). The parallel priority resolving technique is, in general, the best compromise. It allows for many arbiters to be present on the bus while not requiring much logic to implement.

Whatever resolving technique is chosen, it is the highest priority bus arbiter requesting use of the multi-master system bus which obtains the bus. Exceptions do exist with the 8289 Bus Arbiter where a lower priority arbiter may take away the bus from a higher priority arbiter without the need for any additional external logic. This is accomplished through the use of the $\overline{\text{CBRQ}}$ pin, discussed in a later section.

MULTI-MASTER SYSTEM BUS SURRENDER AND REQUEST

The 8289 Bus Arbiter provides an intelligent interface to allow a processor or bus master of the 8086 family to access a multi-master system bus. The arbiter directs the processor onto the bus and allows both higher and lower priority bus masters to acquire the bus. Higher priority masters obtain the bus when the present bus master utilizing the bus completes its transfer cycle (including hold time). Lower priority bus masters obtain the bus when a higher priority bus master is not accessing the system bus and a lower priority arbiter has pulled $\overline{\text{CBRQ}}$ low. This signifies to the arbiter presently holding the multi-processor bus that a lower priority arbiter would like to acquire the bus when it is not being used. A strapping option (ANYRQST) allows the multi-master system bus to be surrendered to any bus master requesting the bus, regardless of its priority. If there are no other bus masters requesting the bus, the arbiter maintains the bus as long as its associated bus master has not entered the HALT state. *The 8289 Bus Arbiter will not voluntarily surrender the system bus and has to be forced off by another bus master.* An exception to this can be obtained by strapping $\overline{\text{CBRQ}}$ low and ANYRQST high. In this configuration the 8289 will release the bus after each transfer cycle.

How the 8289 Bus Arbiter is configured determines the manner in which the arbiter requests and surrenders the system bus. If the arbiter is configured to operate with a processor which has access to both a multi-master system bus and a resident bus, the arbiter requests the use of the multi-master system bus only for system bus accesses (i.e., it is a function of the SYSB/RESB input pin). While the processor is accessing the resident bus, the arbiter permits a lower priority bus master to seize the system bus via $\overline{\text{CBRQ}}$, since it is not being used. A processor configuration with both an I/O peripheral bus and a system bus behaves similarly. If the processor is accessing the peripheral bus, the arbiter permits the surrendering of the multi-master system bus to a lower priority bus master. To request the use of the multi-master system bus, the processor must perform a system memory access (as opposed to an I/O access).

The arbiter decodes the processor status lines to determine what type of access is being performed and behaves correspondingly. For simpler system configurations, such as a processor which accesses only a multi-master system bus, the arbiter requests the use of the system bus when it detects the status lines initiating a transfer cycle. The decoding of these status lines can be referenced in the 8086, 8088 (non-I/O processor) data sheets or the 8089 (I/O processor) data sheet.

There is one condition common to all system configurations where the multi-master system bus is surrendered to a lower priority bus master requesting the bus by pulling $\overline{\text{CBRQ}}$ low. This is the idle or inactive state (TI) which is unique to the 8086 and 8088 processor family. This TI state comes about due to the processor's ability to fetch instructions in advance and store them internally for quick access. The size of the internal queue was optimized so that the processor would make the most eff-

fective use of its resources and be slightly execution bound. Since the processor can fetch code faster than it can execute it, it will fill to capacity its internal storage queue. When this occurs, the processor will enter into idle or inactive states (TI) until the processor has executed some of the code in the storage queue. Once this occurs, the processor will exit the TI state and again start code fetching. Between entering into and exiting from the TI state an indeterminate number of TI states can occur during which the bus arbiter permits the surrendering of the multi-master system bus to a lower priority bus master. As noted earlier and worth repeating here, once the 8289 Bus Arbiter acquires the use of the multi-master system it will not voluntarily surrender the bus and has to be forced off by another bus master. This will be discussed in more detail later.

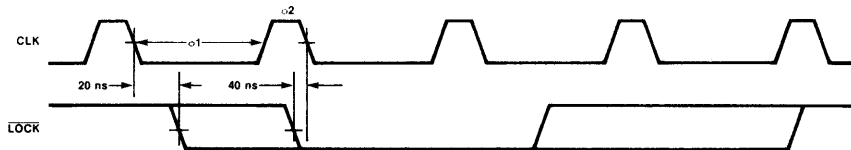
Two other signals, $\overline{\text{LOCK}}$ and $\overline{\text{CRQLCK}}$ (Figure 4), lend to the flexibility of the 8289 Bus Arbiter within system configurations. $\overline{\text{LOCK}}$ is a signal generated by the processor to prevent the bus arbiter from surrendering the multi-master system bus to any other bus master, either higher or lower priority. $\overline{\text{CRQLCK}}$ (common request lock) serves to prevent the bus arbiter from surrendering the bus to a lower priority bus master when conditions warrant it. $\overline{\text{LOCK}}$ is used for implementing software semaphores for critical code sections and real time

critical events (such as refreshing or hard disk transfers).

8289 BUS ARBITER INTERFACING TO THE 8288 BUS CONTROLLER

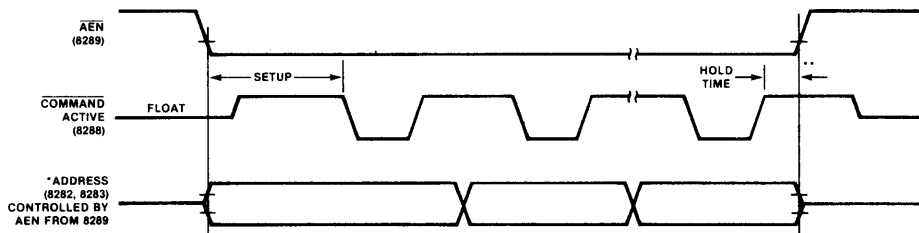
Once the 8289 Bus Arbiter determines to either allow its associated processor onto the multi-master system bus or to surrender the bus, it must guarantee that command setup and hold times are not violated. This is a two part problem. One, guaranteeing hold time and two, guaranteeing setup time. The 8288 Bus Controller performs the actual task of establishing setup time, while the 8289 Bus Arbiter establishes hold time (see Figure 5).

The 8289 Bus Arbiter communicates with the 8288 Bus Controller via the $\overline{\text{AEN}}$ line. When the arbiter allows its associated processor access to the multi-master system bus, it activates $\overline{\text{AEN}}$. $\overline{\text{AEN}}$ immediately enables the address latches and data transceivers. The bus controller responds to $\overline{\text{AEN}}$ by bringing its command output buffers out of high impedance state but keeping all commands disqualified until command setup time is established. Once established, the appropriate command is then issued. $\overline{\text{AEN}}$ is brought to the false state after the command hold time has been established by the arbiter when surrendering the bus.



LOCK TIMING
 THE ONLY CRITICAL LOCK TIMING IS THAT SHOWN ABOVE. $\overline{\text{LOCK}}$ MUST BE ACTIVATED NO SOONER THAN 20 ns INTO ϕ_1 AND NO LATER THAN 40 ns PRIOR TO THE END OF ϕ_2 . $\overline{\text{LOCK}}$ INACTIVE HAS NO CRITICAL TIMING AND CAN BE ASYNCHRONOUS.
 $\overline{\text{CRQLCK}}$ HAS NO CRITICAL TIMING AND IS CONSIDERED AS AN ASYNCHRONOUS INPUT SIGNAL.

Figure 4. Lock Timing



*ADDRESSES ARE ACTIVATED IMMEDIATELY WHILE COMMAND IS DELAY TO ESTABLISH SETUP TIME REQUIREMENTS.
 **THE 8289 ARBITER INTERNALLY TRACKS THE PROCESSOR CYCLE TO ESTABLISH THE PROPER AMOUNT OF HOLD TIME AFTER THE COMMAND HAS GONE INACTIVE.

Figure 5. Single Bus Interface Timing

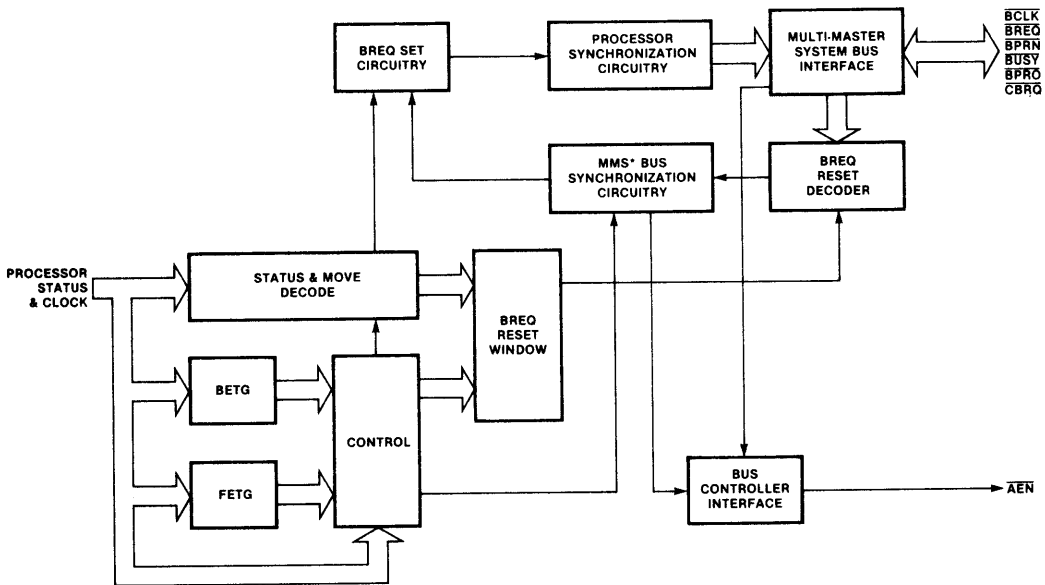
8289 BUS ARBITER INTERNAL ARCHITECTURE

A block diagram of the internal architecture of the 8289 Bus Arbiter is shown in Figure 6. It is useful to understand this block diagram when discussing the different modes of the 8289 and their impact on processor bus operations; however, you may want to skip this section to "8086 family processor types and system configurations" and return to it afterwards, as this section addresses the very involved reader. The front end state generator (FETG) and the back end state generator (BETG) allow the arbiter to track the processor cycle. An examination of an 8086 family processor state timings show that all command and control signals are issued in states T1 and T2 while being terminated in states T3 and T4, with an indeterminate number of wait states (Tw) occurring in between. Note further, that an indeterminate number of idle or inactive states can occur immediately proceeding and following a given transfer cycle. Since an indeterminate number of wait states can occur, two state generators are required; one to generate control signals (the FETG) and one to terminate control signals (the BETG). The FETG is triggered into operation when the processor activates the status lines. The FETG is reset and the BETG is triggered into operation by the status lines going to the passive condition. The BETG is reset when the status lines again go active.

It is necessary for the 8289 Bus Arbiter to track the processor in order that it is properly able to determine where and when to request or surrender the use of the multi-master system bus. In system configurations which access a resident bus, the use of the multi-master system

bus is requested later in order to allow time for the SYSB/ $\overline{\text{RESB}}$ input to become valid. For systems which access a peripheral bus, the arbiter issues a request for the system bus only for memory transfer cycles which it decodes from the status lines (and time must be allowed for the status lines to become valid and then decoded). In a system which accesses only a multi-master system bus, a request is made as soon as the arbiter detects an active-going transition on the processor's status lines. Thus, when the processor initiates a transfer cycle, the FETG is triggered into operation and, depending upon what mode the arbiter is configured in, the STATUS & MODE DECODE circuitry initiates a request for the system bus at the appropriate time. The request enters the BREQ SET circuitry where it is then synchronized to the multi-master system bus clock (BCLK) by the PROCESSOR SYNCHRONIZATION circuitry.* Once synchronized, the multi-master system bus interface circuitry issues a BREQ. When the priority resolving circuitry returns a BPRN (bus priority in), the PROCESSOR SYNCHRONIZATION circuitry seizes the bus the next time it becomes available (i.e., BUSY goes high) by pulling BUSY low one BCLK after it goes high and enables AEN. (See waveform timing diagram in Figure 2). Once the arbiter acquires the use of the system bus and a data exchange has taken place (a transfer acknowledge, XACK, was returned to the processor), the processor status lines go passive and the

*Due to the asynchronous nature of processor transfer request to the multi-master system bus clock, it is necessary to synchronize the processor's transfer request to BCLK.



*MMS = MULTI-MASTER SYSTEM

Figure 6. 8289 Bus Arbiter Block Diagram

BETG is triggered into operation. The BETG provides the timing for the bus surrender circuitries in the event that conditions warrant the surrender of the multi-master bus, i.e., the bus arbiter lost priority to a higher bus master or the processor has entered into TI states and \overline{CBRQ} is pulled low, etc. If such is the case, the BREQ RESET DECODER initiates a bus surrender request. The bus surrender request is synchronized by the MMS BUS SYNCHRONIZATION CIRCUITRY to the processor clock. The MMS BUS SYNCHRONIZATION CIRCUITRY instructs the bus controller interface circuitry to make \overline{AEN} go false and resets the BREQ SET circuitry. Resetting the BREQ SET circuitry will cause its output to go false and be synchronized by the processor synchronization, eventually instructing the MULTIMASTER SYSTEM BUS INTERFACE circuitry to reset BREQ. In the event that a lower priority arbiter has caused the arbiter to surrender the bus, it is necessary that BREQ be reset. Resetting BREQ allows the priority resolving circuitry to generate \overline{BPRN} to the next highest priority bus master requesting the bus. The BREQ RESET WINDOW circuitry provides a 'window' wherein the arbiter allows the multi-master system bus to be surrendered and serves as part of the MMS bus-processor synchronization circuitry.

8086 FAMILY PROCESSOR TYPES AND SYSTEM CONFIGURATIONS

There are two types of processors in the 8086 family — an I/O processor (the 8089 IOP) and a non-I/O processor (the 8086 and 8088 CPUs). Consequently, there are two basic operating modes in the 8289 Bus Arbiter. One, the IOB (I/O peripheral bus) mode, permits the processor access to both an I/O peripheral bus and a multi-master system bus. The second, the RESB (resident bus) mode, permits the processor to communicate over both a resident bus and a multi-master system bus. Even though it is intended for the arbiter to be configured in the IOB mode when interfacing to an I/O processor and for it to be in the RESB mode when interfacing to a non-I/O processor, it is quite possible for the reverse to be true. That is, it is possible for a non-I/O processor to have access to an I/O peripheral bus or for an I/O processor to have access to a resident bus as well as access to a multi-master system bus. The IOB strapping option configures the 8289 Bus Arbiter into the IOB mode and RESB strapping option configures it into the resident bus mode. If both strapping options are strapped false, a third mode of operation is created, the single bus mode, in which the arbiter interfaces the processor to a multi-master system bus only. With both options strapped true, the arbiter interfaces the processor to a multi-master system bus, a resident bus and an I/O bus.

To better understand the 8289 Bus Arbiter, each of the operating modes, along with their respective timings, are examined by means of examples. The simplest configuration, the Single Bus Configuration, (both IOB and RESB strapped inactive) will be considered first, fol-

lowed by the I/O bus Configuration and the Resident Bus Configuration. Finally, brief mention is made of a configuration that allows the processor to interface to two multi-master system buses. This particular configuration is briefly mentioned because, as will be seen, it is simply an extension of the resident bus configuration. When discussing the Single Bus Configuration, processor/arbiter, arbiter/system bus and internal arbiter, considerations are made resulting in a table that illustrates overhead in requesting the system bus. As this applies to the other 8289 configurations, only additional considerations will be given. A summary of when to use the different configurations is given at the end.

8289 SINGLE BUS INTERFACE

Figure 7 shows a block diagram of a bus master which has to interface only to a system bus — preferably the MULTIBUS — where there exists more than one bus master. In later configurations, it will be shown how the processor can be made to interface with more than one bus. Since the processor has only to interface with one bus, this configuration is called "single".

Connecting the 8289 Bus Arbiter to the processor is as simple as it was to connect the 8288 Bus Controller. Namely, the three status lines, $\overline{S0}$, $\overline{S1}$, and $\overline{S2}$ are directly connected from the processor to the arbiter. The clock line from the 8284 Clock Generator is brought down and connected. (Note that both the 8288 Bus Controller and the 8289 Bus Arbiter are connected to the same clock, CLK and not the peripheral clock, PCLK as the 8086 processor.) From the arbiter, \overline{AEN} is connected to the bus controller and to the clock generator. The IOB pin on the arbiter is strapped high and on the controller the IOB pin is strapped low. In addition, the RESB pin on the arbiter is strapped low, finishing the processor interface.

Some flexibility exists with the MULTIBUS or multi-master system bus interface. The system designer must first decide upon the type of priority resolving scheme to be employed, whether it is to be the serial, parallel, or rotating priority scheme. A rotating priority scheme would be employed where the system designer would want to guarantee that every bus master on the bus would be given time on the bus. In the serial and parallel schemes, the possibility exists that the lowest assigned priority bus master may not acquire the bus for long periods of time. This occurs because priority is permanently assigned and if bus demand is high by the higher assigned priorities, then the lower priorities must wait. In most cases, this situation is acceptable because the highest priority is assigned to the bus master that cannot wait. Highest priority is usually assigned to DMA type devices where service requirements occur in real time. CPUs are assigned the lower priorities. For the purpose of this discussion, the parallel priority scheme will be used with brief reference to the serial priority scheme.