

ZERO OVERHEAD COMPUTER INTERRUPTS
WITH TASK SWITCHING

TECHNICAL FIELD

5 This invention relates to interrupt processing and task switching circuitry and methods in data processing systems.

BACKGROUND ART

10 Modern computers use special hardware to automate interrupt processing and task switching. Interrupts include such events as I/O transfers, clocks, hardware faults, software interrupts and software faults such as divide by zero. So that the interrupting program
15 does not alter the running program, all states of the interrupted task must be saved. Then, all of the new interrupting task states must be retrieved from memory and the processing then continued with these new interrupting task states. Once the interrupting task is
20 completed, a Return From Interrupt instruction is usually executed causing the state save process to reverse itself and restore the previous task.

 Most modern computers achieve the task change process using a hardware interrupt mechanism that first
25 disables additional interrupts and saves only the minimum number of state registers before jumping to a special interrupt service program. The special interrupt service program then saves the rest of the task registers as may be required. It then loads the register for the
30 interrupting tasks. This process is reversed for returning from interrupts so that the computer may then return to its original task.

 During the saving and restoring of the registers, any further attempted interrupts must be
35 disabled to prevent corruption of the system's existing data. Thus, the computer is insensitive to interrupts for these periods of time, known as interrupt latency.

Traditional computer architectures must complete current instructions before responding to interrupts. Typically, these instructions vary widely in their durations. In addition, various computer architectures have different interrupt latencies depending upon the modes of the currently executing process and that of the interrupting process. These two factors combine to cause highly non-deterministic interrupt latency delays and overhead.

Besides the latency delay associated with beginning each new interrupt task, in state-of-the-art computers there is also the time spent switching between tasks. This switching time is the interrupt and task change processing overhead. It is the time used for saving and restoring the registers, and thus is time unavailable for use in processing. This interrupt and task change processing overhead amounts to tens or hundreds of cycles in many modern computers. Even at multi-Megahertz clock rates, the combination of both the latency delay and the time required with respect to switching overhead results in computers that can only process interrupts at tens or hundreds of Kiloherzt. Meanwhile, multimedia data transfers take place at tens to hundreds of megahertz.

As a result of these modern computer time limitations in interrupt and task change latency and processing overhead, current computers cannot handle the required data rates without the use of extensive storage buffers and auxiliary specialized processors. The size of these buffers causes multimedia systems to span dozens of chips, each with its additional attendant cost. Even as technological advances place more memory on each integrated circuit, the large size and complexity of current buffered data transfer methods continue to cause excess manufacturing costs while decreasing computer operational performances. Modern applications require an ever greater number of interrupts which simply further aggravates the problem.

Multiple register interrupt and task switch systems have been proposed before, but all have suffered various limitations. In particular, the use of large register sets causes a concomitant increase in wiring and capacitance related delays, slowing down the entire computer. Other methods have used separate register stores for data storage, but this results in a waste of time during the course of the movement of the data to and from these task saving registers before the execution of an interrupt. Then, additional time must be used to load the working data into the general purpose registers before interrupt processing can begin.

A further disadvantage of conventional computer architectures is that they do not incorporate effective methods for tracing processor flow, including task changes, without expensive hardware emulators and logic analyzers. Meanwhile, conventional hardware emulators cannot test computers under actual operational conditions as they cannot run as fast as the computers that they are testing.

Also, computers lack hardware based task linking sequential execution systems and deadline priority "impatience" counters to permit interrupts to automatically increase their priorities with the passage of time.

Meanwhile, multicycle uninterruptable instructions also increase interrupt latency delays. These latency delays can be hundreds or even thousands of cycles long in modern complex computers.

Furthermore, current state-of-the-art CPU's pipeline the processing of data and instructions in order that more data and instructions can be processed in each cycle. This has a number of disadvantages. Whenever an interrupt comes in, the pipeline CPU must abandon whatever is unfinished in the pipeline, and reset the various counters to resume processing when it restarts. This is called a pipeline "stall". When a branch

instruction is executed, the pipeline CPU must either guess at which way the branch will go, stall or try to follow the multiple, possible branch outcomes. All of these various pipeline CPU situations can cause computers to waste processing resources and result in further delays.

Additionally, conventional computer systems require extensive buffering to accommodate high data rates. Therefore, they are not generally as suitable for single-chip, unitary construction fabrication with its advantages in terms of costs, lower energy and cooling expenditures and increased performances.

DISCLOSURE OF THE INVENTION

In light of the above described problems and limitations with regard to the current structures and methods of operating conventional computers, it is an object of the present invention to provide a unique, zero overhead interrupt and task change mechanism for use in computer architectures for detecting interrupts in the background then performing complete state saving and restoring operations between one processor cycle and the next without software intervention while limiting the number of tasks only by the size of the task storage memory system and while preserving the zero overhead aspects of the mechanism regardless of the task memory access time.

It is a further object of this invention to prevent the loss of time by eliminating the need to transfer the running task's data to a separate task state storage element at the instant of the interrupt and task change event.

It is an additional object of the invention to permit the use of part of the system's task memory for the effective tracing of processor flow for full speed, real-time use in computer programming debugging.

It is an additional object of the invention to provide a task linking system to allow the linking of tasks for the mandatory sequential execution of the linked tasks as required in multi-tasking systems.

5 It is also an object of the invention to provide a priority "impatience" counter system to increase the relative priorities of various tasks as their task priority deadlines approach.

10 It is an additional object of the invention's interrupt and task change mechanism to eliminate the need for extensive storage buffers and auxiliary specialized processors with their associated costs and delays.

A further object of the invention is that it allows the interruption of what were previously
15 considered to be uninterruptable instructions by switching the entire state of the task, not just the state of the CPU registers, to eliminate the latency otherwise involved in long multicycle uninterruptable instructions.

20 It is also an object of the invention to create a deterministic computer architecture structure to effectively eliminate latency unpredictability.

It is yet another object of the invention's interrupt and task changing mechanism to eliminate
25 portions of previously required CPU pipeline and branch-prediction circuitries so as to better avoid computer stalls and delays.

Further, it is an object of the invention that the system be designed to be simple and compact enough in
30 its circuitry so as to accomplish its objects within the confines of a single, small, integrated circuit capable of being implemented within a CPU within the limits of conventional fabrication processes so as to reduce the energy and cooling requirements of the invention's
35 circuitry while enhancing computer processing speeds.

A benefit of the present invention is that it allows computers to process both interrupts and task

changes with zero overhead and very few cycles of latency, typically just one or two from local memory. This is accomplished by employing hardware to completely automate the entire interrupt and task change process by
5 detecting interrupts in the background then performing complete state saving and restoring operations between one processor cycle and the next without software intervention while limiting the number of tasks only by the size of the task storage memory system and while
10 preserving the zero overhead of the mechanism regardless of the task memory access time,

The present invention avoids the wiring and capacitance delays associated with large register banks. It uses a relatively small register set which is
15 electrically decoupled from the computer's background process for saving and restoring the register states from a large memory store. This is of substantial benefit due to the much higher density of memory stores in comparison to register banks. In addition, at an expense in
20 latency, external memory stores can be used for storing register states. This substantially increases the capacity of the system for additional available zero overhead tasks.

The interrupt and task change mechanism
25 accomplishes its purpose by using two or more task latches per task register bit. Each such latch may be alternately connected to the large task memory store or the CPU. Thus, while the CPU is using one latch, data in another latch may be saved and further task data may be
30 simultaneously loaded. A small state machine controls the sequencing of the task changes.

The interrupt and task change mechanism's architecture adds less than 2 gate delays over and above a simple register. The access time of the task storage
35 random access memory (RAM) does not add to the interrupt and task change overhead. Current integrated circuit processes allow the task RAM to operate within one cycle,

allowing a "worst case" interrupt and task change latency of 2 cycles.

In contrast to traditional computer systems, deterministic systems have a fixed latency independent of operating or interrupting modes, and independent of the duration of the current instruction.

In the current invention, the system is interruptable at any moment since the full state of the machine is preserved upon interrupt. The interrupt delay is one of two cycles of latency, and is thus far more deterministic than that of traditional computers.

A further benefit of this invention is that systems are much easier to design, test and debug using deterministic interrupt structures.

Yet another benefit of this invention is that it is fault tolerant, it can place itself or be placed in a mode which permits individual step by step debugging.

An additional benefit is that less processing power and less buffering is required to be kept in reserve when using a deterministic as opposed to a non-deterministic architecture resulting in less costly processor configurations.

A further benefit is that because it is deterministic, the invention facilitates new potential interrupt application opportunities.

Since this invention's interrupt and task change mechanism uses one of the task latches as the register connected to the CPU, no data transfer is required to a separate task state storage element as in other methods. The result is yet a further benefit in that time is saved during the course of the process.

An additional benefit of this interrupt and task change system is the ability to create a trace memory out of a portion of the task memory, thus allowing full speed, real-time debugging of computers based on this system.

Yet a further benefit of this invention is its task link system which allows tasks to be linked so that they must execute sequentially. This saves time over software methods.

5 The invention provides an additional benefit in that its task deadline priority counter system allows tasks that have signaled a hardware interrupt to automatically increase their priority with the passage of time. This is particularly appropriate for slower tasks
10 like audio tasks, which can start off at a low priority. Usually the CPU will find time to serve the lower priority tasks, as anytime before the next sample period is fine. Sometimes, however, the CPU has many high priority tasks. In such cases, this invention's task
15 deadline priority counter system gradually increases the priority of the lower priority tasks as their deadlines approach.

A further benefit of the invention is that computer systems designed with this interrupt and task
20 change mechanism can handle multimedia data rates without significant buffering and its attendant control circuits. This allows greater I/O functionality with much less circuitry.

Yet another significant benefit of the
25 invention is that it allows the interruption of what were previously considered to be uninterruptable instructions by switching the entire state of the task, not just the state of the CPU registers, to eliminate the latency otherwise involved in long multicycle interruptable
30 instructions.

An additional benefit of this invention's interrupt and task switching system is its ability to replace pipelines and branch prediction functions in the CPU. The proposed invention solves the current problems
35 which exist with pipelined CPUs. It accomplishes this without pipelines or branch prediction circuitry by simply switching tasks. No branch prediction is

required. The CPU processes branch instructions as they occur. If the branch instruction requires data not in the local memory, the branching task is swapped for the waiting task and processing continues.

5 A further benefit of this invention is that it is simple and compact enough so as to permit its placement upon a single, small, integrated circuit compatible with existing fabrication processes. As a result, energy requirements for the actual circuitry are
10 reduced, possible needs for cooling equipment with concomitant requirements for extra electrical cooling energy are either reduced or eliminated, and the CPU speed of processing is enhanced by the reduction in the size and complexity of the circuitry and the reduction
15 thereby in the system's number of chips and chip interconnections.

The foregoing and other objectives, features, and advantages of the invention will be more readily understood upon consideration of the following detailed
20 description of the invention, taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a circuit diagram of an embodiment of
25 a bit slice of a task register in accordance with the present invention.

FIG. 2 is a flow chart illustrating task switching.

FIG. 3 is a block diagram of an interrupt and
30 task change processing circuit.

FIG. 4 is a circuit diagram of an alternate implementation of the interrupt and task change system using three latches instead of two.

FIG. 5 is a circuit diagram of a task register
35 with additional trace circuitry.

FIG. 6 is a block diagram of an implementation of the trace control system.

FIG. 7 is a timing diagram illustrating the control sequence for sample task switching.

BEST MODES FOR CARRYING OUT THE INVENTION

5 FIG. 1 schematically shows one bit slice of a task register for a computer 1. Thirty two such slices are used for each 32 bit register. Typically, a CPU would use a plurality of 32 bit registers. CPU 1 and static random access memory (SRAM) 2 can alternately be
10 connected to latch 3 and latch 4 (never to the same latch). For discussion purposes we will start with latch 3 connected to the CPU and latch 4 connected to SRAM 2. In this configuration, the CPU is running the task stored in latch 3. While the CPU 1 is running a program using
15 task latch 3, the hardware processes interrupts in the background.

 FIG. 3 schematically illustrates the interrupt processing circuitry. The Hardware interrupts 90 contains all of the possible hardware sources for
20 interrupts including external pins, and internal peripheral devices, and exception generators. In the preferred implementation, each of the major interrupt sources is assigned a fixed hardware task number. Alternatively, a flexible mapping between hardware and
25 sources and tasks is possible. When the Time-out Counter Registers 91 reach their end count, they output an interrupt, just like any hardware interrupt 90 source. The task controlled by a time-out counter register 91 will run after the time-out counter register 91 finishes,
30 and when running will automatically transfer the new time-out value from it's time-out counter register 91. The possible sources of interrupt, time-out counter registers 91, and hardware interrupts 90 combine in OR gate 99 and go to the Interrupt Enable Control 96.

35 The Interrupt Enable Control switch 96 is controlled by the task link system. The task link system uses a Link Register 81 of the running task to toggle the

LINKED latch 83 of a stored task, through task decoder 82 and AND gate 100. Writing to the Link Register 81 also disables the task specified. After the LINKED latch has been set, when the prerequisite task runs again it will re-enable the task specified in the Link Register 81. Then the task specified can run. At the end of the LINKED task running, it will disable its own interrupt enable latch 97 if its LINKED latch 83 was set, through task decoder 102 and AND gate 103 and OR gate 104. A logic diagram for this circuit is shown in the FIG 3A. Thus, this circuit will guarantee that the linked task and the prerequisite task alternate running and that the LINKED task always runs after the prerequisite task and not before. Long chains of linked tasks are possible.

Interrupt signals that make it through the Interrupt Enable Control 96 then go to their task Priority counter register 95. This register is loaded from the task via the CPU BUS, whenever it runs with the starting priority of the task. The Deadline rate latch 94 is also loaded when the task runs, and then this is loaded into the Task Deadline Priority Counter 93. Interrupts that get through the Interrupt Enable Control 96 start the Task Deadline Priority Counter 93. The Task Deadline Priority Counter 93 counts until done, then outputs a signal to the Task Priority counter 95 to increase the task priority by 1. This same signal is used to reload the Deadline Rate Latch 94 into the Task Deadline Priority Counter 93. In this fashion, the task can be set to increase its priority after it has interrupted as its deadline approaches.

Each task feeds its priority from the task priority counter register 95 to the Priority selector 19. This circuit picks the highest priority task and passes the task number and priority to the Task switch controller 20. Since tasks can spend time waiting to run in one of the latches 4 or 3 (FIG. 1), the task priority Deadline Priority Counter 93 for the waiting task is

passed through MUX 92 to the Task switch Controller 20.

In this fashion the task controller 20 can keep increasing the priority of the waiting task until it exceeds the priority of the running task and then
5 initiate a task swap.

In the event that a plurality of equal priority tasks are simultaneously presented to the priority selector 19, the priority selector 19 will determine the priority and select the tasks sequentially. The task
10 switch controller 20 uses a small state machine to sequence the task controller bus to accomplish a task switch. FIG. 2 shows a flow chart explaining the steps performed by the task switch controller 20 to accomplish a task switch.

15 At time zero, in FIG. 1, the circuit is set as follows: the CPU 1 is connected to latch 3; the task control bus A/B signal line is set to B; the latch 3 Q output goes through the multiplexer (MUX) 17 to CPU 1 input; the CPU 1 output goes through MUX 14 to latch 3
20 input D; the latch 4 outputs go through MUX 13 to SRAM 2 input; and, the SRAM 2 output goes through MUX 15 to latch 4 input D.

The first step in switching tasks is to read the task from SRAM 2 into a latch not being used by the
25 CPU 1, in this case latch 4. This is accomplished by placing the interrupting task address on the task control bus TASK TO READ address bus. CLOCK B signal then stores that state into latch 4. (Since our example is from a cold start, there is no task in latch 4. Normally there
30 will be a task there. This task can be written back into the SRAM 2 at the same time the new interrupt task is read out of SRAM 2, by asserting the TASK TO WRITE task destination addresses and then asserting the CLOCK SRAM line.) The new task is loaded into latch 4 at time T2
35 (FIG. 2). On the same clock edge or a fraction of a cycle later, the A/B control signal is changed to B, causing latch 4 to output and take input from CPU 1, and

causing latch 3 to take input and to output to SRAM 2. At this point, if the original task interrupts again, the A/B line is toggled again to switch tasks without requiring a read from SRAM 2.

5 The above example is the start-up case. Normally, once started, there is a task waiting to run, when a new interrupt comes in. Thus, the waiting task could be "flushed" (thrown away). But, in the preferred embodiment, the waiting task is always run for at least
10 one cycle while the new interrupt task is fetched from the SRAM 2 and the previous running task is saved to SRAM 2. In sequence, when latch 3 is connected to the CPU 1 and running a task, and while latch 4 is holding the next task scheduled to be run, if a new interrupt of equal or
15 higher priority is detected by the priority selector 19 (FIG. 3), then at T1 (FIG. 2) the A/B line is toggled to B, causing the task data in latch 4 to become attached to the CPU 1, while latch 3 is now connected to the SRAM 2. The new interrupt task address is then put on the TASK TO
20 READ task control bus, and the previous running task address is put on the TASK TO WRITE lines. At T2 (FIG. 2), the previous interrupt task data is clocked into the SRAM 2, the new interrupt task data is clocked into latch 3 and the CPU 1 will have clocked any data (if
25 it was programmed to) into latch 4. The A/B line is then toggled again to A to run the new interrupt task out of latch 3.

Alternate implementations of the invention are possible. FIG. 4 shows a register bit implemented with 3
30 latches instead of 2. Here 3-input MUXs 40 & 41 replace the 2-input MUXs in the 2 latch design (FIG. 1). An additional MUX 42 is added to select SRAM 2 or CPU 1. The extra latch allows 3 tasks to have latencies of just 1 cycle. However, the circuit is larger per bit and
35 somewhat slower because of the extra wiring and capacitance. There are also additional gate delays through MUXs 40 & 41 compared to the simpler MUXs 13 & 17

(FIG. 1). The optimum configuration for a given application depends on the specific characteristics of the integrated circuit manufacturing process used and the application's timing demands.

5 FIG. 5 shows an implementation of the trace function. In addition to the task switching function described above and shown in FIG. 1, MUX 50 and tri-state switch 51 have been added. Trace SRAM 2a is used as
10 either task switching SRAM 2 or as trace SRAM 2a by changing the addressing and trace/task control line to MUX 50. When CPU 1 stores into latch 3 or 4, if tracing, MUX 50 will pass data to trace SRAM 2a. FIG. 6 shows a block diagram of the trace controller. Various test
15 points are selected by MUX 60 to logical testing by test logic 61. If the test logic indicates the situation matches a user test set point, the trace controller 62
20 begins tracing. The trace controller 62 in it's simplest use, outputs sequential trace write addresses while asserting the trace signal. The trace controller can
25 also be designed to trace continuously until an event occurs, then stop tracing, or trace for a few more steps, before stopping. Readout of the trace data is accomplished by asserting the pass control on tri-state
30 51 (FIG. 5) for each bit to be connected to the trace read bit. This data can be read out a bit at a time, since trace readout is not a time critical task, or it can be grouped together for register output by using a wider bus.

30 While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation. Thus the breadth and scope of the present invention should not be limited by any of the above described exemplary embodiments, but should be
35 defined only in accordance with the following claims and their equivalents. It will be understood by those skilled in the art that various changes in form and

detail may be made therein without departing from the spirit and scope of the invention.

The terms and expressions which have been employed in the foregoing specification are used therein
5 as terms of description and not of limitation, and there is no intention, in the use of such terms and expressions, of excluding equivalents of the features shown and described or portions thereof, it being
10 recognized that the scope of the invention is defined and limited only by the claims which follow.

CLAIMS

1. In a microcomputer having a multiplicity of registers that are selectively multiplexed to communicate with a CPU, and a register set memory for storing a multiplicity of register sets and being dual addressed for reading a first register set simultaneously with writing a second register set, each said register set designated for performing a task, an improvement for permitting the rapid switching between tasks, said improvement comprising:
- a first and a second latch assembly for each bit in said register set, each said latch assembly including:
 - a latch;
 - a first multiplexer having an output connected to the input of said latch, a first input connected to an output of said CPU and a second input connected to an output of said register set memory;
 - a second multiplexer having an input connected to the output of said latch, a first output connected to an input of said CPU and a second output connected to an input of said register set memory, whereby a first register set processed by said CPU may be written into said first latch sets in the same clock cycle as a second register set is read from said register set memory and stored in said second latch sets, and a third register set is read from said second latch sets and written into said register set memory.







FIG. 1

2/7

*Fig. 2*

Set Interrupt Enable latch If: Link Register task decoder asserted & LINKED
Clear Interrupt Enable latch IEL If: Link register task decoder & not LINKED, and write
or: Running task decoder asserted & LINKED

LINKED	write	Link register	Run task	Set/Clear IEL
1	0	1	0	SET
0	1	1	0	CLEAR
1	0	0	1	CLEAR
		Everything Else		Nothing

Fig. 3A




Fig. 3




FIG. 4

5/7





FIG. 5

*FIG. 6*





FIG. 7



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/40	A1	(11) International Publication Number: WO 99/41661 (43) International Publication Date: 19 August 1999 (19.08.99)
(21) International Application Number: PCT/US99/02575 (22) International Filing Date: 5 February 1999 (05.02.99) (30) Priority Data: 09/023,333 13 February 1998 (13.02.98) US (71) Applicant (for all designated States except US): XYRON CORPORATION [US/US]; 18205 S.W. Boones Ferry Road, Portland, OR 97224 (US). (72) Inventor; and (75) Inventor/Applicant (for US only): DONOVAN, Brian [US/US]; 8615D S.W. Curry Drive, Wilsonville, OR 97070 (US). (74) Agents: SIEGEL, Timothy, E. et al.; Chernoff, Vilhauer, McClung & Stenzel, LLP, 600 Benj. Franklin Plaza, One Southwest Columbia, Portland, OR 97258-2007 (US).		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i>

(54) Title: ZERO OVERHEAD COMPUTER INTERRUPTS WITH TASK SWITCHING



(57) Abstract

The invention constitutes a unique hardware zero overhead interrupt and task change system for the reduction or elimination of interrupt latency as task change processing overhead delays in computer architectures. Without loss of time, the system performs complete task state saving and restoration between one cycle and the next without software intervention. For each Central Processing Unit (1) the invention uses one or more auxiliary latches (3, 4) wherein one latch (4) is used as the "running" latch and one of the auxiliary latches is attached to task storage memory. The invention swaps connections between alternate "running" registers and auxiliary registers while transferring other tasks to and from task storage memory (2). The invention provides a task linking system to allow the linking of tasks for mandatory sequential execution of the linked tasks. Further, the invention includes a priority "impatience" counter system to increase the relative priorities of various tasks as they approach their task deadlines.