

INTERRUPT SIGNALING FOR DIRECTED INTERRUPT VIRTUALIZATION

BACKGROUND

[0001] The present disclosure relates, in general, to interrupt processing within a computer
5 system, and in particular, to handling interrupts generated by bus connected modules in a multiprocessor computer system.

[0002] Interrupts are used to signal to a processor that an event needs the attention of the processor. For example, hardware devices, e.g., hardware devices connected with the processors via a bus, use interrupts to communicate that they require attention from an
10 operating system. In case the receiving processor is currently performing some activities, the receiving processor may in response to receiving an interrupt signal suspend its current activities, save its state, and handle the interrupt, e.g., by executing an interrupt handler. The interruption of the processors' current activities resulting from the receipt is only temporary. After having handled the interrupt, the processor may resume its suspended activities. Thus,
15 interrupts may allow for a performance improvement by eliminating unproductive waiting time of a processor in polling loops, waiting for external events.

[0003] In multiprocessor computer systems interrupt routing efficiency issues may arise. The challenge is to forward interrupt signals send by hardware devices, like, e.g., bus connected modules, to a processor of the multiple processors assigned for use by the operating in an
20 efficient way. This may be particularly challenging in case the interrupt is used to communicate with a guest operating system on a virtual machine. A hypervisor or virtual machine monitor (VMM) creates and run one or more virtual machine, i.e., guest machines. A virtual machine provides a guest operating system executed on the same with a virtual operating platform, while hiding the physical characteristics of the underlying platform. Using
25 multiple virtual machines allows to run multiple operating systems parallel. Since being executed on a virtual operating platform, the guest operating system's view of the processors may in general differ from an underlying, e.g., physical view of the processors. The guest operating system uses virtual processor ID to identify processor, which in general do not coincide with underlying logical processor IDs. The hypervisor which manages the execution

of the guest operating system defines a mapping between underlying logical processor IDs and virtual processor IDs used by the guest operating system. However, this mapping and the selection of processors scheduled for use by the guest operating system are not static, but may be changed by the hypervisor, while the guest operating system is running, without the
5 knowledge of the guest operating system.

[0004] Typically, this challenge is solved by forwarding interrupt signal using broadcasting. When using the broadcasting, the interrupt signal is successively forwarded between the multiple processors until a processor suitable for handling the interrupt signal is met. However, in case of multiple processor the probability that a processor receiving the
10 broadcasted interrupt signal first is indeed suitable for handling the interrupt signal may rather be low. Furthermore, being suitable for handling the interrupt signal not necessarily means that the respective processor is the best choice for handling the interrupt.

SUMMARY

[0005] Various embodiments provide for a method for providing an interrupt signal to a
15 guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, a computer system and a computer program product as described by the subject matter of the independent claims. Advantageous embodiments are described in the dependent claims. Embodiments of the present invention can be freely combined with each other if they are not mutually exclusive.

20 [0006] In one aspect, the invention relates to a method for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, the computer system further comprising one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device, the computer system further
25 comprising a memory operationally connected with the bus attachment device, each processor of the plurality of processors being assigned with a logical processor ID used by the bus attachment device to address the respective processor, each processor of the plurality of processors assigned for usage by the guest operating system further being assigned with an interrupt target ID used by the operating system and one or more bus connected modules to
30 address the respective processor, the memory comprising a directed interrupt signal vector per

interrupt target ID assigned to the respective interrupt target ID, each directed interrupt signal vector comprising a directed interrupt signal indicator per bus connected module assigned to the respective bus connected module, each directed interrupt signal vector indicating whether there is an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID to be handled, the method comprising: receiving by the bus attachment device an interrupt signal from one of the bus connected modules of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; selecting by the bus attachment device the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed; selecting by the bus attachment device in the selected directed interrupt signal vector the directed interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal; updating by the bus attachment device the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that there is an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID to be handled; forwarding by the bus attachment device the interrupt signal to the target processor.

[0007] In another aspect, the invention relates to a computer system for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of the computer system assigned for usage by the guest operating system, the computer system further comprising one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device, the computer system further comprising a memory operationally connected with the bus attachment device, each processor of the plurality of processors being assigned with a logical processor ID used by the bus attachment device to address the respective processor, each processor of the plurality of processors assigned for usage by the guest operating system further being assigned with an interrupt target ID used by operating system and the one or more bus connected modules to address the respective processor, the memory comprising a directed interrupt signal vector per interrupt target ID assigned to the respective interrupt target ID, each directed interrupt signal vector comprising a directed interrupt signal indicator per bus connected module assigned to the respective bus connected module, each directed

interrupt signal vector indicating whether an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID is waiting to be handled, the computer system being configured to perform a method comprising: receiving by the bus attachment device an interrupt signal from one of the bus connected modules of the bus

5 connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; selecting by the bus attachment device the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed; selecting by the bus attachment device in the selected directed interrupt signal vector the directed

10 interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal; updating the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID is waiting to be handled; forwarding by the bus attachment device the interrupt signal to the target

15 processor.

[0008] In another aspect, the invention relates to a computer program product for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, the computer system further comprising one or more bus connected modules being

20 operationally connected with the plurality of processors via a bus and a bus attachment device, the computer system further comprising a memory operationally connected with the bus attachment device, each processor of the plurality of processors being assigned with a logical processor ID used by the bus attachment device to address the respective processor, each processor of the plurality of processors assigned for usage by the guest operating system

25 further being assigned with an interrupt target ID used by the operating system and one or more bus connected modules to address the respective processor, the memory comprising a directed interrupt signal vector per interrupt target ID assigned to the respective interrupt target ID, each directed interrupt signal vector comprising a directed interrupt signal indicator per bus connected module assigned to the respective bus connected module, each directed

30 interrupt signal vector indicating whether an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID is waiting to be handled,

the computer program product comprising a computer readable non-transitory medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising: receiving by the bus attachment device an interrupt signal from one of the bus connected modules of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; selecting by the bus attachment device the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed; selecting by the bus attachment device in the selected directed interrupt signal vector the directed interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal; updating the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID is waiting to be handled; forwarding by the bus attachment device the interrupt signal to the target processor.

15 BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0009] In the following, embodiments of the invention are explained in greater detail, by way of example only, making reference to the drawings in which:

- [0010] FIG. 1 depicts a schematic diagram of an exemplary computer system,
- [0011] FIG. 2 depicts a schematic diagram of an exemplary virtualization scheme,
- 20 [0012] FIG. 3 depicts a schematic diagram of an exemplary virtualization scheme,
- [0013] FIG. 4 depicts a schematic diagram of an exemplary virtualization scheme,
- [0014] FIG. 5 depicts a schematic diagram of an exemplary computer system,
- [0015] FIG. 6 depicts a schematic diagram of exemplary vector structures,
- [0016] FIG. 7 depicts a schematic diagram of exemplary vector structures,
- 25 [0017] FIG. 8 depicts a schematic diagram of an exemplary computer system,
- [0018] FIG. 9 depicts a schematic flow diagram of an exemplary method,
- [0019] FIG. 10 depicts a schematic flow diagram of an exemplary method,

- [0020] FIG. 11 depicts a schematic flow diagram of an exemplary method,
[0021] FIG. 12 depicts a schematic diagram of an exemplary computer system,
[0022] FIG. 13 depicts a schematic flow diagram of an exemplary method,
[0023] FIG. 14 depicts a schematic flow diagram of an exemplary method,
5 [0024] FIG. 15 depicts a schematic flow diagram of an exemplary method,
[0025] FIG. 16 depicts a schematic flow diagram of an exemplary method,
[0026] FIG. 17 depicts a schematic flow diagram of an exemplary method,
[0027] FIG. 18 depicts a schematic diagram of exemplary data structures,
[0028] FIG. 19 depicts a schematic flow diagram of an exemplary method,
10 [0029] FIG. 20 depicts a schematic diagram of an exemplary computer system,
[0030] FIG. 21 depicts a schematic diagram of an exemplary computer system,
[0031] FIG. 22 depicts a schematic diagram of an exemplary computer system,
[0032] FIG. 23 depicts a schematic diagram of an exemplary computer system,
[0033] FIG. 24 depicts schematic diagrams of an exemplary units, and
15 [0034] FIG. 25 depicts a schematic diagram of an exemplary computer system.

DETAILED DESCRIPTION

- [0035] The descriptions of the various embodiments of the present invention will be presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of
20 ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.
- 25 [0036] Embodiments may have the beneficial effect of enabling the bus attachment device to address the target processor directly. Thus, the interrupt signal may be targeted by the issuing

bus connected module selecting a target processor ID to a particular processor, i.e. target processor, of a multiprocessor computer system. For example, a processor may be selected as a target processor for the interrupt signal, which has performed activities related to the interrupt before. Handling the interrupt signal by the same processor as the respective activities may result in a performance advantage, since in case the same processor is also handling the interrupt signal, all data in context with this interrupt may already be available to the processor and/or stored in a local cache enabling a fast access for the respective processor without requiring extensive cache traffic.

[0037] Thus, a broadcasting of the interrupt signal may be avoided for which there is no guaranty that the processor which will handle the interrupt in the end is best suited for this task from a performance point of view, like cache traffic minimization. Rather than presenting the interrupt signal to all processors, each processor trying to handle it and one processor wins, the interrupt signal may be provided directly to a target processor increasing the efficiency of the interrupt signal handling.

[0038] An interrupt mechanism may be implemented using directed interrupts. The bus attachment device may be enabled to directly address a target processor using a logical processor ID of the same, when forwarding an interrupt signal for handling to its target processor defined by the issuing bus connected module. Translating interrupt target IDs to logical processor IDs by the bus connected device may further ensure that from point of view of the guest operating system always the same processor is addressed, even though the mapping between interrupt target IDs and logical processor IDs or the selection of processors scheduled for use by the guest operating system may be changed by the hypervisor.

[0039] The directed interrupt signal vectors are ordered depending on the target processor ID, i.e., optimized for taking track of directed interrupts. In other words, the main order criterium are the target processor IDs rather than requestor IDs identifying the issuing bus connected modules. Each directed interrupt signal vector may comprise one or more directed interrupt signal indicators depending on the number of bus connected modules.

[0040] An ordering of interrupt signal indicators, e.g. in form of interrupt signaling bits, indicating an individual interrupt signal has been received, e.g. in form of an MSI-X message, sequentially within a contiguous area of memory, like a cache line, for an individual bus

connected module, like e.g. PCIe function, may thus be avoided. Enabling and/or disabling an interrupt signal indicator, e.g. by setting and/or resetting an interrupt signaling bit, requires the respective contiguous area of memory to be moved to one of the processors to change the respective interrupt signal indicator accordingly.

5 [0041] It may be intended that a processor handle all indicators for which it is responsible from a guest operating system perspective, i.e., in particular all indicators assigned to the respective processor. This may enable a performance advantage, since in case each processor is handling all data assigned to the same a likelihood that data required in this context is provided to the processor and/or stored in a local cache may be high enabling a fast access to
10 the respective data for the processor without requiring extensive cache traffic.

[0042] However, each processor trying to handle all indicators for which it is responsible may nevertheless lead to a high cache traffic between the processors, as each processor needs to write all cache lines for all functions. Since the indicators assigned to each individual processor may be distributed over all contiguous areas, such as cache lines.

15 [0043] The interrupt signaling indicators may be reordered in form of directed interrupt signaling vectors such that all interrupt signaling indicators assigned to the same interrupt target ID are combined in the same contiguous area of memory, e.g., cache line. Thus, a processor intending to handle indicators assigned to the respective processor, i.e., interrupt target ID, may only have to load a single contiguous area of memory. Thus, a contiguous area
20 per interrupt target ID is used rather than a contiguous area per bus connected module. Each processor may only need to scan and update a single contiguous area of memory, e.g., a cache line for all interrupt signal received from all available bus connected modules targeted to that specific processor as a target processor identified by the interrupt target ID.

[0044] According to embodiments, an offset may be applied by the hypervisor for a guest
25 operating system to align bits to different offsets.

[0045]

[0046] In accordance with an embodiment, the interrupt signal is received in form of a message signaled interrupt comprising the interrupt target ID of the target processor. Using message signaled interrupts (MSI) is a method for a bus connected module, such as a
30 Peripheral Component Interconnect (PCI) or Peripheral Component Interconnect express

(PCIe) function, to generate a central processing unit (CPU) interrupt in order to notify a guest operating system using the respective central processing unit of the occurrence of an event or the presence of some status. MSI provides an in-band method of signaling an interrupt, using special in-band messages, thereby avoiding a need for dedicated paths separated from a main data path to send such control information, such as dedicated interruptions pin on each device. MSI rather relies on exchanging special messages indicating interrupts through the main data path. When a bus connected module is configured to use MSI, the respective module requests an interrupt by performing an MSI write operation of a specified number of bytes of data to a special address. The combination of this special address, i.e., MSI address, and a unique data value, i.e. MSI data, is termed an MSI vector.

[0047] Modern PCIe standard adapters have the capability to present multiple interrupts. MSI-X for example permits a bus connected module to allocate up to 2048 interrupts. Thus, targeting individual interrupts to different processors, such as in a high-speed networking application relying on multi-processor systems, is enabled. MSI-X allows to assign a multitude of interrupts each with an individual MSI address and MSI data value.

[0048] For transmitting an interrupt signal an MSI-X message may be used. The required content of the MSI-X message may be determined using an MSI-X data table. The MSI-X data table local to the bus connected module, i.e. PCIe adapter/function, may be indexed by a number assigned to each interrupt signal, also referred to as interrupt request (IRQ). The MSI-X data table content is under control of the guest operating system and may be set to the operating system by guidance of hardware and/or firmware. A single PCIe adapter may comprise multiple PCIe functions, each of which may have an independent MSI-X data table. This may for example be the case for a single root input/output virtualization (SR-IOV) or multi-function devices.

[0049] An interrupt target ID, like e.g., a virtual processor ID, may be directly encoded as part of a message sent by the bus connected module, like e.g., an MSI-X message, comprising the interrupt signal. The message, e.g., an MSI-X message, may comprise a requestor ID, i.e. an ID of the bus connected module, the aforementioned interrupt target ID, an DIBV or AIBV index, an MSI address and MSI data. An MSI-X message may provide 64 bits for the MSI

address and 32 bits for the data. A bus connected module may request an interrupt using MSI by performing an MSI write operation of a specific MSI data value to a special MSI address.

[0050] The device table is a shared table which may be fully indexed by the requestor ID (RID) of the interrupt requestor, i.e., the bus connected module. The bus attachment device
5 remaps and posts the interrupt, i.e., the bus attachment device translates the interrupt target ID and uses the same to directly addresses the target processor.

[0051] A guest operation system may use virtual processor IDs to identify processors in a multi-processor computer system. Thus, the guest operation system's view of processors may not be not identical to a view of an underlying system using logical processor IDs. Bus
10 connected modules providing resources used by a guest operation system may use virtual processors IDs as the resources for communicating with the guest operation system, like e.g., an MSI-X data table, may be under control of the guest operating system. As an alternative to a virtual processor ID, any other ID may be defined for the bus connected module to address the processors.

15 [0052] The interrupt is presented to the guest operating system or other software executed thereon, such as other programs etc. As used herein, the term operating system includes operating system device drivers.

[0053] As used herein, the term bus connected module may comprises any type of bus connected module. According to embodiments, the module may be a hardware module, like
20 e.g. storage function, processing module, network module, cryptographic module, PCI/PCIe adapter, other type of input/output module, etc. According to other embodiments, the module may be a software module, i.e., a function, like e.g. storage function, processing function, network function, cryptographic function, PCI/PCIe function, other type of input/output function, etc. Thus, in the examples presented herein, module is used interchangeably with
25 function, e.g., PCI/PCIe function, and adapter, e.g., PCI/PCIe function, unless otherwise noted.

[0054] Embodiments may have the beneficial that an interrupt signal routing mechanism, e.g., MSI-X message routing mechanism, is provided which allows it to keep bus connected modules, e.g., PCIe adapters and functions, as well as device drivers used to operate or control
30 the bus connected modules unchanged. Furthermore, the hypervisor may be kept from

intercepting the underlying architecture for implementing communication between bus connected modules and guest operating system, e.g., PCIe MSI-X architecture. In other words, changes to the interrupt signal routing mechanism may be implemented outside the hypervisor and bus connected modules.

5

[0055] In accordance with an embodiment, interrupt signal indicators assigned to the same bus connected module are each comprising the same offset within the directed interrupt signal vector comprising the respective interrupt signal indicator. Embodiments may have the beneficial effect of providing simplifying a checking of interrupt signal indicators regarding a specific bus connected module. In accordance with an embodiment, the directed interrupt signal vectors are each being implemented as a contiguous area in memory. Embodiments may have the beneficial effect that the interrupt summary vector is provided in a form which minimum memory space and is fast and efficiently processable. The contiguous area may, e.g., be a cache line. In accordance with an embodiment, the directed interrupt signal indicators are each being implemented as a single bit. Embodiments may have the beneficial effect that the directed interrupt signal indicators are provided in a form which minimum memory space and is fast and efficiently processable.

[0056] In accordance with an embodiment, the method further comprises retrieving by the bus attachment device a copy of an interrupt table entry assigned to the received interrupt target ID from an interrupt table stored in the memory, the interrupt table entry comprising a directed interrupt signal vector address indicator indicating a memory address of the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed, the bus attachment device using the memory address of the respective directed interrupt signal vector for the selecting of the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed.

[0057] In accordance with an embodiment, the method further comprises retrieving by the bus attachment device a copy of a device table entry from a device table stored in the memory, the device table entry comprising an interrupt table address indicator indicating a memory address of the interrupt table, the bus attachment device using the memory address of the interrupt table for the retrieving of the first copy of the interrupt table entry.

[0058] In accordance with an embodiment, the device table entry further comprises a directed interrupt signal offset indicator indicating an offset of the directed interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal.

[0059] In accordance with an embodiment, the memory further comprises a directed
5 interrupt summary vector with a directed interrupt summary indicator per interrupt target ID, each directed interrupt summary indicator being assigned to an interrupt target ID indicating whether there is an interrupt signal addressed to the respective interrupt target ID to be handled, wherein the method further comprises: selecting by the bus attachment device the directed interrupt summary indicator assigned to the target processor ID to which the received
10 interrupt signal is addressed; updating by the bus attachment device the selected directed interrupt summary indicator such that the selected directed interrupt summary indicator indicating that there is an interrupt signal addressed to the respective interrupt target ID to be handled.

[0060] When an interrupt cannot be delivered directly, e.g., because the hypervisor has not
15 scheduled the target processor, the guest operating system may benefit by delivering the interrupt with the originally intended affinity, i.e. information for which processor the interrupt was intended, using broadcast. In this case the bus attachment device may set a bit specifying the target processor in the DISB after setting the DIBV and before delivering a broadcast interrupt request to the guest operating system. If a guest operating system receives
20 a broadcast interrupt request, it may therefore identify by scanning and disabling direct interrupt summary indicators in the DISB, e.g. scanning and resetting direct interrupt summary bits, which target processors have an interrupt signal pending as signaled in a DIBV. The guest operating system may thus be enabled to decide, whether the interrupt signal is handled by the current processor which received the broadcast or further forwarded to the original
25 target processor.

[0061] In accordance with an embodiment, the directed interrupt summary vector is implemented as a contiguous area in memory. Embodiments may have the beneficial effect that the directed interrupt summary vector is provided in a form which minimum memory space and is fast and efficiently processable. The contiguous area may, e.g., be a single cache
30 line. In accordance with an embodiment, the directed interrupt summary indicators are each

being implemented as a single bit. Embodiments may have the beneficial effect that the directed interrupt summary indicators are provided in a form which minimum memory space and is fast and efficiently processable.

- 5 [0062] In accordance with an embodiment, the interrupt table entry is further comprising a directed interrupt summary vector address indicator indicating a memory address of the directed interrupt summary vector, wherein the bus attachment device is using the memory address of the directed interrupt summary vector for selecting the directed interrupt summary indicator assigned to the target processor ID to which the received interrupt signal is addressed.
- 10 [0063] In accordance with an embodiment, the interrupt table entry is further comprising a directed interrupt summary offset indicator indicating an offset of the directed interrupt summary indicator assigned to the target processor ID within the directed interrupt summary vector.
- 15 [0064] In accordance with an embodiment, the method is further comprising translating by the bus attachment device the interrupt target ID of the target processor received with the interrupt signal to a logical processor ID of the target processor and using the logical processor ID of the target processor to address the target processor as the target of the interrupt signal, when forwarding the interrupt signal to the target processor.
- 20 [0065] In accordance with an embodiment, the method is further comprising retrieving by the bus attachment device a copy of an interrupt table entry assigned to the interrupt target ID from an interrupt table stored in the memory, wherein the copy of the interrupt table entry is further comprising a current mapping of the interrupt target ID to a first logical processor ID, the bus attachment device using the copy of the interrupt table entry for the translating of the interrupt target ID of the target processor received with the interrupt signal.
- 25 [0066] Embodiments may have the beneficial effect of providing an interrupt table (IRT) comprising interrupt table entries (IRTE) each entry providing a mapping of an interrupt target ID to a logical processor ID. Thus, the entries may define a unique assignment for each interrupt target ID to a logical processor ID. According to embodiments, the interrupt target ID may be provided in form of a virtual processor ID. According to embodiments, the interrupt

target ID may be any other ID used by the guest operating system to identify individual processors used.

[0067] According to embodiments, the IRT is provided in the memory for use by the bus attachment device to map interrupt target IDs to a logical processor IDs. According to
5 embodiments the IRT may be provided in a single location. An address indicator, e.g., a pointer, indicating a memory address of the IRT may be provided. The address indicator may for example be provided by an entry of a device table fetched from the memory by the bus attachment device. Embodiments may have the beneficial effect that no large mapping tables have to be stored in the bus attachment device. The interrupt table used for mapping may
10 rather be stored in the memory and accessed by the bus attachment device if necessary. Thus, the bus attachment device may only have to handle a working copy of one or more interrupt table entry for each interrupt signal to be forwarded. The number of interrupt table entry may preferably be small, e.g. one.

[0068] According to embodiments, the IRT or individual IRTE may be updated upon a
15 rescheduling of the processors. According to embodiments, the IRT may be stored in an internal section of the memory, i.e., an HSA.

[0069] In accordance with an embodiment, the copy of the device table entry further comprises a direct signaling indicator indicating whether the target processor is to be addressed directly, the direct signaling indicator indicating a direct forwarding of the interrupt
20 signal being a requirement for executing the forwarding of the interrupt signal using logical processor ID of the target processor to address the target processor directly, else the forwarding being executed using broadcasting.

[0070] Embodiments may have the beneficial effect of controlling with the direct signaling indicator whether an interrupt signal is forwarded using direct addressing or broadcasting.
25 Using the direct signaling indicator for each bus connected module an individual predefined selection may be provided whether for interrupt signals received from this bus connected module a direct addressing is to be performed or a broadcasting.

[0071] According to embodiments, a double fetch of the IRTE may be performed, in order to prevent the interrupt signal from being sent to a processor which, e.g., has been deactivated in
30 the meantime. According to embodiments, after forwarding the interrupt signal to the

processor identified by the logical processor ID resulting from the translation of the interrupt target ID using a first copy of IRTE a second copy of the same IRTE may be fetched to check whether in the meantime any changes of the IRTE took place. In case the IRTE has been updated in the meantime, there is a risk that the interrupt signal has been forwarded to a deactivated processor. Therefore, the second copy of the IRTE may be used to translate the interrupt target ID again and forward the interrupt signal to the processor identified by the logical processor ID resulting from the second translation. According to an alternative embodiment, the complete method starting with the fetching of the first copy of the IRTE may be repeated in case the second copy of the IRTE does not match the first copy. For example, a third copy of the IRTE may be fetched replacing the first copy of the IRTE or the second copy of the IRTE may replace the first copy of the IRTE and a third copy of the IRTE may be fetched to implement a double fetch scheme also for the partial repetition of the method. This scheme may be repeated until a match is achieved. According to further alternative embodiments, the interrupt signal may be forwarded using broadcasting in case the second copy of the IRTE does not match the first copy. According to embodiments, the bus attachment device participates in the memory-cache-coherence protocol and detects alternations on the IRTE, e.g., cache snooping, by the same mechanism CPUs may detect cache-line-alternations.

[0072] Embodiments may have the beneficial effect of avoiding cache purges which may have an inefficient scaling. The double fetch may be global or specific to an IRTE, i.e., the whole entry may be subject to the double fetching or restricted to specific information comprised by the respective entry.

[0073] According to embodiments, the race condition resulting from the time required for translating the interrupt target ID and forwarding the interrupt signal to the target processor, until it reaches the processor may be caught by a checking logic on CPU checking whether the receiving processor is still the correct target processor. For the check an interrupt target ID and/or logical partition ID received with the interrupt request may be compared with a current interrupt target ID and/or logical partition ID assigned to the receiving processor as a reference. In case of a match, the receiving processor directly addressed using the logical processor ID resulting from the translation using the copy of the IRTE is indeed the correct target processor. Thus, the information provided by the copy of the IRTE has been up-to-date. In case of a mismatch, the copy of the IRTE has not been up-to-date and the receiving

processor is not the target processor anymore. In case of a mismatch, the interrupt signal may, e.g., be forwarded to the target operating system using broadcasting.

[0074] According to embodiments, there may be three entities operating in parallel, i.e. the bus attachment device and the target processor processing the interrupt signal as well as the
5 hypervisor which may change the assignments between interrupt target IDs and logical processor IDs. According to embodiments, there may be no central synchronization point in a physically distributed systems besides memory which provides a virtual appearance of such a system at a latency cost. Embodiments using a double fetch scheme may have the beneficial effect of providing a method which is optimized for speed vs preventing double deliveries or
10 even misses of interrupt requests.

[0075] In view of the interrupt signal, the following actions may be performed: A1) Reading a first copy of the IRTE, A2) Sending interrupt request to a processor directly addressed, and A3) Reading a second copy of the IRTE. Simultaneously, the following sequence regarding a change of assignment between interrupt target ID and logical processor ID may occur: B1)
15 Activate additional processor with additional logical processor ID and deactivate previous processor with previous logical processor ID, and B2) Update IRTE with additional logical processor ID, i.e., replacing the previous logical processor ID by additional logical processor ID.

[0076] In certain error cases a processor, e.g., the target processor, may be reset to a
20 checkpoint and lose intermediate information. In order to regain the lost information, the processor may scan all IRTE entries for this specific processor, i.e. assigned to its logical processor ID, and deliver direct interrupt requests as indicated by the pending direct interrupt indicators, e.g., dPIA bits, present in memory which does not get impacted by a processor recovery.

[0077] The pending direct interrupt indicator comprised by the IRTE, e.g., IRTE.dPIA bit, may be used as a master copy, i.e. single point of truth, if an interrupt signal should be presented. In order to simplify processor recovery, a pending direct interrupt indicator in the processor may be used, e.g., as a shadow-copy of the IRTE.dPIA bit to keep direct interrupts pending on the processor.

[0078] In case the memory has a strictly ordered nature, only the following sequences may be possible in view of steps A1, A2, and B1: alternative 1 with $A1 \rightarrow A3 \rightarrow B1$ and alternative 2 with $A1 \rightarrow B1 \rightarrow A3$. In case of alternative 1, the first and the second copy of the IRTE may match. Thus, the interrupt signal may be forwarded to the previous processor
5 instead of the current target processor. The previous processor may see a mismatch regarding the interrupt target ID and/or logical partition ID and initiate a broadcast of the received interrupt signal. In case of alternative 2, the bus attachment device may see a mismatch between the first and the second copy of the IRTE. In response to the mismatch the bus attachment device may broadcast the interrupt signal. Due to the broadcasting, the interrupt
10 signal may be received by the additional processor which sees a hit and directly processes the received interrupt request. Embodiments may have the beneficial effect of closing the timing window through an over-initiative-approach.

[0079] In accordance with an embodiment, the copy of the interrupt table entry is further comprising a copy of a running indicator indicating whether the target processor identified by
15 the interrupt target ID is scheduled for usage by the guest operating system, the target processor being scheduled for usage by the guest operating system being a further requirement for executing the forwarding of the interrupt signal using logical processor ID of the target processor to address the target processor directly, else the interrupt signal to the first operating system for handling using broadcasting.

20 [0080] Embodiments may have the beneficial effect of preventing of interrupts targeted at processor not running, i.e., not scheduled for usage by the guest operating system. Embodiments may have the beneficial effect of supporting the hypervisor to reschedule processors.

[0081] A running indicator indicating whether the target processor identified by the interrupt
25 target ID received with an interrupt signal is scheduled for usage by the guest operating system. The running indicator may for example be implemented in form of a running bit, i.e., a single bit indicating whether a processor assigned with the respective bit is running, i.e., is scheduled for usage by the guest operating system. Thus, an enabled running bit may tell the bus attachment device that the target processor is currently scheduled, while a disabled
30 running bit may tell the bus attachment device that the target processor is currently not

scheduled. In case the target processor is not running, the bus attachment device may send a fallback broadcast interrupt request right way, without trying to address one of the processors directly.

[0082] In accordance with an embodiment, the copy of the interrupt table entry further
5 comprises an interrupt blocking indicator indicating whether the target processor identified by the interrupt target ID is currently blocked from receiving interrupt signals, the target processor being unblocked being a further requirement for executing the forwarding of the interrupt signal using logical processor ID of the target processor to address the target processor directly, else the interrupt signal to the first operating system for handling using
10 broadcasting.

[0083] According to embodiments a direct interrupt blocking indicator is introduced in the interrupt entries of an interrupt table in memory. The direct interrupt blocking indicator may be implemented in form of a single bit, i.e., a dIBPIA bit.

[0084] According to embodiments, the IRTE is fetched from the memory, a running
15 indicator is checked in order to determine whether the target processor is scheduled. In case the target processor is scheduled a direct interrupt blocking indicator is enabled, in order to block the target processor from receiving further interrupt signal, while handling the current interrupt signal. Further interrupt signal may otherwise interfere with the handling of the current interrupt signal. In order to ensure that the target processor has not been rescheduled in
20 the meanwhile, the IRTE is re-fetch and the current running indicator checked again in order to determine whether the target processor is still scheduled. In case the target processor is still scheduled, the interrupt signal may be forwarded to the target processor using a logical processor ID of the target processor to address the target processor directly. In addition, it may be checked whether the logical processor ID of the target processor provided by the IRTE for
25 the received interrupt target ID is still the same.

[0085] In accordance with an embodiment, the device table entry is further comprising a logical partition ID identifying a logical partition to which the guest operating system is assigned, wherein the forwarding of the interrupt signal by the bus attachment device is further comprising forwarding with the interrupt signal the logical partition ID. Embodiments may

have the beneficial effect of enabling the receiving processor to check to which target processor an interrupt signal is addressed.

[0086] In accordance with an embodiment, the method further comprises retrieving by the bus attachment device an interrupt subclass ID identifying an interrupt subclass to which the received interrupt signal is assigned, wherein the forwarding of the interrupt signal by the bus attachment device further comprises forwarding with the interrupt signal the interrupt subclass ID.

[0087] According to embodiments, the instructions provided on the computer readable non-transitory medium for execution by a processing circuit are configured for performing any of the embodiments of the method for providing an interrupt signal to a guest operating system as described herein.

[0088] According to embodiments, the computer system further is configured to execute any of the embodiments of the method for providing an interrupt signal to a guest operating system as described herein.

15 [0089]

[0090] Fig. 1 depicts an exemplary computer system 100 for providing an interrupt signal to a guest operating system. The computer system 100 comprises a plurality of processors 130, which are used to execute a guest operating system. The computer system 100 furthermore comprises a memory 140, also referred to as storage memory or main memory. Memory 140 may provide memory space, i.e. memory sections, assigned for use by hardware, firmware, and software components comprised by computer system 100. Memory 140 may be used by hardware and firmware of computer system 100 as well as by software, e.g., hypervisors, host/guest operating systems, application programs etc. One or more bus connected modules 120 are operationally connected with the plurality of processors 130 as well as memory 140 via a bus 102 and a bus attachment device 110. The bus attachment device 110 manages the communication between bus connected modules 120 on one hand and processors 130 as well as memory 140 on the other hand. Bus connected modules 120 may be connected with the bus 102 directly or via one or more intermediate components, like e.g. a switch 104.

[0091] The bus connected modules 120 may, e.g., be provided in form of Peripheral Component Interconnect express (PCIe) modules, also referred to as PCIe adapter or PCIe

functions provided by an PCIe adapter. A PCIe function 120 may issue a request which is send to a bus attachment device 110, e.g., a PCI Host Bridge (PHB) also referred to a PCI Bridge Unit (PBU). The bus attachment device 110 receives the requests from the bus connected modules 120. The requests may for example comprise an input/output address that is used to perform a direct memory access (DMA) to memory 140 by the bus attachment device 110 or an input/output address that that indicates an interrupt signal, e.g., a message signaled interrupt (MSI).

[0092] FIG. 2 depicts an exemplary virtual machine support provided by the computer system 100. The computer system 100 may comprise one or more virtual machines 202 as well as at least one hypervisor 200. A virtual machine support may provide an ability to operate large numbers of virtual machines, each capable of executing a guest operating system 204, such as z/Linux. Each virtual machine 201 may be capable of functioning as a separate system. Thus, each virtual machine may be independently reset, execute a guest operating system, and run different programs, like application programs. An operating system or application program running in a virtual machine appears to have access to the full and complete computer system. However, in reality only a portion of the available resources of the computers system may be available for usage by the respective operating system or application program.

[0093] The virtual machines may use a V=V model, in which the memory assigned to a virtual machine is backed by virtual memory, instead of real memory. Thus, each virtual machine has a virtual linear memory space. The physical resources are owned by hypervisor 200, such as a VM hypervisor, and the shared physical resources are dispatched by the hypervisor to the guest operating systems, as needed, to meet their processing demands. The V=V virtual machine model assumes that the interactions between the guest operating systems and the physical shared machine resources are controlled by the VM hypervisor, since the large number of guests may preclude the hypervisor from simply partitioning and assigning the hardware resources to the configured guest.

[0094] The processors 120 are assignable to virtual machines 202 by the hypervisor 200. A virtual machine 202 may e.g. be assigned with one or more logical processors. Each of the logical processors may represents all or a share of the physical processors 120 that may be

dynamically allocated to the virtual machine 202 by the hypervisor 200. Virtual machines 202 are managed by a hypervisor 200. The hypervisor 200 may for example be implemented in firmware running on the processors 120 or may be a part of an operating system executed on the computer system 100. The hypervisor 200 may for example be a VM hypervisor, such as z/VM® offered by International Business Machines Corporation, Armonk, N.Y.

[0095] FIG. 3 depicts an exemplary multilevel virtual machine support provided by the computer system 100. In addition to the first level virtualization of FIG.2, a second level virtualization is provided with a second hypervisor 210 being executed on one of the first level guest operating systems which acts as a host operating system for the second hypervisor 210.

10 The second hypervisor 210 may manage one or more second level virtual machines 212, each capable of executing a second level guest operating system 212.

[0096] FIG. 4 depicts an exemplary pattern illustrating the use of different types of IDs to identify processors at different hierarchical levels of the computer system 100. An underlying firmware 220 may provide logical processor IDs lCPU 222 to identify the processors 130 of the computer system 100. A first level hypervisor 200 uses the logical processor IDs lCPU 222 to communicate with the processors 130. The first level hypervisor may provide first virtual processor IDs vCPU 224 for use by guest operating systems 204 or second level hypervisors 219 execute on virtual machines managed by the first level hypervisor 200. The hypervisor 200 may group the first virtual processor IDs vCPU 224 to provide logical partitions, also referred to as zones, for the guest operating systems 204 and/or hypervisors 210. First virtual processor IDs vCPU 224 are mapped to the logical processor IDs lCPU 222 by the first level hypervisor 200. On or more of the first virtual processor IDs vCPU 224 provided by the first level hypervisor 200 may be assigned to each guest operating systems 204 or hypervisors 210 executed using the first level hypervisor 200. The second level hypervisor 210 executed on the first level hypervisor 200 may provide one or more virtual machine executing software, like e.g. further guest operating systems 214. For this purpose, the second level hypervisor manages second virtual processor IDs vCPU 226 for use by the second level guest operating system 214 executed on virtual machines of the first level hypervisor 200. The second virtual processor IDs vCPU 226 are mapped to the first virtual processor IDs vCPU 224 by the second level hypervisor 200.

[0097] A bus connected module 120 addressing processors 130 used by a first/second level guest operating system 204 may use a target processor ID in form of the first/second virtual processor IDs vCPU 224, 226 or alternative IDs derived from the first/second virtual processor IDs vCPU 224, 226.

5 [0098] FIG. 5 depicts a simplified schematic setup of the computer system 100 illustrating the main participants in the method for providing interrupt signals to a guest operating system executed on the computer system 100. The simplified setup for the purpose of illustration comprises a bus connected module (BCM) 120 which sends an interrupt signal to a guest operating system executed on one or more of the processors (CPUs) 130. The interrupt signal
10 is sent to the bus attachment device 110 with an interrupt target ID (IT_ID) identifying one of the processors 130 as a target processor. The bus attachment device 110 is an intermediate device managing the communication between the bus connected module 120 and the processors 130 as well as a memory 140 of the computer system 100. The bus attachment device 110 receives the interrupt signal and identifies using the interrupt target ID a logical
15 processor ID of the target processor in order to addressing the respective target processor directly. A directed forwarding to the target processor may improve the efficiency of the data processing, e.g., by reducing cache traffic.

[0099] FIG. 6 depicts a schematic structure of a DISB 160 and a plurality of DIBVs 162. The DISB 160 may be provided in form of a contiguous section of memory, e.g., a cache line, comprising an entry 161, e.g. bit, per interrupt target ID. Each entry indicating whether there is
20 an interrupt request (IRQ) to be handled by the respective processor identified by the interrupt target ID. For each interrupt target ID, i.e., entry of the DISB 160, a DIBV 162 is provided. Each DIBV 162 being assigned to a specific interrupt target ID and comprising one or more entries 163 per bus connected module MN A, MN B. The DIBVs 162 may each be provided in
25 form of a contiguous section of memory, e.g. a cache line, comprising the entries 163 assigned to the same interrupt target ID. Entries of different bus connected modules may be order using different offsets DIBVO per bus connected module.

[0100] FIG. 7 depicts an exemplary DIBV 162. An offset (DIBVO) within a DIBV assigned to the interrupt target ID identifies the start of a section or entry of the vector assigned to a
30 specific bus connected module. The interrupt signal, e.g. MSI-X message, may provide an

DIBV-Idx which is added to the DIBVO to identify a specific entry of the vector assigned to the bus connected module to identify a bit 163 assigned to the respective bus connected module and indicating, if set, that there is an interrupt signal address to the respective interrupt target ID from the bus connected module identified by the bit 163. A directed number of interrupts (NOI) defines a maximum number of bits in the DIBV reserved for the respective bus connected module.

[0101] FIG. 8 depicts the computer system 100 of FIG. 5. The bus attachment device 110 is configured to perform status updates of a status of the bus connected module 120 in a module specific area (MSA) 149 of memory 140. Such a status update may be executed in response to receiving a direct memory access (DMA) write from the bus connected module specifying the status update to be written into memory 140.

[0102] The memory further more comprises a device table (DT) 144 with a device table entry (DTE) 146 for each bus connected module 120. Upon receive of an interrupt signal, e.g., an MSI-X write message with an interrupt target ID identifying a target processor for the interrupt request and a requestor ID identifying an origin of the interrupt request in form of the bus connected module 120, the bus attachment device 110 fetches the DTE 146 assigned to the requesting bus connected module 120. The DTE 146 may indicate whether a directed addressing of the target processor is enabled for the requesting bus connected module 120, e.g., using a DIRQ bit. The bus attachment device updates an entry of a directed interrupt signal vector (DIBV) 162 as well as a directed interrupt summary vector (DISB) 160 in order to keep track for which of the processors 130 an interrupt signal has been received. The DISB 160 may comprise one entry per interrupt target ID, indicating whether there is an interrupt signal to be handled for this processor 130 from any bus connected module 120. Each DIBV 162 is assigned to one of the interrupt target IDs, i.e. processors 130, and may comprise one or more entries. Each entry being assigned to one of the bus connected modules 120. Thus, the DIBV indicates from which bus connected modules there are interrupt signals for a specific processor to 130 be handled. This may have the advantage that in order to checks whether there is any interrupt signal or from which bus connected module 120 there are interrupt signals for a specific processor to be handled. Only a signal entry, e.g., bit, or only a signal vector, e.g. bit vector, has to be read from the memory 140.

[0103] The bus attachment device 110 uses a mapping table 112 provided on the bus connected module 110 to translate the interrupt target ID (IT_ID) to a logical processor ID (ICPU) and forwarding the received interrupt signal to the target processor using the logical processor ID to address the target processor directly. Each processor comprises firmware, e.g. millicode 132, to receive and process direct interrupt signals. Firmware may further comprise e.g. microcode and/or macrocode of the processors 130. It may comprise hardware-level instructions and/or data structures used in implementation of higher-level machine code. According to embodiments, it may comprise proprietary code that may be delivered as microcode that includes trusted software or microcode specific to the underlying hardware and controls operating system access to the system hardware.

[0104] The firmware of the processors 130 comprises a check logic 134 to check whether the receiving processor is identical with the target processor according to the interrupt target ID which is forwarded by the bus attachment device 110 to the receiving processor 130. In case the receiving processor 130 is not the target processor, i.e., in case of a mismatch of the received interrupt target ID and a reference interrupt target ID of the receiving processor 130, the interrupt signal is broadcasted to the logical partition in order to find a processor for handling the interrupt signal.

[0105] FIG. 9 is a flowchart of an exemplary method for performing a status update of a bus connected module 120 via the bus attachment device 110 using a DMA write request. In step 300, a bus connected module may decide to update its status and trigger an interrupt, e.g., in order to indicate that a signal completion. In step 310, the bus connected module initiates a direct memory access (DMA) write via the bus attachment device to a section of the memory assigned to a host running on the computer system, i.e. host memory, in order to update the status of the bus connected module. A DMA is the hardware mechanism that allows peripheral components of a computer system to transfer their I/O data directly to and from main memory without a need to involve system processors. In order to perform a DMA, the bus connected module sends a DMA write request, e.g., in form of an MSI-X message, to the bus attachment device. In case of PCIe, the bus connected module may, e.g., refer to a PCIe function provided on a PCIe adapter. In step 320, the bus connected module receives the DMA write request with the status update of the bus connected module and updates the memory using the

received update. The update may be executed in an area of the host memory reserved for the respective bus connected module.

[0106] FIG. 10 is a flowchart of an exemplary method for providing an interrupt signal to a guest operating system using the computer system 100 of FIG. 8. In step 330, the bus
5 attachment device receives an interrupt signal, e.g., in form of an MSI-X write message, sent by the bus connected module. This transmission of the interrupt signal may be executed according to the specifications of the PCI architecture. The MSI-X write message comprises an interrupt target ID identifying a target processor of the interrupt. The interrupt target ID may, e.g., be a virtual processor ID used by the guest operating system to identify processors
10 of the multiprocessor computer system. According to embodiments, the interrupt target ID may be any other ID agreed upon by the guest operating system and the bus connected module in order to be able to identify processors. Such another ID may for example be a result of a mapping of an virtual processor ID: In addition the MSI-X write message may further comprise an interrupt requestor ID (RID), i.e. the ID of PCIe function issuing the interrupt
15 request, a vector index defining an offset of a vector entry within a vector, an MSI address, e.g., 64 bit address, as well as MSI data, e.g., 32 bit data. The MSI address and MSI data may indicate that the respective write message is in fact an interrupt request in the form of an MSI message.

[0107] In step 340, the bus attachment device fetches a copy of an entry of a device table
20 stored in the memory. The device table entry (DTE) provides address indicators of one or more vectors or vector entries to be updated in order to indicate that an interrupt signal has been received for the target processor. An address indicator for a vector entry may, e.g., comprise an address of the vector in the memory as well as an offset within the vector. Furthermore, the DTE may provide a direct signaling indicator indicating, whether the target
25 processor is to be addressed directly by the bus attachment device using the interrupt target ID provided with the interrupt signal. Furthermore, the DTE may provide a logical partition ID, also referred to as zone ID, and an interrupt subclass ID. The respective copy of the device table entry may be fetched from a cache or from memory.

[0108] In step 350, the bus attachment device updates the vectors specified in the DTE. In
30 step 360, the bus attachment device checks the direct signaling indicator provided with the

interrupt signal. In case the direct signaling indicator indicates no direct signaling, the bus attachment device forwards the interrupt signal by broadcasting using the zone identifier and the interrupt subclass identifier, in order to provide the interrupt signal to a processor used by the guest operating system.

5 [0109] In case the direct signaling indicator indicates no direct signaling, in step 370, the interrupt signal is forwarded to the processor via broadcasting. The broadcast message comprising the zone ID and/or the interrupt subclass ID. When received by a processor, where the interrupt request is enabled for the zone, a status bit is set atomically, e.g., according to a nest communication protocol. Furthermore, the firmware, e.g., millicode, on this processor
10 interrupts its activities, e.g., program execution, and switches to execute an interrupt handler of the guest operating system.

[0110] In case the direct signaling indicator indicates direct signaling, in step 380, the bus attachment device translates the interrupt target ID provided with the interrupt signal to a logical processor ID of a processor assigned for usage by the guest operating system. For the
15 translation the bus attachment device may use a mapping table comprised by the bus attachment device. The bus attachment device may comprise a mapping table or sub-table per zone, i.e. logical partition.

[0111] In step 390, the bus attachment device forwards the interrupt signal to the target processor using the logical processor ID to address the respective processor directly, i.e.,
20 sending a direct message. The direct message comprises the interrupt target ID. The direct message may further comprise the zone ID and/or interrupt subclass ID. The receiving processor comprises an interrupt target ID checking logic. The checking logic may furthermore take a logical partition ID into account, in case the interrupt target ID is unique only per logical partition.

25 [0112] In step 392, the checking logic checks, whether the received interrupt target ID and/or logical partition ID match an interrupt target ID and/or logical partition currently assigned to the receiving processor and accessible for the checking logic. In case of a mismatch, the receiving firmware in step 394 initiates a broadcast and broadcast the received interrupt request to the remaining processors using the logical partition ID and/or an interrupt subclass
30 ID to identify a valid target processor for handling the interrupt. In case of a positive match,

the receiving firmware, e.g., millicode, of the target processor accepts the directly addressed interrupt for presentation to the guest operating system in step 396. In response, the firmware may interrupt its activities, e.g., program execution, and switches to execute an interrupt handler of the guest operating system. The interrupt may be presented to the guest operating system with a direct signaling indication.

[0113] FIG. 11 is an additional flowchart further illustrating the method of FIG. 8. In step 400 the interrupt message is sent to the bus attachment device. In step 402, the interrupt message is received. In step 404, it is checked, whether the DTE assigned to the interrupt requestor, i.e., the bus connected module, is cached in a local cache operationally connected with the bus attachment device. In case the DTE is not cached, in step 406 the respective DTE is fetched from memory by the bus attachment device. In step 408, vector address indicator provided by the DTE are used to set the vector bits in the memory. In step 410, it is checked using the direct signaling indicator provided by the DTE, whether the target processor is to be addressed directly by the bus attachment device using the interrupt target ID provided with the interrupt signal. In case the target processor is not to be targeted directly, the method continues with broadcasting the interrupt request to the processors in step 412. In case the target processor is to be targeted directly, the method continues with translating the interrupt target ID to a logical processor ID in step 414 and sending a message forwarding the interrupt signal to the target processor in step 416. The logical processor ID is used to address the target processor directly. The message comprises the interrupt target ID, a logical partition ID and an interrupt subclass ID. In step 418, the processor receives the message. In step 419, the processor checks, whether the interrupt target ID and/or logical partition ID match the current interrupt target ID and/or logical partition ID provided as references for the check. In case of a match, the processor presents the interrupt request to the guest operating system in step 420. In case of a mismatch, the processor rather broadcast the interrupt request to the other processors in step 422. Then, the processor continues its activities until the next interrupt message is received.

[0114] FIG. 12 depicts a further embodiment of the computer system 100 of FIG. 5. Instead of using a mapping table 112 stored on the bus attachment device 110, the bus attachment device 110 rather fetches a copy 114 of an entry (IRTE) 152 of an interrupt table (IRT) 150 stored in memory 140. The copy may either be fetched from a local cache or from memory

140 using an address of the interrupt table 150 (IRT@) provided by the copy of the DTE 146. The IRTE 152 provides a mapping of the interrupt target ID to a logical processor ID which is used by the bus attachment device 110 to address the target processor directly in case of a directed interrupt forwarding. The IRTE 152 may in addition provide a running indicator 154 indicating whether the target processor identified by the interrupt target ID is scheduled, i.e. running, at all and/or a blocking indicator 146 indicating whether the target processor is currently blocked from receiving interrupt signals. In case the target processor is not scheduled or temporarily block, a broadcast may be initiated in order enable a timely interrupt handling.

[0115] FIG. 13 is a flowchart of an exemplary method for providing an interrupt signal to a guest operating system using the computer system 100 of FIG. 100. The method of FIG. 11 continues after step 340 of FIG. 8 with step 342. In step 342, the bus attachment fetches a copy of an IRTE form the memory using the interrupt target ID received with the interrupt signal as well as an address indicator provided by the DTE indicating a memory of the IRT. In step 350, the bus attachment device updates the vectors specified in the DTE.

[0116] In step 360, the bus attachment device checks the direct signaling indicator provided with the interrupt signal. In case the direct signaling indicator indicates no direct signaling, the bus attachment device forwards the interrupt signal in step 370 by broadcasting using the zone identifier and the interrupt subclass identifier, in order to provide the interrupt signal to a processor used by the guest operating system. In case the direct signaling indicator indicates direct signaling, the bus attachment device further checks in step 362, whether a running indicator comprised by the copy of the IRTE indicates that the target processor identified by the interrupt target ID is running.

[0117] In case the target processor is not running, the bus attachment device in step 364 sends a broadcast interrupt as a fallback using, e.g., a logical partition ID and/or an interrupt subclass ID to identify processor suitable for handling the interrupt. In case no suitable processor matching the logical partition ID and/or the interrupt subclass ID is found, the hypervisor, i.e. a processor assigned for usage by the hypervisor, may receive the interrupt request instead of a processor assigned to the guest operation system. The hypervisor may decide on broadcasting the interrupt request again if one or more processors assigned to the guest operation system get scheduled. On an entry of a processor assigned to the operating

system, the hypervisor may check for direct interrupt pending indicators, e.g., dPIA bits, to be presented to the entering processor. According to embodiments, the hypervisor may, e.g., selectively re-schedule, i.e., wake up, the target processor.

[0118] In case the target processor is running, in step 366 it is checked whether a direct
5 interrupt blocking indicator is enabled, e.g., a dIBPIA bit. An enabled direct interrupt blocking indicator indicates that an interrupt delivery is currently not desired by the guest operating system interrupt handler. Therefore, in case the direct interrupt blocking indicator is enabled, in step 368 interrupt signal may be broadcasted.

[0119] If the direct interrupt blocking indicator is disabled indicating that the target
10 processor is currently unblocked, the delivery of the current interrupt signal is continued in step 380 with translating the received interrupt target ID in order to forward the interrupt to the target processor directly using the logical processor ID provided by the IRTE for the received interrupt target ID.

[0120] In step 380, the bus attachment device translates the interrupt target ID provided with
15 the interrupt signal to a logical processor ID of a processor assigned for usage by the guest operating system. For the translation the bus attachment device may use a mapping table comprised by the bus attachment device. The bus attachment device may comprise a mapping table or sub-table per zone, i.e. logical partition.

[0121] In step 390, the bus attachment device forwards the interrupt signal to the target
20 processor using the logical processor ID to address the respective processor directly, i.e., sending a direct message. The direct message comprises the interrupt target IS. The direct message may further comprise the zone ID and/or interrupt subclass ID. The receiving processor comprises an interrupt target ID checking logic. The checking logic may furthermore take a logical partition ID into account, in case the interrupt target ID is unique
25 only per logical partition.

[0122] In step 392, the checking logic checks, whether the received interrupt target ID and/or
logical partition ID match an interrupt target ID and/or logical partition currently assigned to the receiving processor and accessible for the checking logic. In case of a mismatch, the receiving firmware in step 394 initiates a broadcast and broadcast the received interrupt
30 request to the remaining processors using the logical partition ID and/or an interrupt subclass

ID to identify a valid target processor for handling the interrupt. In case of a positive match, the receiving firmware, e.g., millicode, of the target processor accepts the directly addressed interrupt for presentation to the guest operating system in step 396. In response, the firmware may interrupt its activities, e.g., program execution, and switches to execute an interrupt handler of the guest operating system. The interrupt may be presented to the guest operating system with a direct signaling indication.

[0123] FIG. 14 is an additional flowchart further illustrating the method of FIG. 11. The method shown in FIG. 12 continues after step 410 of FIG. 9, in case the target processor is to be targeted directly. The method continues with fetching a copy of the IRTE assigned to the received interrupt target ID from memory in step 413. In step 413a, it is check whether a running indicator comprised by the IRTE is enabled. In case the running indicator is disabled, the interrupt signal may be forwarded by the bus attachment device using broadcasting in step 413b. In case the running indicator is enabled, the bus attachment device continues in step 413c checking whether the directed blocking indicator is enabled. In case the directed blocking indicator is not enabled, the bus attachment device continues in step 414 with translating the interrupt target ID to a logical processor ID using the fetched copy of the IRTE. Else the interrupt signal may be suppressed in step 413d.

[0124] FIG. 15 depicts a method for performing an exemplary double fetch scheme, in order to ensure that the IRTE used is up-to-date. In step 500, the interrupt signal, e.g., an MSI-X message, is sent from the bus connected module 120, e.g., a PCIe adapter or a PCIe function on a PCIe adapter, to the bus attachment device 110, e.g., a PCIe host bridge (PHB). In step 502, the bus attachment device 110 requests from memory 140 a first copy of an IRTE assigned to an interrupt target ID provided with the interrupt signal. In step 504, the memory 140 sends a copy of the IRTE in response to the request. The point of time of the sending of the copy of the IRTE marks the last point of time at which the IRTE was up-to-date for sure. At that point of time a time windows starts during which the IRTE may be updated and the data provided by the first copy of the IRTE may become outdated. The time window ends with the interrupt being handled by the target processor 130. From that point of time on any changes of the IRTE are of no influence for the handling of the received interrupt signal anymore. In step 506, the bus attachment unit 110 sends a request to the IRTE to enable a directed pending interrupt indicator, e.g., set a directed pending interrupt array (dPIA) bit. The

enabled directed pending interrupt indicator indicating that a directed interrupt is pending for the interrupt target ID. In step 508, the setting of the directed pending interrupt indicator is confirmed by the memory 140. In step 510, the interrupt signal is forwarded in form of a directed interrupt request using direct addressing to the target processor 130 identified by the logical processor ID resulting from translating the interrupt target ID using the IRTE. With the target processor 130 receiving the directed interrupt request, the time window is closed. In step 512, upon closure of the time window, a second copy of the IRTE is read by the bus attachment device 110 from the IRTE provided in the memory 140. In step 514, upon receipt of the requested second copy of the IRTE, the bus attachment device 110 checks, whether the second copy of the IRTE matches the first copy of the IRTE, i.e., whether the IRTE, in particular the mapping of the interrupt target ID has changed. In case of a match, the method ends with resetting the directed pending interrupt indicator in the IRTE by the target processor 130 after having presented the interrupt request to the guest operating system and processing the request. In case of a mismatch, the method may continue with step 502. Alternatively, the method may continue with a broadcasting of the received interrupt signal by the bus attachment device 110.

[0125] FIG. 16 depicts a further method for performing a double fetch of an IRTE in order to ensure that information provided by the IRTE is up-to-date. In step 600, the interrupt signal, e.g., an MSI-X message is sent from the bus connected module 120, e.g., a PCIe adapter or a PCIe function on a PCIe adapter, to the bus attachment device 110, e.g., a PCIe host bridge (PHB). In step 602, the bus attachment device 110 requests from memory 140 a copy of an IRTE assigned to an interrupt target ID provided with the interrupt signal. In step 604, the memory 140 sends a first copy of the IRTE in response to the request. The first copy comprises a running indicator, e.g., a running bit $R = 1$, indicating that the target processor is scheduled, a directed interrupt blocking indicator, e.g., a directed blocking bit $dIBPIA = 0$, indicating that the target processor is currently not blocked from receiving interrupt signals, as well as a logical processor ID $ICPU$. The logical processor ID $ICPU$ is used by the bus attachment device 110 to address the target processor 130 directly. Since the running indicator indicates that the target processor 130 is running, the bus attachment device 110 enables the directed interrupt pending indicator, e.g., setting $dPIA = 1$, and blocks the target processor from receiving further interrupts, e.g., setting $dIBPIA = 1$, in the IRTE in step 606. In order to

check that meanwhile the content of IRTE has not been changed, e.g., the target processor 130 been deactivated, the critical time window is closed by requesting a re-read of the IRTE in step 608. In step 610, the memory 140 sends a second current copy of the IRTE in response to the request. The second copy comprises the running indicator, e.g., running bit $R = 1$,
5 indicating that the target processor 130 is still scheduled, the directed interrupt blocking indicator enabled by the bus attachment device, as well as a logical processor ID ICPU identical with the ICPU provided by the first copy of the IRTE. Since the running indicator as well as the ICPU did not change, the method continues in step 612 with sending an interrupt request directly addressed to the target processor 130 using ICPU. The target processor 130
10 presents the interrupt to the guest operating system and processes the interrupt. When handling the interrupt is finished, the target processor 130 disables the directed interrupt pending indicator, e.g., resetting $dPIA = 0$, as well as the directed interrupt blocking indicator, e.g., resetting $dIBPIA = 0$.

[0126] FIG. 17 depicts an alternative flowchart of the method of FIG. 14 illustrating a case
15 in which information comprised by the IRTE changes meanwhile. In step 600, the interrupt signal, e.g., an MSI-X message is sent from the bus connected module 120, e.g., a PCIe adapter or a PCIe function on a PCIe adapter, to the bus attachment device 110, e.g., a PCIe host bridge (PHB). In step 602, the bus attachment device 110 requests from memory 140 a copy of an IRTE assigned to an interrupt target ID provided with the interrupt signal. In step
20 604, the memory 140 sends a first copy of the IRTE in response to the request. The first copy comprises a running indicator, e.g., a running bit $R = 1$, indicating that the target processor 130 is scheduled as well as a logical processor ID ICPU. The logical processor ID ICPU is used by the bus attachment device 110 to address the target processor 130 directly. Since the running indicator indicates that the target processor 130 is running, the bus attachment device
25 110 enables the directed interrupt pending indicator, e.g., setting $dPIA = 1$, and blocks the target processor from receiving further interrupts, e.g., setting $dIBPIA = 1$, in the IRTE in step 606. In order to check that meanwhile the content of IRTE has not been changed, e.g., the target processor 130 been deactivated, the critical time window is closed by requesting a re-read of the IRTE in step 608. In step 610, the memory 140 sends a second current copy of the
30 IRTE in response to the request. In this example, the target processor 130 has meanwhile been deactivated for the guest operating system. Therefore, the second copy comprises a running

indicator, e.g., running bit R = 0, indicating that the target processor 130 is not scheduled any more. The logical processor ID ICPU may or may not be identical with the ICPU provided by the first copy of the IRTE. The directed interrupt blocking indicator is still enabled by the bus attachment device. Since the running indicator and/or the ICPU did change, the method
 5 continues in step 612 with sending an interrupt request to the processors using broadcasting.

[0127] FIG. 18 depicts an exemplary DTE 146 comprising a memory address IRT@ of the IRT, logical partition ID (zone) and an offset (DIBVO) within a DIBV assigned to the interrupt target ID. The DIBVO identifies the start of a section or entry of the vector assigned to a specific bus connected module. The interrupt signal, e.g. MSI-X message, may provide an
 10 DIBV-Idx which is added to the DIBVO to identify a specific entry of the vector assigned to the bus connected module. Furthermore, a directed number of interrupts (NOI) is provided defining a maximum number of bits in the DIBV reserved for the respective bus connected module. Further details of the DIBV are illustrated in FIG. 19A. In case of an AIBV, the DTE may provide corresponding AIBV specific parameters as illustrated in FIG. 19B.

15 [0128] Furthermore, an exemplary IRTE 152 is depicted. The IRTE 152 may comprise a logical partition ID (zone), an interrupt subclass ID (DISC), a memory address DISB@ of the DISB, an offset DISBO within the DISB as well as a memory address DIBV of a DIBV assigned to the interrupt target ID of the target processor.

[0129] FIGS. 19A and 19B illustrate an exemplary method for providing an interrupt signal to a guest operating system. In step 704, a bus connected module (BCM) like, e.g. a virtual
 20 function on a PCI adapter, i.e. PCI-Adapter (VF), sends an interrupt signal. The interrupt signal may, e.g., be sent in form of an MSI-X message MSI-X(VF, vCPU, DIBV-Idx), comprising an identifier of the virtual function VF, an interrupt target ID, e.g., in form of a virtual processor ID vCPU, and an offset within a directed interrupt signal vector, e.g. DIBV-
 25 Idx, identifying an entry, e.g. bit, comprised by the vector. In step 706, the bus attachment device (BAD), e.g., a PCI Host Bridge (PHB) also referred to a PCI Bridge Unit (PBU), receives the interrupt signal.

[0130] In step 708, PBU reads an entry of a device table (DT) assigned to VF. Entries of DT stored in a hardware system area (HSA) of the memory are illustrated as rows of the table. The
 30 entry of DT may comprise an address of an interrupt table (IRT@) as well as a directed

- signaling bit (S) indicating, whether directed signaling is to be performed. The PBU uses IRT@ to fetch an entry of IRT assigned to vCPU from HSA comprising a running bit (R) indicating whether vCPU is running, a directed interrupt blocking bit (dIBPIA) indicating whether vCPU is blocked from receiving interrupts, as well as a directed interrupt pending bit (dPIA) indicating whether an interrupt directed to vCPU is pending. At an earlier point in time in step 700, a start interpretive execution instruction (SIE-entry) has been issued initiating a state change of the target processor from hypervisor mode to guest mode. In step 701, R is set to 1 in IRTE assigned to the target processor as well as a logical processor ID (TrgtPU#) of the target processor provided. Then, the method ends with 702. For firmware and hardware
- 10 TrgtPU# refers to a physical ID of a processing unit (1 physical PU), while for zOS and a logical partition (LPAR) TrgtPU# refers to a logical ID of a processing unit (logical PU).
- [0131] In step 710, PBU sets a bit in DIBV assigned to vCPU using DIBV-Idx from MSI-X in order to indicate that there is an interrupt signal targeted to vCPU from VF. In step 712, PBU checks whether IRTE is blocked, i.e. whether $\text{IRTE.dIBIA} = 1$. In case IRTE assigned to vCPU and thus vCPU is blocked from receiving further interrupts, the method ends with
- 15 step 714. In case IRTE is unblocked, the method continues with step 716 in which it is checked by PBU whether vCPU is running, i.e. whether R is set in IRTE.
- [0132] If R is set, the method continues with step 718 executing a directed addressing. In step 718, dIBPIA and dPIA are set to 1 in IRTE indicating that vCPU is currently blocked from receiving interrupt signals and that an interrupt addressed to vCPU is pending. In step
- 20 720, it is checked whether IRTE, more precisely the state of R and/or TrgtPU# of IRTE changed compared to IRTE in step 78. Thus, a double fetch scheme is implemented reading IRTE twice in order to ensure that no relevant changes have taken place between the readings, e.g., due to a SIE-entry of another guest like shown in step 722.
- [0133] In step 722, a SIE-entry instruction for another guest is executed on the target processor. In step 724, the other guest reads IRTE of previous guest and issues in step 726 an atomic reset command for R, i.e. setting $R = 0$ and indicating that vCPU is not running anymore. Furthermore, dPIA is read from IRTE. In step 728, it is checked whether dPIA is set ($\text{IRTE.dPIA} = 1$) indicating that an interrupt for vCPU is still pending. If no interrupt is
- 30 pending, the method ends with step 730. If an interrupt is still pending, in step 732, a pending

interrupt indicator PU.dPIA is reset on the target PU as well as IRTE.dPIA of IRTE and a broadcast for the pending interrupt is initiated. Thus, in case a relevant change of IRTE is determined in step 720, the interrupt is broadcasted.

[0134] In case no relevant change of IRTE is determined in step 720, the method continuous
 5 with step 734. In step 734, the interrupt signal (Directed PCI-Interrupt SYSOP) is forwarded directed to the target PU, also referred to as directed PU. In step 736, directed PU receives directed PCI interrupt and sets a pending interrupt indicator PU.dPIA on the directed PU in step 738. In step 739 it is checked whether directed PU is masked, i.e. in general prevented from receiving and executing interrupts. If directed PU is masked the method ends with step
 10 740. If directed PU is unmasked, e.g., due to an unmasking as shown in step 742, the method continues with an interrupt executed by the firmware, e.g. millicode, of directed PU (mCode IO-Irpt) in step 744. In step 746, PU.dPIA and IRTE.dPIA are reset to indicate that the interrupt is not pending anymore.

[0135] In step 748, the operating system interrupt handler (OS IO-Irpt) is called and in step
 15 750 reads and resets the DIBV bit set in step 710. In step 752, it is looped over all DIBV bits of the DIBV assigned to the target PU, i.e. directed PU. Thus, all interrupts for the target PU may successively be handled. In case all DIBV bits have been handled, the target PU is unblocked in step 754 (SIC.OC17) by resetting IRTE.dIBPIA. Furthermore, DIBV is re-read in order to determine in step 756, whether meanwhile another DIBV bit has been set. If this is
 20 the case, the respective interrupt is handled, else the method ends with step 758.

[0136] If the result of the check in step 716 is that R is not set, the method continues with step 760 executing a broadcast as a fallback. In step 760, a directed interrupt summary indicator is enabled, e.g. a bit is set, in a directed interrupt summary vector. Each bit of the interrupt summary vector is assigned to a CPU, indicating whether there is any interrupt to be
 25 handled by the respective CPU. In step 764, the interrupt is broadcasted (SIGI.enq.IBPIA) and received by any PU in step 766. In step 768, a blocking bit is set in IBPIA for the respective PU, indicating that the PU is currently blocked from receiving interrupts. In step 770, it is checked whether IBPIA has been changed by setting the blocking bit, i.e., whether IBPIA $0 \rightarrow 1$. If IBPIA has not been changed, i.e., has already been blocked, the method ends
 30 with step 772. If IBPIA has been changed, in step 774, a pending bit is set in PIA for the

respective PU. In step 776 it is checked whether the PU is masked, i.e. in general prevented from receiving and executing interrupts. If the PU is masked the method ends with step 778. If the PU is unmasked, e.g., due to an unmasking as shown in step 780, the method continues with an interrupt executed by the firmware, e.g. millicode, of the PU (mCode IO-Irpt) in step 5 782. In step 784, the pending bit in PIA is reset to indicate that the interrupt is not pending anymore.

[0137] In step 786, the operating system interrupt handler (OS IO-Irpt) is called and in step 788 reads and resets the DISB bit set in step 760. In steps 790 and 792, it is signaled to the corresponding directed PU that the interrupt is handled. In step 794, it is looped over all DISB 10 bits of the DISB array, each bit assigned to another PU. DISB summarizes all interrupts to be handled by broadcasting. The interrupts are ordered according to the PU to which they are targeted. Thus, all interrupts to be handled by broadcasting may successively be handled by PU. In case all DISB bits have been handled, the PU is unblocked in step 796 (SIC.OC1) by resetting IBPIA. Furthermore, DISB is re-read in order to determine in step 798, whether 15 meanwhile another DISB bit has been set. If this is the case, the respective interrupt is handled, else the method ends with step 799.

[0138] The guest operating system may for example be implemented using a pageable storage mode guest. A pageable guest, e.g., in a z/Architecture®, may be interpretively executed via an Start Interpretive Execution (SIE) instruction, at level 2 of interpretation. For 20 instance, a logical partition (LPAR) hypervisor executes the SIE instruction to begin a logical partition in physical, fixed memory. An operating system in that logical partition, e.g., z/VM®, may issue the SIE instruction to execute its guests (virtual) machines in its virtual storage. Thus, the LPAR hypervisor may use level-1 SIE and the z/VM® hypervisor may use level-2 SIE.

[0139] According to embodiments, a computer system is a System z® server offered by International Business Machines Corporation. System z® is based on the z/Architecture® offered by International Business Machines Corporation. Details regarding the z/Architecture® are described in an IBM® publication entitled, “z/Architecture Principles of Operation,” IBM® Publication No. SA22-7832-11, August 25, 2017, which is hereby 30 incorporated herein by reference in its entirety. IBM®, System z® and z/Architecture® are

registered trademarks of International Business Machines Corporation, Armonk, N.Y. Other names used herein may be registered trademarks, trademarks or product names of International Business Machines Corporation or other companies.

[0140] According to embodiments, computer systems of other architectures may implement and use one or more aspects of the present invention. As examples, servers other than System z® servers, such as Power Systems servers or other servers offered by International Business Machines Corporation, or servers of other companies implement, use and/or benefit from one or more aspects of the present invention. Further, although in the example herein, the bus connected modules and bus attachment device are considered a part of a server, in other
5
10
15
embodiments, they do not have to necessarily be considered a part of the server, but may simply be considered as being coupled to system memory and/or other components of a computer system. The computer system need not be a server. Further, although the bus connected modules may be PCIe, one or more aspects of the present invention are usable with other bus connected modules. PCIe adapter and PCIe functions are just examples. Further, one or more aspects of the present invention may be applicable to interrupt schemes other than PCI MSI or PCI MSI-X. Yet further, although examples are described in which bits are set, in other embodiments, bytes or other type of indicators may be set. Moreover, the DTE and other structures may include more, less or different information.

[0141] Further, other types of computer systems may benefit from one or more aspects of the present invention. As an example, a data processing system suitable for storing and/or
20
executing program code is usable that includes at least two processors coupled directly or indirectly to memory elements through a system bus. The memory elements include, for instance, local memory employed during actual execution of the program code, bulk storage, and cache memory which provide temporary storage of at least some program code in order to
25
reduce the number of times code must be retrieved from bulk storage during execution.

[0142] Input/Output or I/O devices including, but not limited to, keyboards, displays, pointing devices, DASD, tape, CDs, DVDs, thumb drives and other memory media, etc., may be coupled to the system either directly or through intervening I/O controllers. Network adapters may also be coupled to the system to enable the data processing system to become
30
coupled to other data processing systems or remote printers or storage devices through

intervening private or public networks. Modems, cable modems, and Ethernet cards are just a few of the available types of network adapters.

[0143] Referring to FIG. 20, representative components of a host computer system 800 to implement one or more aspects of the present invention are portrayed. The representative host computer 800 comprises one or more processors, e.g., CPUs, 801 in communication with computer memory 802, as well as I/O interfaces to storage media devices 811 and networks 810 for communicating with other computers or SANs and the like. The CPUs 801 is compliant with an architecture having an architected instruction set and architected functionality. The CPUs 801 may have dynamic address translation (DAT) 803 for transforming program addresses, virtual addresses, into real addresses of memory. A DAT may comprise a translation lookaside buffer (TLB) 807 for caching translations so that later accesses to the block of computer memory 802 do not require the delay of address translation. A cache 809 may be employed between computer memory 802 and the CPUs 801. The cache 809 may be hierarchically structured providing a large, high level cache available to more than one CPU and smaller, faster, lower level caches between the high-level cache and each CPU. In some implementations, the lower level caches may be split to provide separate low-level caches for instruction fetching and data accesses. According to embodiments, an instruction may be fetched from memory 802 by an instruction fetch unit 804 via a cache 809. The instruction may be encoded in an instruction decode unit 806 and dispatched, in some embodiments with other instructions, to instruction execution unit or units 808. Several execution units 808 may be employed, for example an arithmetic execution unit, a floating-point execution unit and a branch instruction execution unit. The instruction is executed by the execution unit, accessing operands from instruction specified registers or memory as needed. If an operand is to be accessed, e.g., loaded or stored, from memory 802, a load/store unit 805 may handle the access under control of the instruction being executed. Instructions may be executed in hardware circuits or in internal microcode, i.e., firmware, or by a combination of both.

[0144] A computer system may comprise information in local or main storage, as well as addressing, protection, and reference and change recording. Some aspects of addressing include the format of addresses, the concept of address spaces, the various types of addresses, and the manner in which one type of address is translated to another type of address. Some of

main storage includes permanently assigned storage locations. Main storage provides the system with directly addressable fast-access storage of data. Both data and programs are to be loaded into main storage, e.g., from input devices, before they may be processed.

[0145] Main storage may include one or more smaller, faster-access buffer storages, sometimes called caches. A cache may be physically associated with a CPU or an I/O processor. The effects, except on performance, of the physical construction and use of distinct storage media may generally not be observable by programs executed.

[0146] Separate caches may be maintained for instructions and for data operands. Information within a cache may be maintained in contiguous bytes on an integral boundary called a cache block or cache line. A model may provide an EXTRACT CACHE ATTRIBUTE instruction which returns the size of a cache line in bytes. A model may also provide PREFETCH DATA and PREFETCH DATA RELATIVE LONG instructions which effects the prefetching of storage into the data or instruction cache or the releasing of data from the cache.

[0147] Storage may be viewed as a long horizontal string of bits. For most operations, accesses to storage may proceed in a left-to-right sequence. The string of bits is subdivided into units of eight bits. An eight-bit unit is called a byte, which is the basic building block of all information formats. Each byte location in storage may be identified by a unique nonnegative integer, which is the address of that byte location, also referred to as the byte address. Adjacent byte locations may have consecutive addresses, starting with 0 on the left and proceeding in a left-to-right sequence. Addresses are unsigned binary integers and may e.g. be 24, 31, or 64 bits.

[0148] Information is transmitted between memory and CPUs one byte, or a group of bytes, at a time. Unless otherwise specified, in, for instance, the z/Architecture®, a group of bytes in memory is addressed by the leftmost byte of the group. The number of bytes in the group is either implied or explicitly specified by the operation to be performed. When used in a CPU operation, a group of bytes is called a field. Within each group of bytes, in, for instance, the z/Architecture®, bits are numbered in a left-to-right sequence. In the z/Architecture®, the leftmost bits are sometimes referred to as the “high-order” bits and the rightmost bits as the “low-order” bits. Bit numbers are not storage addresses, however. Only bytes may be

addressable. To operate on individual bits of a byte in storage, the entire byte may be accessed. The bits in a byte may be numbered 0 through 7, from left to right in, e.g., the z/Architecture. The bits in an address may be numbered 8-31 or 40-63 for 24-bit addresses, or 1-31 or 33-63 for 31-bit addresses; they are numbered 0-63 for 64-bit addresses. Within any other fixed-length format of multiple bytes, the bits making up the format may be consecutively numbered starting from 0. For purposes of error detection, and in preferably for correction, one or more check bits may be transmitted with each byte or with a group of bytes. Such check bits are generated automatically by the machine and cannot be directly controlled by the program. Storage capacities are expressed in number of bytes. When the length of a storage-operand field is implied by the operation code of an instruction, the field is said to have a fixed length, which may be one, two, four, eight, or sixteen bytes. Larger fields may be implied for some instructions. When the length of a storage-operand field is not implied but is stated explicitly, the field is said to have a variable length. Variable-length operands may vary in length by increments of one byte or with some instructions, in multiples of two bytes or other multiples. When information is placed in storage, the contents of only those byte locations are replaced that are included in the designated field, even though the width of the physical path to storage may be greater than the length of the field being stored.

[0149] Certain units of information are to be on an integral boundary in storage. A boundary is called integral for a unit of information when its storage address is a multiple of the length of the unit in bytes. Special names are given to fields of 2, 4, 8, and 16 bytes on an integral boundary. A halfword is a group of two consecutive bytes on a two-byte boundary and is the basic building block of instructions. A word is a group of four consecutive bytes on a four-byte boundary. A doubleword is a group of eight consecutive bytes on an eight-byte boundary. A quadword is a group of 16 consecutive bytes on a 16-byte boundary. When storage addresses designate halfwords, words, doublewords, and quadwords, the binary representation of the address contains one, two, three, or four rightmost zero bits, respectively. Instructions are to be on two-byte integral boundaries. The storage operands of most instructions do not have boundary-alignment requirements.

[0150] On devices that implement separate caches for instructions and data operands, a significant delay may be experienced if the program stores into a cache line from which

instructions are subsequently fetched, regardless of whether the store alters the instructions that are subsequently fetched.

[0151] In one embodiment, the invention may be practiced by software, sometimes referred to licensed internal code, firmware, micro-code, milli-code, pico-code and the like, any of which would be consistent with the present invention. Referring to FIG. 20, software program code which embodies the present invention may be accessed from long-term storage media devices 811, such as a CD-ROM drive, tape drive or hard drive. The software program code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from computer memory 802 or storage of one computer system over a network 810 to other computer systems for use by users of such other systems.

[0152] Software program code may comprise an operating system which controls the function and interaction of the various computer components and one or more application programs. Program code may be paged from storage media device 811 to the relatively higher-speed computer storage 802 where it is available for processing by processor 801. Well known techniques and methods for embodying software program code in memory, on physical media, and/or distributing software code via networks may be used. Program code, when created and stored on a tangible medium, including but not limited to electronic memory modules (RAM), flash memory, Compact Discs (CDs), DVDs, Magnetic Tape, and may be referred to as a “computer program product”. The computer program product medium may be readable by a processing circuit preferably in a computer system for execution by the processing circuit.

[0153] FIG. 21 illustrates a representative workstation or server hardware system in which embodiments of the present invention may be implemented. The system 820 of FIG. 21 comprises a representative base computer system 821, such as a personal computer, a workstation or a server, including optional peripheral devices. The base computer system 821 includes one or more processors 826 and a bus employed to connect and enable communication between the processor(s) 826 and the other components of the system 821 in accordance with known techniques. The bus connects the processor 826 to memory 825 and long-term storage 827 which may include a hard drive, including any of magnetic media, CD, DVD and Flash Memory for example, or a tape drive for example. The system 821 might also

include a user interface adapter, which connects the microprocessor 826 via the bus to one or more interface devices, such as a keyboard 824, a mouse 823, a printer/scanner 830 and/or other interface devices, which may be any user interface device, such as a touch sensitive screen, digitized entry pad, etc. The bus also connects a display device 822, such as an LCD
5 screen or monitor, to the microprocessor 826 via a display adapter.

[0154] The system 821 may communicate with other computers or networks of computers by way of a network adapter capable of communicating 828 with a network 829. Example network adapters are communications channels, token ring, Ethernet or modems. Alternatively, the system 821 may communicate using a wireless interface, such as a cellular
10 digital packet data (CDPD) card. The system 821 may be associated with such other computers in a Local Area Network (LAN) or a Wide Area Network (WAN), or the system 821 may be a client in a client/server arrangement with another computer, etc.

[0155] FIG. 22 illustrates a data processing network 840 in which embodiments of the present may be implemented. The data processing network 840 may include a plurality of
15 individual networks, such as a wireless network and a wired network, each of which may include a plurality of individual workstations 841, 842, 843, 844. Additionally, as those skilled in the art will appreciate, one or more LANs may be included, where a LAN may comprise a plurality of intelligent workstations coupled to a host processor.

[0156] Still referring to FIG. 22, the networks may also include mainframe computers or
20 servers, such as a gateway computer, e.g., client server 846, or application server, e.g., remote server 848 which may access a data repository and may also be accessed directly from a workstation 845. A gateway computer 846 may serve as a point of entry into each individual network. A gateway may be needed when connecting one networking protocol to another. The gateway 846 may be preferably coupled to another network, like the Internet 847 for example,
25 by means of a communications link. The gateway 846 may also be directly coupled to one or more workstations 841, 842, 843, 844 using a communications link. The gateway computer may be implemented utilizing an IBM eServer™ System z® server available from International Business Machines Corporation.

[0157] Referring concurrently to FIG. 21 and FIG. 22, software programming code which
30 may embody the present invention may be accessed by the processor 826 of the system 820

from long-term storage media 827, such as a CD-ROM drive or hard drive. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users 850, 851 from the memory or storage of one
5 computer system over a network to other computer systems for use by users of such other systems.

[0158] Alternatively, the programming code may be embodied in the memory 825, and accessed by the processor 826 using the processor bus. Such programming code may include an operating system which controls the function and interaction of the various computer
10 components and one or more application programs 832. Program code may be paged from storage media 827 to high-speed memory 825, where it is available for processing by the processor 826. Well known techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks may be used.

15 [0159] The cache that is most readily available to a processor, i.e., which may be faster and smaller than other caches of the processor, is the lowest cache, also referred to as L1 or level one cache, and main memory is the highest-level cache, also referred to as Ln, e.g., L3, if there are n, e.g., n = 3, levels. The lowest-level cache may be divided into an instruction cache, also referred to as I-cache, holding machine readable instructions to be executed and a data cache,
20 also referred to as D-Cache, holding data operands.

[0160] Referring to FIG. 23, an exemplary processor embodiment is depicted for processor 826. One or more levels of cache 853 may be employed to buffer memory blocks in order to improve processor performance. The cache 853 is a high-speed buffer holding cache lines of memory data that are likely to be used. Cache lines may for example be 64, 128 or 256 bytes
25 of memory data. Separate caches may be employed for caching instructions and for caching data. Cache coherence, i.e., synchronization of copies of lines in memory and the caches, may be provided by various suitable algorithms, e.g., “snoop” algorithms. Main memory storage 825 of a processor system may be referred to as a cache. In a processor system having 4 levels of cache 853, main storage 825 is sometimes referred to as the level 5 (L5) cache, since it may
30 be faster and only holds a portion of the non-volatile storage that is available to a computer

system. Main storage 825 “caches” pages of data paged in and out of the main storage 825 by the operating system.

[0161] A program counter (instruction counter) 861 keeps track of the address of the current instruction to be executed. A program counter in a z/Architecture® processor is 64 bits and
5 may be truncated to 31 or 24 bits to support prior addressing limits. A program counter may be embodied in a program status word (PSW) of a computer such that it persists during context switching. Thus, a program in progress, having a program counter value, may be interrupted by, for example, the operating system resulting in a context switch from the program environment to the operating system environment. The PSW of the program maintains the
10 program counter value while the program is not active, and the program counter in the PSW of the operating system is used while the operating system is executing. The program counter may be incremented by an amount equal to the number of bytes of the current instruction. Reduced Instruction Set Computing (RISC) instructions may be fixed length, while Complex Instruction Set Computing (CISC) instructions may be variable length. Instructions of the IBM
15 z/Architecture® are CISC instructions having a length of 2, 4 or 6 bytes. The Program counter 861 may be modified by either a context switch operation or a branch taken operation of a branch instruction for example. In a context switch operation, the current program counter value is saved in the program status word along with other state information about the program being executed, such as condition codes, and a new program counter value is loaded
20 pointing to an instruction of a new program module to be executed. A branch taken operation may be performed in order to permit the program to make decisions or loop within the program by loading the result of the branch instruction into the program counter 861.

[0162] An instruction fetch unit 855 may be employed to fetch instructions on behalf of the processor 826. The fetch unit either fetches “next sequential instructions”, target instructions
25 of branch taken instructions, or first instructions of a program following a context switch. Modern instruction fetch units may employ prefetch techniques to speculatively prefetch instructions based on the likelihood that the prefetched instructions might be used. For example, a fetch unit may fetch 16 bytes of instruction that includes the next sequential instruction and additional bytes of further sequential instructions.

[0163] The fetched instructions may then be executed by the processor 826. According to embodiments, the fetched instruction(s) may be passed to a dispatch unit 856 of the fetch unit. The dispatch unit decodes the instruction(s) and forwards information about the decoded instruction(s) to appropriate units 857, 858, 860. An execution unit 857 may receive
5 information about decoded arithmetic instructions from the instruction fetch unit 855 and may perform arithmetic operations on operands according to the opcode of the instruction. Operands may be provided to the execution unit 857 preferably either from memory 825, architected registers 859 or from an immediate field of the instruction being executed. Results of the execution, when stored, may be stored either in memory 825, registers 859 or in other
10 machine hardware, such as control registers, PSW registers and the like.

[0164] A processor 826 may comprise one or more units 857, 858, 860 for executing the function of the instruction. Referring to FIG. 24A, an execution unit 857 may communicate with architected general registers 859, a decode/dispatch unit 856, a load store unit 860, and other 865 processor units by way of interfacing logic 871. An execution unit 857 may employ
15 several register circuits 867, 868, 869 to hold information that the arithmetic logic unit (ALU) 866 operates on. The ALU performs arithmetic operations such as add, subtract, multiply and divide as well as logical function such as And, Or, Exclusive-or (XOR), Rotate and Shift. Preferably the ALU may support specialized operations that are design dependent. Other circuits may provide other architected facilities 872 including condition codes and recovery
20 support logic for example. The result of an ALU operation may be held in an output register circuit 870 which is configured to forward the result to a variety of other processing functions. There are many arrangements of processor units, the present description is only intended to provide a representative understanding of one embodiment.

[0165] An ADD instruction for example may be executed in an execution unit 857 having
25 arithmetic and logical functionality while a floating-point instruction for example would be executed in a floating-point execution having specialized floating point capability. Preferably, an execution unit operates on operands identified by an instruction by performing an opcode defined function on the operands. For example, an ADD instruction may be executed by an execution unit 857 on operands found in two registers 859 identified by register fields of the
30 instruction.

- [0166] The execution unit 857 performs the arithmetic addition on two operands and stores the result in a third operand where the third operand may be a third register or one of the two source registers. The execution unit preferably utilizes an Arithmetic Logic Unit (ALU) 866 that is capable of performing a variety of logical functions such as Shift, Rotate, And, Or and XOR as well as a variety of algebraic functions including any of add, subtract, multiply, divide. Some ALUs 866 are designed for scalar operations and some for floating point. Data may be big endian, where the least significant byte is at the highest byte address, or little endian, where the least significant byte is at the lowest byte address, depending on architecture. The IBM z/Architecture® is big endian. Signed fields may be sign and magnitude, 1's complement or 2's complement depending on architecture. A 2's complement number may be advantageous in that the ALU does not need to design a subtract capability since either a negative value or a positive value in 2's complement requires only an addition within the ALU. Numbers may be described in shorthand, where a 12-bit field defines an address of a 4,096-byte block and described as a 4 Kbyte (Kilo-byte) block, for example.
- [0167] Referring to FIG. 24B, branch instruction information for executing a branch instruction may be sent to a branch unit 858 which often employs a branch prediction algorithm such as a branch history table 882 to predict the outcome of the branch before other conditional operations are complete. The target of the current branch instruction will be fetched and speculatively executed before the conditional operations are complete. When the conditional operations are completed the speculatively executed branch instructions are either completed or discarded based on the conditions of the conditional operation and the speculated outcome. A branch instruction may test condition codes and branch to a target address if the condition codes meet the branch requirement of the branch instruction, a target address may be calculated based on several numbers including ones found in register fields or an immediate field of the instruction for example. The branch unit 858 may employ an ALU 874 having a plurality of input register circuits 875, 876, 877 and an output register circuit 880. The branch unit 858 may communicate with general registers 859, decode dispatch unit 856 or other circuits 873, for example.
- [0168] The execution of a group of instructions may be interrupted for a variety of reasons including a context switch initiated by an operating system, a program exception or error causing a context switch, an I/O interrupt signal causing a context switch or multi-threading

activity of a plurality of programs in a multi-threaded environment, for example. Preferably a context switch action saves state information about a currently executing program and then loads state information about another program being invoked. State information may be saved in hardware registers or in memory for example. State information preferably comprises a
5 program counter value pointing to a next instruction to be executed, condition codes, memory translation information and architected register content. A context switch activity may be exercised by hardware circuits, application programs, operating system programs or firmware code, like e.g. microcode, pico-code or licensed internal code (LIC), alone or in combination.

[0169] A processor accesses operands according to instruction defined methods. The
10 instruction may provide an immediate operand using the value of a portion of the instruction, may provide one or more register fields explicitly pointing to either general purpose registers or special purpose registers, like e.g. floating-point registers for example. The instruction may utilize implied registers identified by an opcode field as operands. The instruction may utilize memory locations for operands. A memory location of an operand may be provided by a
15 register, an immediate field, or a combination of registers and immediate field as exemplified by the z/Architecture® long displacement facility wherein the instruction defines a base register, an index register and an immediate field, i.e., displacement field, that are added together to provide the address of the operand in memory for example. Location herein may imply a location in main memory unless otherwise indicated.

[0170] Referring to FIG. 24C, a processor accesses storage using a load/store unit 860. The load/store unit 860 may perform a load operation by obtaining the address of the target operand in memory 853 and loading the operand in a register 859 or another memory 853 location, or may perform a store operation by obtaining the address of the target operand in memory 853 and storing data obtained from a register 859 or another memory 853 location in
25 the target operand location in memory 853. The load/store unit 860 may be speculative and may access memory in a sequence that is out-of-order relative to instruction sequence, however the load/store unit 860 is to maintain the appearance to programs that instructions were executed in order. A load/store unit 860 may communicate with general registers 859, decode/dispatch unit 856, cache/memory interface 853 or other elements 883 and comprises
30 various register circuits, ALUs 885 and control logic 890 to calculate storage addresses and to provide pipeline sequencing to keep operations in-order. Some operations may be out of order

but the load/store unit provides functionality to make the out of order operations to appear to the program as having been performed in order.

[0171] Preferably addresses that an application program “sees” are often referred to as virtual addresses. Virtual addresses are sometimes also referred to as “logical addresses” and “effective addresses”. These virtual addresses are virtual in that they are redirected to physical memory location by one of a variety of dynamic address translation (DAT) technologies including, but not limited to, simply prefixing a virtual address with an offset value, translating the virtual address via one or more translation tables, the translation tables preferably comprising at least a segment table and a page table alone or in combination, preferably, the segment table having an entry pointing to the page table. In the z/Architecture®, a hierarchy of translation is provided including a region first table, a region second table, a region third table, a segment table and an optional page table. The performance of the address translation is often improved by utilizing a translation lookaside buffer (TLB) which comprises entries mapping a virtual address to an associated physical memory location. The entries are created when the DAT translates a virtual address using the translation tables. Subsequent use of the virtual address may then utilize the entry of the fast TLB rather than the slow sequential translation table accesses. TLB content may be managed by a variety of replacement algorithms including Least Recently used (LRU).

[0172] Each processor of a multi-processor system has the responsibility to keep shared resources, such as I/O, caches, TLBs and memory, interlocked for coherency. So-called “snoop” technologies may be utilized in maintaining cache coherency. In a snoop environment, each cache line may be marked as being in any one of a shared state, an exclusive state, a changed state, an invalid state and the like in order to facilitate sharing.

[0173] I/O units 854 may provide the processor with means for attaching to peripheral devices including tape, disc, printers, displays, and networks for example. I/O units are often presented to the computer program by software drivers. In mainframes, such as the System z® from IBM®, channel adapters and open system adapters are I/O units of the mainframe that provide the communications between the operating system and peripheral devices.

[0174] Further, other types of computer systems may benefit from one or more aspects of the present invention. As an example, the computer system may comprise an emulator, e.g.,

- software or other emulation mechanisms, in which a particular architecture including, for example, instruction execution, architected functions, such as address translation, and architected registers, is emulated or in which a subset thereof is emulated, e.g., on a native computer system having a processor and memory. In such an environment, one or more
- 5 emulation functions of the emulator may implement one or more aspects of the present invention, even though a computer executing the emulator may have a different architecture than the capabilities being emulated. For example, in emulation mode, the specific instruction or operation being emulated may be decode, and an appropriate emulation function may be built to implement the individual instruction or operation.
- 10 [0175] In an emulation environment, a host computer may for example comprise a memory to store instructions and data, an instruction fetch unit to fetch instructions from memory and to optionally provide local buffering for the fetched instruction, an instruction decode unit to receive the fetched instructions and to determine the type of instructions that have been fetched, and an instruction execution unit to execute the instructions. Execution may include
- 15 loading data into a register from memory, storing data back to memory from a register, an/or performing some type of arithmetic or logical operation, as determined by the decode unit. For example, each unit may be implemented in software. Operations being performed by the units may be implemented as one or more subroutines within emulator software.
- [0176] More particularly, in a mainframe, architected machine instructions are used by
- 20 programmers, like e.g. “C” programmers, e.g., by way of a compiler application. These instructions stored in the storage medium may be executed natively in a z/Architecture® IBM® Server, or alternatively in machines executing other architectures. They may be emulated in the existing and in future IBM® mainframe servers and on other machines of IBM® (e.g., Power Systems servers and System x® Servers). They may be executed in
- 25 machines running Linux on a wide variety of machines using hardware manufactured by IBM®, Intel®, AMD™, and others. Besides execution on that hardware under a z/Architecture®, Linux may be used as well as machines which use emulation by Hercules, UMX, or FSI (Fundamental Software, Inc), where generally execution is in an emulation mode. In emulation mode, emulation software is executed by a native processor to emulate the
- 30 architecture of an emulated processor.

[0177] A native processor may execute emulation software comprising either firmware or a native operating system to perform emulation of the emulated processor. The emulation software is responsible for fetching and executing instructions of the emulated processor architecture. The emulation software maintains an emulated program counter to keep track of instruction boundaries. The emulation software may fetch one or more emulated machine instructions at a time and convert the one or more emulated machine instructions to a corresponding group of native machine instructions for execution by the native processor. These converted instructions may be cached such that a faster conversion may be accomplished. Notwithstanding, the emulation software is to maintain the architecture rules of the emulated processor architecture so as to assure operating systems and applications written for the emulated processor operate correctly. Furthermore, the emulation software is to provide resources identified by the emulated processor architecture including, but not limited to, control registers, general purpose registers, floating point registers, dynamic address translation function including segment tables and page tables for example, interrupt mechanisms, context switch mechanisms, Time of Day (TOD) clocks and architected interfaces to I/O subsystems such that an operating system or an application program designed to run on the emulated processor, may be run on the native processor having the emulation software.

[0178] A specific instruction being emulated is decoded, and a subroutine is called to perform the function of the individual instruction. An emulation software function emulating a function of an emulated processor is implemented, for example, in a “C” subroutine or driver, or some other method of providing a driver for the specific hardware.

[0179] In FIG. 25, an example of an emulated host computer system 892 is provided that emulates a host computer system 800' of a host architecture. In the emulated host computer system 892, the host processor, i.e., CPU, 891 is an emulated host processor or virtual host processor and comprises an emulation processor 893 having a different native instruction set architecture than that of the processor 891 of the host computer 800'. The emulated host computer system 892 has memory 894 accessible to the emulation processor 893. In the example embodiment, the memory 894 is partitioned into a host computer memory 896 portion and an emulation routines 897 portion. The host computer memory 896 is available to programs of the emulated host computer 892 according to host computer architecture. The

- emulation processor 893 executes native instructions of an architected instruction set of an architecture other than that of the emulated processor 891, the native instructions obtained from emulation routines memory 897, and may access a host instruction for execution from a program in host computer memory 896 by employing one or more instructions obtained in a sequence and access/decode routine which may decode the host instructions accessed to determine a native instruction execution routine for emulating the function of the host instruction accessed. Other facilities that are defined for the host computer system 800' architecture may be emulated by architected facilities routines, including such facilities as general purpose registers, control registers, dynamic address translation and I/O subsystem support and processor cache, for example. The emulation routines may also take advantage of functions available in the emulation processor 893, such as general registers and dynamic translation of virtual addresses, to improve performance of the emulation routines. Special hardware and off-load engines may also be provided to assist the processor 893 in emulating the function of the host computer 800'.
- 15 [0180] It is understood that one or more of the aforementioned embodiments of the invention may be combined as long as the combined embodiments are not mutually exclusive. Ordinal numbers, like e.g. 'first' and 'second', are used herein to indicate different element assigned with the same name, but do not necessarily establish any order of the respective elements.
- [0181] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.
- 20 [0182] The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.
- [0183] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage
- 30

- medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable
- 5 computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any
- 10 suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.
- 15 [0184] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers,
- 20 firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.
- 25 [0185] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as
- 30 Smalltalk, C++ or the like, and conventional procedural programming languages, such as the 'C' programming language or similar programming languages. The computer readable

program instructions may execute entirely on the user computer system's computer, partly on the user computer system's computer, as a stand-alone software package, partly on the user computer system's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user computer system's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0186] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0187] These computer readable program instructions may be provided to a processor of a general-purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0188] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the
5 computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0189] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard,
10 each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes
15 be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

20 [0190] Possible combinations of features described above may be the following:

1. A method for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, the computer system further comprising one or more bus connected modules being operationally connected with the plurality of processors via a bus
25 and a bus attachment device the computer system further comprising a memory operationally connected with the bus attachment device,
each processor of the plurality of processors being assigned with a logical processor ID used by the bus attachment device to address the respective processor,

each processor of the plurality of processors assigned for usage by the guest operating system further being assigned with an interrupt target ID used by the operating system and one or more bus connected modules to address the respective processor,

the memory comprising a directed interrupt signal vector per interrupt target ID assigned to the respective interrupt target ID, each directed interrupt signal vector comprising a directed interrupt signal indicator per bus connected module assigned to the respective bus connected module, each directed interrupt signal vector indicating whether there is an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID to be handled,

the method comprising:
receiving by the bus attachment device an interrupt signal from one of the bus connected modules of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal,
selecting by the bus attachment device the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed,
selecting by the bus attachment device in the selected directed interrupt signal vector the directed interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal,
updating by the bus attachment device the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that there is an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID to be handled,
forwarding by the bus attachment device the interrupt signal to the target processor.

2. The method of item 1, the interrupt signal indicators assigned to the same bus connected module each comprising the same offset within the directed interrupt signal vector comprising the respective interrupt signal indicator.

3. The method of any of the previous items, the directed interrupt signal vectors each being implemented as a contiguous area in memory.

4. The method of any of the previous items, the directed interrupt signal indicators each being implemented as a single bit.

5. The method of any of the previous items, the method further comprising retrieving by the bus attachment device a copy of an interrupt table entry assigned to the received interrupt target ID from an interrupt table stored in the memory, the interrupt table entry comprising a directed interrupt signal vector address indicator indicating a memory address of the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed, the bus attachment device using the memory address of the respective directed interrupt signal vector for the selecting of the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed.

6. The method of item 5, the method further comprising retrieving by the bus attachment device a copy of a device table entry from a device table stored in the memory, the device table entry comprising an interrupt table address indicator indicating a memory address of the interrupt table, the bus attachment device using the memory address of the interrupt table for the retrieving of the first copy of the interrupt table entry.

7. The method of item 6, the device table entry further comprising a directed interrupt signal offset indicator indicating an offset of the directed interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal.

8. The method of any of the previous items, the memory further comprising a directed interrupt summary vector with a directed interrupt summary indicator per interrupt target ID, each directed interrupt summary indicator being assigned to an interrupt target ID indicating whether there is an interrupt signal addressed to the respective interrupt target ID to be handled,

the method further comprising:

selecting by the bus attachment device the directed interrupt summary indicator assigned to the target processor ID to which the received interrupt signal is addressed,

updating by the bus attachment device the selected directed interrupt summary indicator such that the selected directed interrupt summary indicator indicating that there is an interrupt signal addressed to the respective interrupt target ID to be handled.

5 9. The method of item 8, the directed interrupt summary vector being implemented as a contiguous area in memory.

10 10. The method of any of items 8 to 9, the directed interrupt summary indicators each being implemented as a single bit.

10

11. The method of any of items 8 to 10, the interrupt table entry further comprising a directed interrupt summary vector address indicator indicating a memory address of the directed interrupt summary vector, the bus attachment device using the memory address of the directed interrupt summary vector for selecting the directed interrupt summary indicator
15 assigned to the target processor ID to which the received interrupt signal is addressed.

12. The method of any of items 8 to 11, the interrupt table entry further comprising a directed interrupt summary offset indicator indicating an offset of the directed interrupt summary indicator assigned to the target processor ID within the directed interrupt summary
20 vector.

13. The method of any of the previous claims, the method further comprising translating by the bus attachment device the interrupt target ID of the target processor received with the interrupt signal to a logical processor ID of the target processor and using the logical processor
25 ID of the target processor to address the target processor as the target of the interrupt signal, when forwarding the interrupt signal to the target processor.

14. The method of any of items 11 to 13, the method further comprising retrieving by the bus attachment device a copy of an interrupt table entry assigned to the interrupt target ID
30 from an interrupt table stored in the memory, the copy of the interrupt table entry further comprising a current mapping of the interrupt target ID to a first logical processor ID, the bus

attachment device using the copy of the interrupt table entry for the translating of the interrupt target ID of the target processor received with the interrupt signal.

15 15. The method of any of the previous items, the copy of the device table entry further comprising a direct signaling indicator indicating whether the target processor is to be addressed directly, the direct signaling indicator indicating a direct forwarding of the interrupt signal being a requirement for executing the forwarding of the interrupt signal using logical processor ID of the target processor to address the target processor directly, else the forwarding being executed using broadcasting.

10

16. The method of any of the previous items, the copy of the interrupt table entry further comprising a copy of a running indicator indicating whether the target processor identified by the interrupt target ID is scheduled for usage by the guest operating system, the target processor being scheduled for usage by the guest operating system being a further requirement for executing the forwarding of the interrupt signal using logical processor ID of the target processor to address the target processor directly, else the interrupt signal to the first operating system for handling using broadcasting.

15

17. The method of any of the previous items, the copy of the interrupt table entry further comprising an interrupt blocking indicator indicating whether the target processor identified by the interrupt target ID is currently blocked from receiving interrupt signals, the target processor being unblocked being a further requirement for executing the forwarding of the interrupt signal using logical processor ID of the target processor to address the target processor directly, else the interrupt signal to the first operating system for handling using broadcasting.

20

25

18. The method of any of the previous items, the device table entry further comprising a logical partition ID identifying a logical partition to which the guest operating system is assigned, the forwarding of the interrupt signal by the bus attachment device further comprising forwarding with the interrupt signal the logical partition ID.

30

19. The method of any of the previous items, the method further comprising retrieving by the bus attachment device an interrupt subclass ID identifying an interrupt subclass to which the received interrupt signal is assigned, the forwarding of the interrupt signal by the bus attachment device further comprising forwarding with the interrupt signal the interrupt subclass ID.

20. A computer system for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of the computer system assigned for usage by the guest operating system, the computer system further comprising one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device the computer system further comprising a memory operationally connected with the bus attachment device,

each processor of the plurality of processors being assigned with a logical processor ID used by the bus attachment device to address the respective processor,

each processor of the plurality of processors assigned for usage by the guest operating system further being assigned with an interrupt target ID used by operating system and the one or more bus connected modules to address the respective processor,

the memory comprising a directed interrupt signal vector per interrupt target ID assigned to the respective interrupt target ID, each directed interrupt signal vector comprising a directed interrupt signal indicator per bus connected module assigned to the respective bus connected module, each directed interrupt signal vector indicating whether an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID is waiting to be handled,

the computer system being configured to perform a method comprising:

receiving by the bus attachment device an interrupt signal from one of the bus connected modules of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal,

selecting by the bus attachment device the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed,

- selecting by the bus attachment device in the selected directed interrupt signal vector the directed interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal,
- updating the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID is waiting to be handled,
- forwarding the interrupt signal to the target processor.
21. A computer program product for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, the computer system further comprising one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device the computer system further comprising a memory operationally connected with the bus attachment device,
- each processor of the plurality of processors being assigned with a logical processor ID used by the bus attachment device to address the respective processor,
- each processor of the plurality of processors assigned for usage by the guest operating system further being assigned with an interrupt target ID used by the operating system and one or more bus connected modules to address the respective processor,
- the memory comprising a directed interrupt signal vector per interrupt target ID assigned to the respective interrupt target ID, each directed interrupt signal vector comprising a directed interrupt signal indicator per bus connected module assigned to the respective bus connected module, each directed interrupt signal vector indicating whether an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID is waiting to be handled,
- the computer program product comprising a computer readable non-transitory medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising:
- receiving by the bus attachment device an interrupt signal from one of the bus connected modules of the bus connected modules with an interrupt target ID, the interrupt target ID

identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal,
selecting by the bus attachment device the directed interrupt signal vector assigned to the interrupt target ID to which the received interrupt signal is addressed,
5 selecting by the bus attachment device in the selected directed interrupt signal vector the directed interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal,
updating the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that an interrupt signal issued by the respective bus
10 connected module and addressed to the respective interrupt target ID is waiting to be handled,
forwarding the interrupt signal to the target processor.

CLAIMS

1. A method for providing an interrupt signal to a guest operating system (204) executed using one or more processors of a plurality of processors (130) of a computer system (100) assigned for usage by the guest operating system (204), the computer system (100) further comprising one or more bus connected modules (120) being operationally connected with the plurality of processors (130) via a bus (102) and a bus attachment device (110), the computer system (100) further comprising a memory (140) operationally connected with the bus attachment device (110),

each processor of the plurality of processors (130) being assigned with a logical processor ID (222) used by the bus attachment device (110) to identify and address the respective processor (130),

each processor of the plurality of processors (130) assigned for usage by the guest operating system (204) further being assigned with an interrupt target ID used by the guest operating system (204) and one or more bus connected modules (120) to address the respective processor (130),

the memory (140) comprising a directed interrupt signal vector (162) per interrupt target ID assigned to the respective interrupt target ID,

the method comprising:

receiving by the bus attachment device (110) an interrupt signal from one of the bus connected modules (120) of the one or more bus connected modules (120), the interrupt signal including an interrupt target ID, the interrupt target ID identifying one of the processors (130) assigned for usage by the guest operating system (204) as a target processor (130) for handling the interrupt signal,

selecting by the bus attachment device (110) the directed interrupt signal vector (162) assigned to the interrupt target ID to which the received interrupt signal is addressed, wherein each directed interrupt signal vector (162) comprising a directed interrupt signal indicator per bus connected module (120) assigned to the respective bus connected module (120), each directed interrupt signal vector (162) indicating whether there is an interrupt signal issued by the respective bus connected module (120) and addressed to the respective interrupt target ID to be handled,

the directed interrupt signal vectors (162) each being implemented as a contiguous area in memory (140),

selecting by the bus attachment device (110) in the selected directed interrupt signal vector (162) the directed interrupt signal indicator assigned to the bus connected module (120) which issued the received interrupt signal,

updating by the bus attachment device (110) the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that there is an interrupt signal issued by the respective bus connected module (120) and addressed to the respective interrupt target ID to be handled,

translating by the bus attachment device (110) the interrupt target ID assigned to the target processor (130) received with the interrupt signal to a logical processor ID (222) of the target processor (130), and

forwarding by the bus attachment device (110) the interrupt signal to the target processor (130), the forwarding using the logical processor ID (222) of the target processor (130) to address the target processor (130) as the target of the interrupt signal, wherein the interrupt signal comprises an indication of the selected directed interrupt signal vector (162).

2. The method of claim 1, the interrupt signal indicators assigned to the same bus connected module (120) each comprising the same offset within the directed interrupt signal vector (162) comprising the respective interrupt signal indicator.

3. The method of claim 1, the directed interrupt signal indicators each being implemented as a single bit.

4. The method of claim 1, the method further comprising retrieving by the bus attachment device (110) a copy of an interrupt table entry (152) assigned to the received interrupt target ID from an interrupt table (150) stored in the memory (140), the interrupt table entry (152) comprising a directed interrupt signal vector address indicator indicating a memory address of the directed interrupt signal vector (162) assigned to the interrupt target ID to which the received interrupt signal is addressed, the bus attachment device (110) using the memory address of the respective directed interrupt signal vector (162) for the selecting of the directed interrupt signal vector (162) assigned to the interrupt target ID to which the received interrupt signal is addressed.

5. The method of claim 4, the method further comprising retrieving by the bus attachment device (110) a copy of a device table entry (146) from a device table (144) stored in the memory (140), the device table entry (146) comprising an interrupt table address indicator indicating a memory address of the interrupt table (150), the bus attachment device (110) using the memory address of the interrupt table (150) for the retrieving of the copy of the interrupt table entry (152).

6. The method of claim 5, the device table entry (146) further comprising a directed interrupt signal offset indicator indicating an offset of the directed interrupt signal indicator assigned to the bus connected module (120) which issued the received interrupt signal.

7. The method of claim 1, the memory (140) further comprising a directed interrupt summary vector (160) with a directed interrupt summary indicator per interrupt target ID, each directed interrupt summary indicator being assigned to an interrupt target ID indicating whether there is an interrupt signal addressed to the respective interrupt target ID to be handled,

the method further comprising:

selecting by the bus attachment device (110) the directed interrupt summary indicator assigned to the interrupt target ID to which the received interrupt signal is addressed,

updating by the bus attachment device (110) the selected directed interrupt summary indicator such that the selected directed interrupt summary indicator indicating that there is an interrupt signal addressed to the respective interrupt target ID to be handled.

8. The method of claim 7, the directed interrupt summary vector (160) being implemented as a contiguous area in memory (140).

9. The method of claim 7, the directed interrupt summary indicators each being implemented as a single bit.

10. The method of claim 7, the interrupt table entry (152) further comprising a directed interrupt summary vector address indicator indicating a memory address of the directed interrupt summary vector (160), the bus attachment device (110) using the memory address of the

directed interrupt summary vector (160) for selecting the directed interrupt summary indicator assigned to the target processor ID to which the received interrupt signal is addressed.

11. The method of claim 7, the memory (140) further comprising an interrupt table entry (152) assigned to the received interrupt target ID, the interrupt table entry (152) comprising a directed interrupt summary offset indicator indicating an offset of the directed interrupt summary indicator assigned to the interrupt target ID within the directed interrupt summary vector (160).

12. The method of claim 1, the method further comprising retrieving by the bus attachment device (110) a copy of an interrupt table entry (152) assigned to the interrupt target ID from an interrupt table (150) stored in the memory (140), the copy of the interrupt table entry (152) comprising a current mapping of the interrupt target ID to a first logical processor ID (222), the bus attachment device (110) using the copy of the interrupt table entry (152) for the translating of the interrupt target ID of the target processor (130) received with the interrupt signal.

13. The method of claim 1, the method further comprising retrieving by the bus attachment device (110) a copy of a device table entry (146) from a device table (144) stored in the memory (140), the device table entry (146) comprising a direct signaling indicator indicating whether the target processor (130) is to be addressed directly, the direct signaling indicator indicating a direct forwarding of the interrupt signal being a requirement for executing the forwarding of the interrupt signal using the logical processor ID (222) of the target processor (130) to address the target processor (130) directly, else the forwarding being executed using broadcasting.

14. The method of claim 1, the method further comprising retrieving by the bus attachment device (110) a copy of an interrupt table entry (152) from an interrupt table stored (150) in the memory (140), the interrupt table entry (152) comprising a copy of a running indicator (154) indicating whether the target processor (130) identified by the interrupt target ID is scheduled for usage by the guest operating system (204), the target processor (130) being scheduled for usage by the guest operating system (204) being a further requirement for executing the forwarding of the interrupt signal using the logical processor ID (222) of the target processor (130) to address the target processor (130) directly, else the interrupt signal to the guest operating system for handling using broadcasting.

15. The method of claim 1, the method further comprising retrieving by the bus attachment device (110) a copy of an interrupt table entry (152) from an interrupt table (150) stored in the memory (140), the interrupt table entry (152) comprising an interrupt blocking indicator indicating whether the target processor (130) identified by the interrupt target ID is currently blocked from receiving interrupt signals, the target processor (130) being unblocked being a further requirement for executing the forwarding of the interrupt signal using the logical processor ID (222) of the target processor (130) to address the target processor (130) directly, else the interrupt signal to the guest operating system for handling using broadcasting.

16. The method of claim 1, the method further comprising retrieving by the bus attachment device (110) a device table entry (146) from a device table (144) stored in the memory (140), the device table entry (146) comprising a logical partition ID identifying a logical partition to which the guest operating system (204) is assigned, the forwarding of the interrupt signal by the bus attachment device (110) further comprising forwarding with the interrupt signal the logical partition ID.

17. The method of claim 1, the method further comprising retrieving by the bus attachment device (110) an interrupt subclass ID identifying an interrupt subclass to which the received interrupt signal is assigned, the forwarding of the interrupt signal by the bus attachment device (110) further comprising forwarding with the interrupt signal the interrupt subclass ID.

18. A computer system (100) for providing an interrupt signal to a guest operating system (204) executed using one or more processors of a plurality of processors (130) of the computer system (100) assigned for usage by the guest operating system (204), the computer system (100) further comprising one or more bus connected modules (120) being operationally connected with the plurality of processors (130) via a bus (102) and a bus attachment device (110) the computer system (100) further comprising a memory (140) operationally connected with the bus attachment device (110),

each processor of the plurality of processors (130) being assigned with a logical processor ID (222) used by the bus attachment device (110) to identify and address the respective processor (130),

each processor of the plurality of processors (130) assigned for usage by the guest operating system (204) further being assigned with an interrupt target ID used by the guest

operating system (204) and the one or more bus connected modules (120) to address the respective processor (130),

the memory (140) comprising a directed interrupt signal vector (162) per interrupt target ID assigned to the respective interrupt target ID,

the computer system (100) being configured to perform a method comprising:

receiving by the bus attachment device (110) an interrupt signal from one of the bus connected modules (120) of the one or more bus connected modules (120), the interrupt signal including an interrupt target ID, the interrupt target ID identifying one of the processors (130) assigned for usage by the guest operating system (204) as a target processor (130) for handling the interrupt signal,

selecting by the bus attachment device (110) the directed interrupt signal vector (162) assigned to the interrupt target ID to which the received interrupt signal is addressed, wherein each directed interrupt signal vector (162) comprising a directed interrupt signal indicator per bus connected module (120) assigned to the respective bus connected module (120), each directed interrupt signal vector (162) indicating whether there is an interrupt signal issued by the respective bus connected module (120) and addressed to the respective interrupt target ID to be handled; the directed interrupt signal vectors (162) each being implemented as a contiguous area in memory (140);

selecting by the bus attachment device (110) in the selected directed interrupt signal vector (162) the directed interrupt signal indicator assigned to the bus connected module (120) which issued the received interrupt signal,

updating the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that an interrupt signal issued by the respective bus connected module (120) and addressed to the respective interrupt target ID is waiting to be handled,

translating the interrupt target ID assigned to the target processor (130) received with the interrupt signal to a logical processor ID (222) of the target processor (130),

forwarding the interrupt signal to the target processor (130), the forwarding using the logical processor ID (222) of the target processor (130) to address the target processor (130) as the target of the interrupt signal, wherein the interrupt signal comprises an indication of the selected directed interrupt signal vector (162).

19. A computer program product for providing an interrupt signal to a guest operating system (204) executed using one or more processors of a plurality of processors (130) of a computer

system (100) assigned for usage by the guest operating system (204), the computer system (100) further comprising one or more bus connected modules (120) being operationally connected with the plurality of processors (130) via a bus (102) and a bus attachment device (110) the computer system (100) further comprising a memory (140) operationally connected with the bus attachment device (110),

each processor of the plurality of processors (130) being assigned with a logical processor ID (222) used by the bus attachment device (110) to identify and address the respective processor (130),

each processor of the plurality of processors (130) assigned for usage by the guest operating system (204) further being assigned with an interrupt target ID used by the operating system and one or more bus connected modules (120) to address the respective processor (130), the memory (140) comprising a directed interrupt signal vector (162) per interrupt target ID assigned to the respective interrupt target ID,

the computer program product comprising a computer readable non-transitory medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising:

receiving by the bus attachment device (110) an interrupt signal from one of the bus connected modules (120) of the one or more bus connected modules (120), the interrupt signal including an interrupt target ID, the interrupt target ID identifying one of the processors (130) assigned for usage by the guest operating system (204) as a target processor (130) for handling the interrupt signal,

selecting by the bus attachment device (110) the directed interrupt signal vector (162) assigned to the interrupt target ID to which the received interrupt signal is addressed, wherein each directed interrupt signal vector (162) comprising a directed interrupt signal indicator per bus connected module (120) assigned to the respective bus connected module (120), each directed interrupt signal vector (162) indicating whether there is an interrupt signal issued by the respective bus connected module (120) and addressed to the respective interrupt target ID to be handled; the directed interrupt signal vectors (162) each being implemented as a contiguous area in memory (140);

selecting by the bus attachment device (110) in the selected directed interrupt signal vector (162) the directed interrupt signal indicator assigned to the bus connected module (120) which issued the received interrupt signal,

updating the selected directed interrupt signal indicator such that the respective directed interrupt signal indicator indicates that an interrupt signal issued by the

respective bus connected module (120) and addressed to the respective interrupt target ID is waiting to be handled,

translating the interrupt target ID assigned to the target processor (130) received with the interrupt signal to a logical processor ID (222) of the target processor (130), forwarding the interrupt signal to the target processor (130), the forwarding using the logical processor ID (222) of the target processor (130) to address the target processor (130) as the target of the interrupt signal, wherein the interrupt signal comprises an indication of the selected directed interrupt signal vector (162).