

Description

Background of the invention

The invention relates generally to computer systems. In particular, the invention relates to the sharing of an interrupt level by multiple interrupt sources.

A popular type of computer architecture allows multiple users to asynchronously request the service of a resource that is to be used exclusively by the source. One example of the resource is the processor and I/O channel of the computer system. Several I/O devices are attached to the I/O channel and have sufficient intelligence to execute tasks independently of the processor. However, at random times the I/O devices require communication with the processor. In this case, a request is sent to the processor requesting service by the processor. When the processor interrupts its own processing and honors the request, the details of the service request and possibly a response are conveyed on the I/O channel. It is, however, to be understood that the architecture associated with service requests and interrupt is more general than the situation just described.

One method of servicing these asynchronous requests for service is to have the processor periodically poll all the devices attached to the I/O channel to determine if the device requires service. This method though has several drawbacks. There is a high overhead associated with the polling because the poll must be periodically performed even if there are no outstanding requests. Furthermore, in order to reduce the overhead to a reasonable level, the repetition rate for the polling is made relatively long. As a result, though, the response to a service request becomes slow. In the usual situation, the I/O device making the service request cannot continue its own processing until the service request is serviced by the processor. As a result, system utilization falls with increased time between polls.

An alternative approach to polling is the use of interrupts or interrupt requests. A dedicated interrupt line is connected between the I/O device and the processor. Whenever the I/O device requires service, that I/O device outputs an interrupt signal onto the interrupt line. A separate interrupt terminal is provided on the processor and the processor, upon detecting an active interrupt line, goes into an interrupt servicing mode to service the request indicated by the interrupt signal. It should be emphasized that the interrupt signal is asynchronous with the operation of the processor and, indeed, is not necessarily synchronized with the processor clock if separate clocks are provided for the processor and the I/O device.

The computer architecture described above is satisfactory if there is only a single I/O device that is making interrupt request. However, it is more typical that there are several such I/O devices, each asynchronously requesting service by an interrupt signal. Indeed, the trend in modern

computer architecture is for increasing intelligence contained in a large number of I/O devices. The communication to the processor, however, is performed over the I/O channel and requires the exclusive utilization of the processor and the I/O channel for the communication between the I/O device and the processor.

One conventional method of providing for multiple interrupt sources is to use a programmable interrupt controller. A popular programmable interrupt controller is the 8259 manufactured by the Intel Corporation and described in its publication "The 8086 Family Users Manual" dated October 1979 and available from the Intel Corporation in Santa Clara, California. The INTEL controller 8259 allows for eight interrupt levels. Each interrupt level has its own interrupt request line and the eight levels are prioritized among themselves. Whenever an interrupt request line indicates a request for service and the request does not conflict with a higher priority request, the INTEL 8259 outputs an interrupt signal on a single line to the processor that it supports. Along with the interrupt signal to the processor, the INTEL 8259 sends to the processor an interrupt vector associated with that interrupt level. The processor uses the interrupt vector to properly service the interrupt request.

For multiple interrupt sources, there is always the possibility that two interrupts for a service request are outstanding at a particular time. the INTEL 8259 prioritizes the outstanding requests and, possibly, the request currently being serviced. The highest priority request is serviced first, via the interrupt signal and interrupt vector to the processor, while the outstanding requests are held by the INTEL 8259. Thus, it is seen that the eight interrupt request lines to the 8259 function as separate interrupt levels. Apart from the prioritizing in the case of conflict, the eight different input levels are separable and each is operated similarly to the single interrupt level described above.

An important aspect of the use of the INTEL 8259 is the convention that an interrupt request signal to the 8259 transitions from a low to a high level to determine that an interrupt request has been received by the INTEL 8259. The convention further states that the interrupt request signal remains high until the request indicated by the interrupt has been serviced. It is a finite time T_{IO} after the above described transition that the interrupt request line having the upward transition is measured to determine that it is still high. This measurement indicates which of the interrupt levels are currently asserted. It is to be noted that this high level measurement is redundant.

A single INTEL 8259 is thus designed to interface the interrupt request line for up to eight I/O devices. However, for modern computer systems this number turns out to be inadequate. For example, the Personal Computer (PC) of the IBM Corporation uses an INTEL 8259 as an interface for the interrupt request lines from the various

additional boards plugged into the system or planar board housing the processor. Two of the eight interrupt request lines are, however, used for interrupts originating from the system board. The remaining six interrupt levels are lead to a bus connected to all the slots for additional I/O devices. Thus, there are an insufficient number of interrupt levels for a PC whose I/O slots are completely filled. In fact, some boards require more than one interrupt level. The problem is even worse because a particular board needs to be identified with a particular interrupt level. There are a large number of I/O boards, well in excess of six varieties, and the selection of the particular boards for a PC depends upon the user's needs. If the I/O boards are hard-wired for a particular interrupt level, the possibility arises that the selection of I/O boards is such that a large number of them are using the same interrupt level.

It would be possible, of course, to expand the number of interrupt levels, even with the continued use of the INTEL 8259. For instance, two or more INTEL 8259s can be used, as described in the previously cited Intel publication and also by Khera in U.S. Patent 4,275,458. The difficulty with this approach is that it requires the number of interrupt lines that is equal to the number of interrupt levels. It is desirable to provide for additional interrupt sources without modifying the bus connecting the I/O slots to the system board, that is, to continue with the use of only six interrupt levels.

It is possible to share an interrupt level, that is, that more than one interrupt source is somehow using the same interrupt level. In one version of the PC, the computer system can turn off the interrupt function on one or more boards. Thereby two or more boards may be sharing an interrupt level but only one of them has an operable interrupt at any one time. This approach has the obvious disadvantage that some boards cannot always use their interrupt functions. Furthermore, its use depends upon the predictability of interrupts from particular boards. Such predictability is not assured and runs counter to the architecture of an interrupt driven system.

Another approach for interrupt level sharing is to provide the possibility of multiple boards outputting active interrupt level signals to a single interrupt request line and then to further provide software routines to determine which of the boards connected to that interrupt level is actually making the request. It is well known that TTL open-collector outputs can be used to share a single electrical line. Unfortunately, open collector outputs can assert a line only if the high level is non-asserted or inactive and the low level is asserted or active. Then, any low open-collector output will pull its connected line low regardless of the high output of any other open-collector output connected to that line. Such a convention is, unfortunately, inconsistent with the requirements of the INTEL 8259 which requires that the low state is inactive and the high state is active. A

solution has been proposed by others to accommodate line sharing on the interrupt request inputs to the INTEL 8259. This circuit is illustrated in Figure 1 in which the signal input to a tri-state buffer 10 is grounded. An enable input to the tri-state buffer 10 is connected to an internal interrupt signal 12. The output of the tri-state buffer 10 is connected to an external interrupt line 16 that has a pull-up resistor 14 to a voltage source equal to a high level 19 for the logic circuit, for example, +5V for TTL logic. When the tri-state buffer is not enabled, the pull-up resistor 14 pulls the output up to a high level, assuming that no other buffer connected to the external interrupt line 16 is pulling that line 16 low. However, when the tri-state buffer 10 is enabled by the internal interrupt signal 20, the output line 16 is grounded or put into the low state. The result can also be accomplished with an open-collector output of a TTL circuit.

Although it may appear that the negative pulse transition 17 on the output line 16 has the wrong polarity for the INTEL 8259, in fact the INTEL 8259 responds to the positive transition 18 of the output signal and the output signal remains at a high level 19 following the positive transition to allow the identification of the interrupt level. Once the INTEL 8259 has detected the positive transition 18, it waits a time T_{ID} before measuring the high level on the external interrupt line 16 to identify the interrupt level making the service request to the INTEL 8259. When that interrupt level is finally serviced, polling is performed by the system to determine which of the interrupt sources of that interrupt level has caused the interrupt.

There are several disadvantages to this approach. One of the most serious disadvantages is the possibility of a phantom interrupt. A phantom interrupt occurs when two interrupt sources of a particular level make closely timed interrupt requests as illustrated in Figure 2. If the first interrupt source produces a negative pulse 20 on the external interrupt line 16 and then, in a time less than T_{ID} following the upward transition 18, a second interrupt source produces a second negative pulse 22, when the INTEL 8259 measures the input levels on its interrupt request inputs, it detects a low or inactive level on the interrupt line 16 of this interrupt level. This situation of a measured low level 22 following an upward transition 18 on an interrupt request input does not conform to the convention for interrupts to an INTEL 8259. The INTEL 8259 is designed so that in such a situation, interrupt level 7 is indicated by default as the level requesting service even though this may not be the case.

The open-collector output approach has the further disadvantage that there may be two outstanding interrupt requests on the same external interrupt line. The second and subsequent requests are not explicitly handled by the INTEL 8259. Once the positive transition 18 has been detected by the INTEL 8259, subsequent positive transitions have no effect until the INTEL 8259

recognizes the interrupt request on the interrupt line and forwards it to the processor. As a result, when the processor is responding to an interrupt request from a particular interrupt level, it is possible that more than one interrupt source of that interrupt level have an outstanding interrupt request. This the processor, once it has received an interrupt from the INTEL 8259, must poll every interrupt source of that level to determine which of them have outstanding requests. Furthermore, it is possible that the processor may be servicing one interrupt request on a particular interrupt level while a second interrupt source requests service. Thus, the only way that it is possible to determine that all outstanding interrupt requests on a particular interrupt level has been serviced is to not only poll each interrupt source of that level but also to completely poll the sources of that level without finding an outstanding interrupt request. Needless to say, this approach exacts a high overhead for the processor.

The article published in the IBM Technical Disclosure Bulletin Vol. 18 No. 3, August 1975, pages 767 and 768 describes an interface circuit which allows an interrupt line to be shared by several devices. This interface circuit cannot be used with the INTEL 8259 programmable interrupt controller since an interrupt request is indicated on the interrupt line by a voltage level (high) and not by a transition from a first level (low) to a second level (high).

Summary of the invention

Accordingly, it is an object of this invention to provide interrupt level sharing.

It is a further object of this invention to provide interrupt level sharing with a minimum of processor overhead.

It is yet another object of this invention to provide interrupt level sharing in which there are no phantom interrupts and in which interrupt requests are not lost.

A further object of this invention is to provide a method of interrupt level sharing that is transparent to the user and to the hardware of the user.

The shared-interrupt interface circuit according to the subject invention is defined in Claim 1. It is to be used in an interrupt system having an interconnection line (30). The shared-interrupt interface circuit is associated to a device, permits sharing of interrupt levels by a plurality of devices, and comprises interrupt state means (32) having an interrupt state and a noninterrupt state in response to an internal interrupt signal generated by the associated device requesting an interrupt. It is characterized in that it comprises:

pulse generating means (36, 38, 40) for impressing on said interconnection line a first pulse transition to a first level from a second level when said interrupt state means are in said interrupt state, said interconnection line being held at said first level if any such pulse generating means in the other shared interrupt interface circuits is impressing a first level on said interconnection line, and if not the pulse generating means pro-

vides a second pulse transition from the first level to the second level, which second transition is the external interrupt signal,

re-enabling means (50) responsive to a signal indicating an end of interrupt servicing for generating a release signal which is active at the end of interrupt servicing,

inhibiting means (44) responsive to the second transition from the said first level to said second level on said interconnection line and to the release signal for inhibiting the generation of any subsequent first transition from the second to the first level by the pulse generating means in all interface circuits until the release signal becomes active.

Brief description of the drawings

Figure 1 is a circuit diagram of a prior art interrupt interface circuit.

Figure 2 is a timing diagram illustrating the cause of phantom interrupts in the prior art.

Figure 3 is a schematic diagram of a clocked embodiment of the interrupt interface circuit of the present invention.

Figure 4 is a timing diagram illustrating the operation of Figure 3.

Figure 5 is an illustration of the sequence of interrupt servicing by interrupt handlers.

Figure 6 is a schematic diagram of an unlocked embodiment of the present invention.

Brief description of the preferred embodiments

Figure 3 is a schematic diagram of the preferred embodiment of the interrupt interface circuit of the present invention. The illustrated embodiment is designed for operation with a clocked system that is clocked by the system clock signal SYS CLK. At this point, assume that the signal IRQ on an external interrupt line 30 is in its high state 19, as illustrated in the timing diagram of Figure 4. An internal interrupt signal CARD INT is in a low state when no interrupt is being requested by the card. However, CARD INT transitions to a high state when the card is requesting an interrupt and is held there by a latch 32, shown in Figure 3. The latching function is usually provided by the logic circuitry providing the internal interrupt signal. The internal interrupt signal CARD INT and the external interrupt signal IRQ are combined in an AND gate 34. The third input to the AND gate 34 INT ENABLE is an extra feature possible with this invention, which will be described later. For the present, assume that INT ENABLE is high. In the presence of a high internal interrupt signal CARD INT and with the external interrupt signal IRQ being high, the AND gate 34 passes a positive signal to the data input of a D-type flip-flop 36 that has its clock input connected to the system clock SYS CLK. As a result, at the first upward transition of the system clock SYS CLK following the upward transition of the internal interrupt signal CARD INT, the flip-flop 36 outputs on its positive output Q₁ a high signal, as illustrated in Figure 4. This signal is passed to the data input of another D-type flip-flop 38, also with its clock input con-

connected to the system clock SYS CLK. As a result, one clock period later a negative transition to a low level appears on the complemented output $\bar{Q}_2$ of the flip-flop 38. The complemented output $\bar{Q}_2$ is connected to the complemented control input of a tri-state buffer 40, such as a 74LS125, with its input grounded. The low level from $\bar{Q}_2$ thereby opens the tri-state buffer 40 and grounds the external interrupt line 30 to produce the first transition 17 of signal IRQ shown in Figure 4.

A feedback circuit is provided between the external interrupt line 30 to the AND gate 34 of this circuit as well as to all other similar interrupt interface circuits of this interrupt level. As a result, the low signal IRQ on the external interrupt line 30 disables the AND gate 34 and, one clock period later, as shown in Figure 4, the Q_1 output of the flip-flop 36 goes low.

Then one clock cycle later, the complemented output $\bar{Q}_2$ of the flip-flop 38 goes high and, as a result, the tri-state buffer 40 is disabled and presents a high impedance to the external interrupt line 30. Assuming that no other interrupt circuit is grounding the external interrupt line 30, this line is left floating. However, a pull-up resistor 42 is connected between the external interrupt line 30 and a logic high voltage source, +5V for TTL logic. As a result, the pull-up resistor 42 pulls the signal IRQ on the external interrupt line 30 to the positive high level 19 as illustrated in Figure 4 (assuming that the external interrupt line 30 is not grounded by another interrupt interface circuit). The pull-up resistor 42 preferably has a resistance value of 8.2 kohms which allows a reasonably rapid pull-up of the external interrupt line 30 yet still allows multiple pull-up resistors on the external interrupt line 30 associated with the other interrupt circuits. The positive transition 18 of the IRQ signal is the positive transition which is recognized by the INTEL 8259 as indicating a new interrupt request.

The signal IRQ on the external interrupt line 30 is also fed back to the clock input of yet another D-type flip-flop 44 that has its data input connected to a logic high voltage source, once again +5V. Thus, the positive transition 18 that is recognized by the INTEL 8259 causes the complemented output $\bar{Q}_3$ of flip-flop 44 to go low. This signal $\bar{Q}_3$ is connected to the complemented clear inputs of the flip-flops 36 and 38. The continued low value of the $\bar{Q}_3$ signal disables the flip-flops 36 and 38 from producing any further pulses on the external interrupt line 30 regardless of possible new internal interrupt signals CARD INT. It is to be noted that this disabling occurs in the interrupt circuit of Figure 3 even if this particular circuit was not the interrupt circuit producing the interrupt request IRQ on the external interrupt line 30. Thus, as soon as any one of the interrupt interface circuits attached to the external interrupt line 30, defined as the interrupt level of interest, outputs a negative pulse to the external interrupt line 30, no further interrupt requests are produced on the external interrupt line 30. However, the latch 32 continues to hold the internal interrupt signal

CARD INT pending further action. It is also to be noted that the low pulse of the IRQ signal on the external interrupt line 30 disables the AND gate 34 so that any new internal interrupt request CARD INT appearing during the negative IRQ pulse is prevented from propagating to the pulse generator embodied in the flip-flop 38.

The external interrupt line 30 is connected to one of the interrupt request inputs of the INTEL 8259 controller. All interrupt circuits so connected to a particular INTEL 8259 interrupt input comprise a single interrupt level. The INTEL 8259 recognizes the positive transition 18 of the IRQ signal and makes the proper measurement of the high IRQ signal so as to determine which of the interrupt levels is requesting service. Because only a single IRQ pulse is allowed on the external interrupt line 30, there is no possibility of a phantom interrupt. Furthermore, the circuit of Figure 3 does not process second and further pulses for that interrupt level. When an interrupt request is received by the INTEL 8259, it is possible that the request is not immediately serviced because there are higher priority requests outstanding. However, eventually the interrupt request of the particular interrupt level is serviced. Software must be provided which handles the interrupt for each interrupt source. For example, a series of interrupt handlers, illustrated in Figure 5 are provided. There must be at least one interrupt handler for each interrupt interface circuit. Each interrupt handler queries its particular interrupt source to determine if it has an outstanding interrupt request. If the interrupt source does have an outstanding interrupt request, the interrupt handler performs some preassigned task. It should be noted that while only a single interrupt source produced the interrupt request IRQ, there may be several outstanding interrupt requests for the interrupt level. Once the interrupt handler has serviced an outstanding interrupt request, the interrupt handler turns off the internal interrupt signal CARD INT. At this point, two sequences of action are possible. The next interrupt handler in the chain of Figure 5 can query its associated interrupt source for the existence of an outstanding interrupt request. Alternatively, the processor can terminate its handling of an interrupt signal from the INTEL 8259 after a single handling of an interrupt request. The choice between the two sequences depends upon the likelihood of two interrupt requests of the same interrupt level being outstanding concurrently, as balanced against the additional overhead of the polling of the additional levels of interrupt handlers.

The final step of servicing the interrupt request for a particular level is to reenables the interrupt interface circuits for that level. This can be performed by an I/O write to device address 02FX_H where X corresponds to the interrupt level. For instance, a command OUT(02F7_H) causes a reenables signal to all the interrupt interface circuits of interrupt level 7.

The interrupt interface circuit of Figure 3 has a

port to the I/O channel containing the address bus for the addressed I/O device. The address bus is connected to an address comparator 46 which compares the address A_0-A_9 against an internal value, for example, 02F7_H, for interrupt level 7. If all the comparisons are true, a high signal is sent to an AND gate 48. The address comparator 46 can be replaced by inverters attached to selected address lines A_0-A_9 for those address lines which would be zero or low for the present device. Then the address lines A_0-A_9 , possibly passing through inverters, are all led to inputs of the AND gate 48. There are two additional inputs to the AND gate 48, an IOW signal indicating an I/O write operation and an AEN signal indicating that the write operation is the type used for the required I/O write rather than a direct memory access write.

The output of the AND gate 48 is connected through an OR gate 50 to the complemented clear input of the flip-flop 44. The function of the other input LOCAL REENABLE to the OR gate will be described later and is an improvement to the invention. The OR gate 50 has a complemented output so that, when the OUT(02FX_H) operation addresses the interrupt level of the interrupt interface circuit, the flip-flop 44 is cleared. Barring any further pulses on the external interrupt line 30, the $\bar{Q}_3$ output goes high to remove the clear of the flip-flops of 36 and 38, thus reenabling them. Following the reenabling, any active internal interrupt signals CARD INT causes the production of an IRQ pulse on the external interrupt line 30.

The two extra features in the interrupt interface circuit of Figure 3 will now be described. The interrupt enable input INT ENABLE to the AND gate 34 provides the means for disabling the interrupt of a particular interrupt interface circuit. Whenever the INT ENABLE signal is low, the output of the AND gate 34 is always low, thereby preventing the interrupt interface circuit from pulsing an interrupt request on the external interrupt line 30. This additional input makes the interrupt interface circuit compatible with software previously written for the IBM PC microcomputer in which there was no interrupt level sharing and the interrupt sources were selectively disabled from producing interrupt requests. This control is necessary in a compatible system where the interrupt mask register of the previously written software would disable all the shared interrupts. Of course, for full sharing of the interrupt level, the INT ENABLE signal is kept high.

The LOCAL REENABLE input to the OR gate 50 allows a single interrupt interface circuit to monopolize the shared interrupt facility. If an I/O device requires exclusive use of the shared interrupt level, it may effect this control by having its associated interrupt interface circuit disable all other devices. It is of course preferable that this exclusive use of the interrupt level be effected without knowing the details of the other I/O devices sharing the interrupt level.

To gain exclusive use of the shared interrupt, an

I/O device causes an interrupt, perhaps under the control of the device handler in the system software. The interrupt handler, in servicing the interrupt request, does not reenable all the interrupt interface circuits of the interrupt level with the previously described OUT(02FX_H), but rather causes a high LOCAL REENABLE. The LOCAL REENABLE is specific to a particular interface circuit and is generally not commonly conveyed to other interrupt interface circuits of the same interrupt level. The means of communication between the processor and the interrupt interface circuit for the LOCAL REENABLE signal can be an I/O port using components similar to the address comparator 46 and the AND gate 48. The interrupt handler would then issue an I/O write to a specific address to initiate the LOCAL REENABLE signal. The interrupt handler may also choose to revector the interrupt in the INTEL 8259 so that all interrupt requests detected by the INTEL 8259 are directed to the monopolizing interrupt interface circuit. As long as the interrupt handler does not issue the OUT(02FX_H) write command, the monopolizing interrupt interface circuit will have exclusive use of the interrupts because the other interrupt circuits remain inhibited by the first negative IRQ pulse on the external interrupt line 30 and have not been subsequently reenabled. In order to reenable the other I/O devices, the device handler must reestablish the original interrupt vector and do the global reenable with an OUT(02FX_H) write instruction.

It should be pointed out that the external interrupt line 30 is acting as a bidirectional line. The interrupt interface circuits both transmit IRQ pulses on the external interrupt line and also receive these pulses for feedback to the AND gate 34 and the flip-flop 44. As a result, the IRQ lines 30 cannot be redriven between the interrupt interface circuits comprising a shared interrupt level. All cards using the interrupt sharing mechanism must be in the same unit in the case where drivers are used between units. However, cards for different interrupt levels may occupy different boxes.

In the IBM PC microcomputer, the external interrupt line 30 is one of six lines of an interrupt bus. Each line of the interrupt bus is equivalent and defines the interrupt level.

The interrupt interface circuit of Figure 3 is for a clocked system. An equivalent interrupt interface circuit for an unclocked system is illustrated in the schematic diagram of Figure 6. The flip-flops 36 and 38 of Figure 3 are replaced by a one-shot 56 that has a complemented enable input connected to the complemented output $\bar{Q}_3$ of the flip-flop 44. The remaining elements are similar to those of Figure 3. A positive output of the AND gate 34 causes the one-shot 56 to produce a negative pulse on the external interrupt line 30 as long as the one-shot 56 is not disabled by the flip-flop 44. A negative pulse on the external interrupt line 30 disables the one-shot 56 until the flip-flop 44 is reenabled. Thus, the operation of the interrupt interface circuit of Figure 6 is seen to be the same

as that of Figure 3 except for the clocking. The I/O port for the reenabling signal is not explicitly shown nor the added features of Figure 3, which all can be copied from the clocked system.

Thus it is seen that the interrupt interface system of the present invention allows the sharing of an interrupt level by multiple interrupt sources. The interrupt level sharing can be accomplished at the card level, that is, at the interrupt source. Furthermore, the interrupt sharing is accomplished without hardware change to the external interrupt bus or to the INTEL 8259 programmable interrupt controller. However, the software of the interrupt handlers does require modification to provide the polling determination of which interrupt source has an outstanding request for service.

Claims

1. In an interrupt system having an interconnection line (30) for an external interrupt signal (IRQ), a shared-interrupt interface circuit associated to a device for permitting sharing of interrupt levels by a plurality of devices, said shared interrupt interface circuit comprising interrupt state means (32) having an interrupt state and a noninterrupt state in response to an internal interrupt signal (CARD INT) generated by the associated device requesting an interrupt, and being characterized in that it comprises:

pulse generating means (36, 38, 40) for impressing on said interconnection line (30) a first pulse transition (17) to a first level (20) from a second level (19) when said interrupt state means (32) are in said interrupt state, said interconnection line (30) being held at said first level (20) if any such pulse generating means (36, 38, 40) in other shared interrupt interface circuits is impressing a first level on said interconnection line, (30) and if not the pulse generating means (36, 38, 40) provides a second pulse transition (18) from the first level (20) to the second level (19), which second transition (18) is the external interrupt signal,

re-enabling means (50) responsive to a signal indicating an end of interrupt servicing for generating a release signal which is active at the end of interrupt servicing,

inhibiting means (44) responsive to the second transition (18) from the said first level (20) to said second level (19) on said interconnection line (30) and to the release signal for inhibiting the generation of any subsequent first transition (17) from the second level (19) to the first level (20) by the pulse generating means (36, 38, 40) in all shared interrupt interface circuits until the release signal becomes active.

2. A shared interrupt interface circuit according to Claim 1, characterized in that it comprises signal combining means (34) connected to said interrupt state means (32) and said interconnection line (30) for preventing said pulse generating means from initiating the generation of said first level signal when a first level signal is present on said interconnection line (30).

3. A shared interrupt interface circuit according

to Claim 1 or 2, wherein said inhibiting means (44) and said reenabling means (50) together comprise a flip-flop (44), said flip-flop being connected to said interconnection line (30) and to a first line carrying said end of interrupt servicing signal, said flip-flop (44) being set by said transition (17) to said first level (20) on said interconnection line 30 and being cleared by said end of interrupt service signal on said first line, the output of said flip-flop (44) being connected to an enabling input (CLR) of said pulse generating means.

4. A shared interrupt interface circuit according to Claim 1, 2 or 3 characterized in that:

said pulse generating means (36, 38, 40) comprises output means (40) having a first state of said first level, said first state being a voltage state, and a second state of high impedance; and

further comprising voltage pull-up means (42) connected to said interconnection line (30) for pulling the voltage on said interconnection line (30) to said second level when only high impedance output means are connected to said interconnection line (30).

5. A shared interrupt interface circuit according to any one of Claims 1 to 4, characterized in that it comprises an input port means (IO, 46) accessible by an addressable output function of a computer system, said input port receiving a signal from said output function producing said end of interrupt servicing signal.

6. A shared interrupt interface circuit according to claim 1 to be used with a computer system operating with a clock (SYS CLOCK) and having an addressable output function (IO), said computer system responding to an interrupt when the signal on the interconnection line (30) for the external interrupt signal changes from the first level (20) to the second level (19) and being characterized in that it comprises:

a signal combining gate (34) receiving said interconnection line (30) and an internal interrupt line (CARD INT), the signal on said internal interrupt line changing state when an interrupt is requested;

a first flip-flop (36) receiving the output of said signal combining gate (34) and said clock;

the pulse generating circuit (38) receiving the output of said first flip-flop (36) and clocked by said clock for producing a negative pulse on said interconnection line (30), said pulse generating circuit (38) having an output circuit (40, 42) that pulls said interconnection line (30) to said first level regardless of the outputs of other similar output circuits in other shared interrupt interface circuits connected to said interconnection line (30), maintains the interconnection line (30) at said first level if any pulse generating circuit is pulling the interconnection line (30) to the first level and if not pulls up the interconnection line (30) toward the second level,

a second flip-flop (44) receiving the interconnection line and having an output connected to disabling inputs (CLR) of said first flip-flop and of said pulse generating circuit; and

an input port means (IO, 46, 48, 50) accessible by

as that of Figure 3 except for the clocking. The I/O port for the reenable signal is not explicitly shown nor the added features of Figure 3, which all can be copied from the clocked system.

Thus it is seen that the interrupt interface system of the present invention allows the sharing of an interrupt level by multiple interrupt sources. The interrupt level sharing can be accomplished at the card level, that is, at the interrupt source. Furthermore, the interrupt sharing is accomplished without hardware change to the external interrupt bus or to the INTEL 8259 programmable interrupt controller. However, the software of the interrupt handlers does require modification to provide the polling determination of which interrupt source has an outstanding request for service.

Claims

1. In an interrupt system having an interconnection line (30) for an external interrupt signal (IRQ), a shared-interrupt interface circuit associated to a device for permitting sharing of interrupt levels by a plurality of devices, said shared interrupt interface circuit comprising interrupt state means (32) having an interrupt state and a noninterrupt state in response to an internal interrupt signal (CARD INT) generated by the associated device requesting an interrupt, and being characterized in that it comprises:

pulse generating means (36, 38, 40) for impressing on said interconnection line (30) a first pulse transition (17) to a first level (20) from a second level (19) when said interrupt state means (32) are in said interrupt state, said interconnection line (30) being held at said first level (20) if any such pulse generating means (36, 38, 40) in other shared interrupt interface circuits is impressing a first level on said interconnection line, (30) and if not the pulse generating means (36, 38, 40) provides a second pulse transition (18) from the first level (20) to the second level (19), which second transition (18) is the external interrupt signal,

re-enabling means (50) responsive to a signal indicating an end of interrupt servicing for generating a release signal which is active at the end of interrupt servicing,

inhibiting means (44) responsive to the second transition (18) from the said first level (20) to said second level (19) on said interconnection line (30) and to the release signal for inhibiting the generation of any subsequent first transition (17) from the second level (19) to the first level (20) by the pulse generating means (36, 38, 40) in all shared interrupt interface circuits until the release signal becomes active.

2. A shared interrupt interface circuit according to Claim 1, characterized in that it comprises signal combining means (34) connected to said interrupt state means (32) and said interconnection line (30) for preventing said pulse generating means from initiating the generation of said first level signal when a first level signal is present on said interconnection line (30).

3. A shared interrupt interface circuit according

to Claim 1 or 2, wherein said inhibiting means (44) and said reenabling means (50) together comprise a flip-flop (44), said flip-flop being connected to said interconnection line (30) and to a first line carrying said end of interrupt servicing signal, said flip-flop (44) being set by said transition (17) to said first level (20) on said interconnection line 30 and being cleared by said end of interrupt service signal on said first line, the output of said flip-flop (44) being connected to an enabling input (CLR) of said pulse generating means.

4. A shared interrupt interface circuit according to Claim 1, 2 or 3 characterized in that:

said pulse generating means (36, 38, 40) comprises output means (40) having a first state of said first level, said first state being a voltage state, and a second state of high impedance; and

further comprising voltage pull-up means (42) connected to said interconnection line (30) for pulling the voltage on said interconnection line (30) to said second level when only high impedance output means are connected to said interconnection line (30).

5. A shared interrupt interface circuit according to any one of Claims 1 to 4, characterized in that it comprises an input port means (IO, 46) accessible by an addressable output function of a computer system, said input port receiving a signal from said output function producing said end of interrupt servicing signal.

6. A shared interrupt interface circuit according to claim 1 to be used with a computer system operating with a clock (SYS CLOCK) and having an addressable output function (IO), said computer system responding to an interrupt when the signal on the interconnection line (30) for the external interrupt signal changes from the first level (20) to the second level (19) and being characterized in that it comprises:

a signal combining gate (34) receiving said interconnection line (30) and an internal interrupt line (CARD INT), the signal on said internal interrupt line changing state when an interrupt is requested;

a first flip-flop (36) receiving the output of said signal combining gate (34) and said clock;

the pulse generating circuit (38) receiving the output of said first flip-flop (36) and clocked by said clock for producing a negative pulse on said interconnection line (30), said pulse generating circuit (38) having an output circuit (40, 42) that pulls said interconnection line (30) to said first level regardless of the outputs of other similar output circuits in other shared interrupt interface circuits connected to said interconnection line (30), maintains the interconnection line (30) at said first level if any pulse generating circuit is pulling the interconnection line (30) to the first level and if not pulls up the interconnection line (30) toward the second level,

a second flip-flop (44) receiving the interconnection line and having an output connected to disabling inputs (CLR) of said first flip-flop and of said pulse generating circuit; and

an input port means (IO, 46, 48, 50) accessible by



said output function of said computer system, the output of said input port means being connected to a disabling input of said second flip-flop (44).

7. A shared interrupt interface circuit according to Claim 6, wherein said output circuit (40, 42) selectively produces a low level voltage signal on said interconnection line (30) or presents a high impedance to said interconnection line (30) and further comprises a pull-up resistor (42) connected between said interconnection line (30) and a positive voltage.

Patentansprüche

1. In einem Interruptsystem mit einer Verbindungsleitung zu einem externen Interruptsignal (IRQ), eine Schaltung mit geteilter Interruptschnittstelle, verbunden mit einer Vorrichtung, die die Teilung der Interruptpegel für eine Mehrzahl Vorrichtungen gestattet, wobei besagte Schaltung mit geteilter Interruptschnittstelle Zustandsmittel (32) mit einem Interrupt-Zustand und einem Nicht-Interrupt-Zustand aufweist, je nach einem internen Interruptsignal (INT-Karte), erzeugt durch die zugeordnete Vorrichtung, die einen Interrupt anfordert, dadurch gekennzeichnet, daß sie enthält:

Impulserzeugungsmittel (36, 38, 40) für das Anlegen eines ersten Impulsübergangs (17) an besagte Verbindungsleitung (30) auf einen ersten Pegel (20) von einem zweiten Pegel (19), wenn besagte Interruptzustandsmittel (32) in besagtem Interrupt-Zustand sind, und besagte Verbindungsleitung (30) auf besagtem erstem Pegel (20) gehalten wird, wenn irgendwelche solche Impulserzeugungsmittel (36, 38, 40) in den anderen Schaltungen mit geteilter Interruptschnittstelle der besagten Verbindungsleitung (30) einen ersten Pegel mitteilen, und wenn die Impulserzeugungsmittel (36, 38, 40) nicht einen zweiten Impulsübergang (18) vom ersten Pegel (20) auf den zweiten Pegel (19) bereitstellen, wobei dieser zweite Übergang (18) das externe Interruptsignal ist,

Mittel (50) für die Wiederherstellung, ansprechend auf ein Signal, das das Ende des Interruptzustands anzeigt, um ein Freigabesignal zu erzeugen, das am Ende des Interruptzustands aktiv ist,

Sperrmittel (44), die auf den zweiten Übergang (18) vom besagten ersten Pegel (20) auf den besagten zweiten Pegel (19) auf besagter Verbindungsleitung (30) und auf das Freigabesignal ansprechen, um die Erzeugung irgendeines nachfolgenden ersten Übergangs (17) vom zweiten Pegel (19) auf den ersten Pegel (20) durch die Impulserzeugungsmittel (36, 38, 40) in allen Schaltungen mit geteilter Interruptschnittstelle zu verhindern, bis das Freigabesignal aktiv wird.

2. Schaltung mit einer geteilten Interruptschnittstelle gemäß Anspruch 1, dadurch gekennzeichnet, daß sie Signalkombinationsmittel (34) enthält, die an besagte Interruptzustandsmittel (32) und an die besagte Verbindungsleitung (30) angeschlossen sind, um die besagten Impulserzeugungsmittel daran zu hindern, die Erzeugung

des Signals des besagten ersten Pegels auszulösen, wenn ein Signal des ersten Pegels auf der besagten Verbindungsleitung (30) ansteht.

3. Schaltung mit einer geteilten Interruptschnittstelle gemäß Anspruch 1 oder 2, in der besagte Sperrmittel (44) und besagte Freigabemittel (50) zusammen eine Flip-Flop-Schaltung (44) enthalten, wobei besagte Flip-Flop-Schaltung an die besagte Verbindungsleitung (30) und an eine erste Leitung, die das besagte Ende des Interruptzustandsignals überträgt, angeschlossen ist, und besagte Flip-Flop-Schaltung (44) durch den besagten Übergang (17) auf den besagten ersten Pegel (20) auf der besagten Verbindungsleitung (30) gesetzt wird, und durch das besagte Ende des Interruptzustandsignals auf der besagten ersten Leitung gelöscht wird, wobei der Ausgang der besagten Flip-Flop-Schaltung (44) an einen Freigabeeingang (CLR) der besagten Impulserzeugermitten angeschlossen ist.

4. Schaltung mit einer geteilten Interruptschnittstelle gemäß Anspruch 1, 2 oder 3, dadurch gekennzeichnet, daß:

die besagten Impulserzeugungsmittel (36, 38, 40) Ausgangsmittel (40) enthalten, die einen ersten Zustand des besagten ersten Pegels aufweisen, wobei der besagte erste Zustand ein Spannungszustand ist, und einen zweiten Zustand mit hoher Impedanz; und

ferner Spannungssteigerungsmittel (42) enthaltend, die an besagte Verbindungsleitung (30) angeschlossen sind, um die Spannung auf der besagten Verbindungsleitung (30) auf den besagten zweiten Pegel zu bringen, wenn an besagte Verbindungsleitung (30) nur Mittel mit hohem Impedanzausgang angeschlossen sind.

5. Schaltung mit einer geteilten Interruptschnittstelle gemäß einem der Ansprüche 1 bis 4, dadurch gekennzeichnet, daß sie ein Eingangsschlußmittel (10, 46) enthält, das über eine adressierbare Ausgangsfunktion eines Rechnersystems zugänglich ist, wobei besagter Eingangsschluß ein Signal von der besagten Ausgabefunktion erhält und das Ende des Interruptzustandsignals bewirkt.

6. Schaltung mit einer geteilten Interruptschnittstelle gemäß Anspruch 1, zu verwenden mit einem Rechnersystem, das mit einem Taktgeber arbeitet (SYS CLOCK) und eine adressierbare Ausgangsfunktion (IO) besitzt, wobei besagtes Rechnersystem auf einen Interrupt antwortet, wenn das Signal auf der Verbindungsleitung (30) für das externe Interruptsignal vom ersten Pegel (20) auf den zweiten Pegel (19) übergeht, dadurch gekennzeichnet, daß sie enthält:

eine Signalkombinationsverknüpfung, die die besagte Verbindungsleitung (30) und eine interne Interruptleitung (CARD INT) empfängt, wobei das Signal auf der besagten internen Interruptleitung den Zustand wechselt, wenn ein Interrupt angefordert wird;

eine erste Flip-Flop-Schaltung (36), die den Ausgang der besagten Signalkombinationsverknüpfung (34) und den besagten Takt empfängt; die Impulserzeugungsschaltung (38), die den



Ausgang der besagten ersten Flip-Flop-Schaltung (36) empfängt und durch besagten Taktgeber getaktet wird, um einen negativen Impuls auf der besagten Verbindungsleitung (30) zu erzeugen, wobei die besagte Impulsgenerierschaltung (38) eine Ausgangsschaltung (40, 42) besitzt, die die besagte Verbindungsleitung (30) auf den besagten ersten Pegel bringt, ohne Berücksichtigung anderer änderlicher Ausgangsschaltungen in anderen Schaltungen mit geteilter Interrupt-Schnittstelle, die an die besagte Verbindungsleitung (30) angeschlossen sind, die Verbindungsleitung (30) auf der besagten ersten Ebene hält, wenn eine Impulserzeugungsschaltung die Verbindungsleitung (30) auf den ersten Pegel bringt, und wenn nicht, die Verbindungsleitung (30) auf den zweiten Pegel bringt;

eine zweite Flip-Flop-Schaltung, die die Verbindungsleitung aufnimmt und einem Ausgang besitzt, der an die Abschalteneingänge (CLR) der besagten ersten Flip-Flop-Schaltung und der besagten Impulserzeugerschaltung angeschlossen ist; und

ein Eingangsanschlußmittel (10, 46, 48, 50), das über die besagte Ausgabefunktion des besagten Rechnersystems zugänglich ist, wobei der Ausgang des besagten Eingangsanschlußmittels an einen Abschalteneingang der besagten Flip-Flop-Schaltung angeschlossen ist.

7. Schaltung mit einer geteilten Interrupt-Schnittstelle gemäß Anspruch 6, in der die besagte Ausgangsschaltung (40, 42) wahlweise ein Signal mit niedrigem Pegel auf der besagten Verbindungsleitung (30) erzeugt oder eine hohe Impedanz an die besagte Verbindungsleitung (30) anlegt und außerdem einen Pull-up-Widerstand (42) enthält, der zwischen der Verbindungsleitung (30) und einer positiven Spannung angeschlossen ist.

Revendications

1. Dans un système d'interruption comprenant une ligne d'interconnexion (30) pour un signal d'interruption externe (IRQ), un circuit d'interface d'interruption partagée associé à un dispositif permettant le partage de niveaux d'interruption par une pluralité de dispositifs, ledit circuit d'interface d'interruption partagée comprenant des moyens d'état d'interruption (32) ayant un état d'interruption et un état de non-interruption en réponse à un signal d'interruption interne (CARD INT) généré par le dispositif associé demandant une interruption, et étant caractérisé en ce qu'il comprend:

des moyens de génération d'impulsions (36, 38, 40) pour appliquer sur ladite ligne d'interconnexion (30) une première transition d'impulsion (17) à un premier niveau (20) à partir d'un second niveau (19) lorsque lesdits moyens d'état d'interruption (32) sont dans ledit état d'interruption, ladite ligne d'interconnexion (30) étant maintenue audit premier niveau (20) si aucun de ces moyens de génération d'impulsions (36, 38, 40) dans les autres circuits d'interface d'interruption

partagée n'est en train d'appliquer un premier niveau sur ladite ligne d'interconnexion (30) et si les moyens de génération d'impulsions (36, 38, 40) n'appliquent pas une seconde transition d'impulsion (18) à partir du premier niveau (20) au second niveau (19), laquelle seconde transition (18) est le signal d'interruption externe,

des moyens de reconditionnement (50) sensibles à un signal indiquant la fin d'un service d'interruption pour générer un signal de dégagement qui est actif à la fin d'un service d'interruption,

des moyens de déconditionnement (44) sensibles à la seconde transition (18) dudit premier niveau (20) audit second niveau (19) sur ladite ligne d'interconnexion (30) et au signal de dégagement pour déconditionner la génération de toute première transition ultérieure (17) du second niveau (19) au premier niveau (20) par les moyens de génération d'impulsions (36, 38, 40) dans tous les circuits d'interface d'interruption partagée jusqu'à ce que le signal de dégagement devienne actif.

2. Un circuit d'interface d'interruption partagée selon la revendication 1 caractérisé en ce qu'il comprend des moyens de combinaison de signaux (34) connectés auxdits moyens d'état d'interruption (32) et à ladite ligne d'interconnexion (30) pour interdire auxdits moyens de génération d'impulsions d'initier la génération dudit signal de premier niveau lorsqu'un signal de premier niveau est présent sur ladite ligne d'interconnexion (30).

3. Un circuit d'interface d'interruption partagée selon la revendication 1 ou 2 dans lequel lesdits moyens de déconditionnement (44) et lesdits moyens de reconditionnement (50) comprennent ensemble une bascule (44), ladite bascule étant connectée à ladite ligne d'interconnexion (30) et à une première ligne portant ledit signal de service de fin d'interruption, ladite bascule (44) étant conditionnée par ledit transition (17) audit premier niveau (20) sur ladite ligne d'interconnexion (30) et étant effacée par ledit signal de service de fin d'interruption sur ladite première ligne, la sortie de ladite bascule (44) étant connectée à une entrée de conditionnement (CLR) desdits moyens de génération d'impulsions.

4. Un circuit d'interface d'interruption partagé selon la revendication 1, 2 ou 3 caractérisé en ce que:

lesdits moyens de génération d'impulsions (36, 38, 40) comprennent des moyens de sortie (40) ayant un premier état dudit premier niveau, ledit premier état étant un état de tension, et un second état d'impédance élevée, et

comprenant en outre des moyens de relèvement de tension (42) connectés à ladite ligne d'interconnexion (30) pour relever la tension sur ladite ligne d'interconnexion (30) audit second niveau lorsqu'il n'y a que les moyens de sortie d'impédance élevée qui sont connectés à ladite ligne d'interconnexion (30).

5. Un circuit d'interface d'interruption partagée selon l'une quelconque des revendications 1 à 4,



caractérisé en ce qu'il comprend un moyen à port d'entrée (10, 46) accessible par une fonction de sortie adressable d'un système d'ordinateur, ledit port d'entrée recevant un signal de ladite fonction de sortie produisant ledit signal de service de fin d'interruption.

6. Un circuit d'interface d'interruption partagée selon la revendication 1 à utiliser avec un système d'ordinateur fonctionnant avec une horloge (SYS CLOCK) et ayant une fonction de sortie adressable (10), ledit système d'ordinateur répondant à une interruption lorsque le signal sur la ligne d'interconnexion (30) pour le signal d'interruption externe passe du premier niveau (20) au second niveau (19) et étant caractérisé en ce qu'il comprend:

une porte de combinaison de signaux (34) recevant ladite ligne d'interconnexion (30) et une ligne d'interruption interne (CARD INT), le signal sur ladite ligne d'interruption interne changeant d'état lorsqu'une interruption est demandée,

une première bascule (36) recevant la sortie de ladite porte de combinaison de signaux (34) et ladite horloge,

le circuit de génération d'impulsions (38) recevant la sortie de ladite première bascule (36) et synchronisé par ladite horloge pour produire une impulsion négative sur ladite ligne d'interconnexion (30), ledit circuit de génération d'impulsions (38) ayant un circuit de sortie (40, 42) qui

relève ladite ligne d'interconnexion (30) audit premier niveau indépendamment des sorties d'autres circuits de sortie similaires dans d'autres circuits d'interface d'interruption partagée connectés à ladite ligne d'interconnexion (30), maintient la ligne d'interconnexion (30) audit premier niveau si un quelconque circuit de génération d'impulsions est en train de relever la ligne d'interconnexion (30) au premier niveau et ne relève pas la ligne d'interconnexion (30) vers le second niveau,

une seconde bascule (44) recevant la ligne d'interconnexion et ayant une sortie connectée aux entrées de déconditionnement (CLR) de ladite première bascule et dudit circuit de génération d'impulsions, et

un moyen de port d'entrée (10, 46, 48, 50) accessible par ladite fonction de sortie dudit système d'ordinateur, la sortie desdits moyens de port d'entrée étant connectée à une entrée de déconditionnement de ladite seconde bascule (44).

7. Un circuit d'interface d'interruption partagé selon la revendication 6 dans lequel ledit circuit de sortie (40, 42) produit sélectivement un signal de tension de faible niveau sur ladite ligne d'interconnexion (30) ou présente une impédance élevée à ladite ligne d'interconnexion (30) et comprend en outre une résistance de relèvement (42) connectée entre ladite ligne d'interconnexion (30) et une tension positive.

35

40

45

50

55

60

65

10

