

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
9 December 2004 (09.12.2004)

PCT

(10) International Publication Number
WO 2004/107713 A2

- (51) International Patent Classification⁷: **H04L 29/06**
- (21) International Application Number:
PCT/US2004/014851
- (22) International Filing Date: 12 May 2004 (12.05.2004)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/444,164 23 May 2003 (23.05.2003) US
- (71) Applicant: **INTEL CORPORATION (a Delaware Corporation)** [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).
- (72) Inventor: **MUNGUIA, Peter**; 1451 West Cardinal Way, Chandler, AZ 85248 (US).
- (74) Agents: **MALLIE, Michael, J.** et al.; Blakely Sokoloff Taylor & Zafman, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, CA 90025 (US).
- (81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— *without international search report and to be republished upon receipt of that report*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*



WO 2004/107713 A2

(54) Title: PACKET COMBINING ON PCI EXPRESS

(57) Abstract: A method and device for combining packets is disclosed. In various embodiments, a plurality of packets that are compliant with a particular protocol are combined into a single packet of the same protocol. The single packet also preserves identification information for each of the combined packets within the single packet. The identification information can be used to partition and/or route the combined packets to their respective destinations.

PACKET COMBINING ON PCI EXPRESS**FIELD**

[0001] The embodiments disclosed herein relate generally to packet transmission, and more particularly to combining
5 packets prior to transmission.

BACKGROUND

[0002] Data can be sent in many different forms in a computer system or between computer systems (e.g., over a network). In some systems, data is grouped together and
10 sent as a packet. PCI Express (Peripheral Component Interconnect Special Interest Group's Peripheral Component Interconnect Express 1.0 specification, published July 22, 2002) is one example of a protocol for packetizing and transmitting data. Regardless of the protocol being used,
15 the transmission of packets requires some level of overhead and minimum packet size.

[0003] For example, some packets require a header (e.g., overhead) that contains information regarding the contents and/or destination of the packet. If a minimum packet size
20 is required before sending the packet, some systems will "stuff" additional bits into a packet that does not meet the minimum size requirement. The stuffed bits contain no useful information and reduce the efficiency of the system.

[0004] Following are some examples of PCI Express
25 transactions that typically require bit stuffing to meet minimum packet size requirements. First, a simple sequence to an IO-mapped (e.g., input/output mapped) device requires an 8-bit IO write transaction to an index register followed by an 8-bit IO cycle to the corresponding data register.
30 This translates into two full PCI Express transactions, both of minimum length and efficiency.

[0005] Second, reading a FIFO (e.g., first-in, first-out) or doorbell register consists of reading the same address repeatedly with a separate PCI Express transaction each time. Third, some special memory transactions (e.g., programming a flash memory device) require byte or word writes in command sequence order. Each of these writes requires a separate PCI Express transaction.

[0006] Thus, each of these three processes require multiple PCI Express transactions of minimum length, which is likely to involve the inefficient transmission of useless bits due to bit stuffing. Moreover, each additional PCI Express transaction requires system resources in order to form packets, to transmit packets, and to acknowledge packet transmission.

[0007] The drawbacks of packet overhead and minimum packet size are particularly noticeable in situations where data traffic is primarily routed to low bandwidth peripheral devices and IO devices since utilization of these devices can be low. Low utilization translates into wasted power, particularly in low power systems. For switches, low utilization can result in inefficient performance and arbitration inefficiencies.

DESCRIPTION OF THE DRAWINGS

[0008] Various embodiments are illustrated by way of example and not by way of limitation in the figures of the accompanying drawings in which like references indicate similar elements. It should be noted that references to "an," "one," "the," "other," "alternative," or "various" embodiments in this disclosure are not necessarily to the same embodiment, and such references mean at least one.

[0009] **Figure 1** is a flow chart for one embodiment of a method for combining a plurality of packets into a single packet.

[0010] **Figure 2** shows one embodiment of a packet that
5 contains multiple combined packets.

[0011] **Figure 3** is a flow chart for one embodiment of a method for delaying resynchronization of a link until combining of packets into a single packet is complete.

[0012] **Figure 4** is a diagram of a first PCI Express device
10 coupled to a second PCI Express device, each device capable of combining and partitioning packets, according to the various embodiments disclosed herein.

[0013] **Figure 5** is a system including multiple PCI Express devices capable of combining and partitioning packets.

[0014] **Figure 6** is a diagram of different data streams based
15 on the enablement of the combining function of the devices in the system of **Figure 5**.

DETAILED DESCRIPTION

[0015] In the following description, for the purposes of
20 explanation, numerous specific details are set forth in order to provide a thorough understanding of the various embodiments. It will be apparent to one skilled in the art that the embodiments may be practiced without some of these specific details. In other instances, certain structures
25 and devices are omitted or simplified in order to avoid obscuring the details of the various embodiments.

[0016] The following description and the accompanying drawings provide examples for the purposes of illustration. However, these examples should not be construed in a
30 limiting sense as they are not intended to provide an exhaustive list of all possible implementations.

[0017] Various methods disclosed herein can be implemented by using a machine to read instructions stored on a machine-readable medium and to perform functions in accordance with those instructions. A machine-readable medium includes any
5 mechanism that provides (e.g., stores and/or transmits) information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory (ROM); random access memory (RAM); magnetic
10 disk storage media; optical storage media; flash memory devices; and electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, and digital signals).

[0018] Referring now to **Figure 1**, a method for combining packets is shown. At block 10, a plurality of packets that
15 are compliant with a particular protocol (e.g., PCI Express) are combined into a single packet. The single packet is compliant with the same protocol as the plurality of combined packets.

[0019] Identification information for each of the combined
20 packets is preserved in the single packet, at block 12. In various embodiments, the preservation of identification information includes maintaining the headers of each of the combined packets within the single packet. In one
embodiment, the preservation of identification information
25 occurs simultaneously with the combining of packets.

[0020] **Figure 2** shows an example of a single packet
containing several combined packets. Specifically, single
packet 14 includes header 16, combined packets 18, and
cyclic redundancy check ("CRC") 20. In various embodiments,
30 header 16 contains an indication that single packet 14
contains combined packets 18. In other embodiments, the
indication that single packet 14 contains combined packets

18 may be located elsewhere (e.g., in other fields) within single packet 14.

[0021] Each combined packet 18 includes header 22, data field 24, and CRC 26. Combined packets 18 may be identical to each other or distinct from one another. For example, PCI Express utilizes different types of packets, which are transaction dependent, and may be combined in any manner since the various embodiments place no restrictions on the manner in which different packet types may be combined within the single packet.

[0022] Although not shown in **Figure 1**, various embodiments of the method can include additional features. For example, once the single packet is created, the single packet can be sent to a device capable of combining and partitioning packets. In one embodiment, the single packet could be sent over a PCI Express bus. The receiving device can be, for example, a PCI Express root device, a PCI Express switch, or a PCI Express peripheral device. Examples of a PCI Express peripheral device include, among others, IO devices such as Gigabit Ethernet and Accelerated Graphics Port.

[0023] If the receiving device is a switch, several different events can occur. For example, once the switch receives the single packet, the switch may simply route or send the single packet to another device. Alternatively, the switch can partition the single packet (e.g., extract the combined packets) and send the extracted packets to an appropriate destination.

[0024] In various embodiments, the switch uses the header of the single packet to determine whether to re-route the single packet or whether to begin extracting the combined packets. If the packets are extracted, the switch can use the headers of the extracted packets (e.g., that were

previously preserved) to determine where each extracted packet should be sent.

[0025] However, if the receiving device is a peripheral device, the single packet would simply be received and the
5 combined packets extracted for use/transmission by the peripheral device.

[0026] The method shown in **Figure 1** can also be extended by coordinating the combining with the resynchronization of a communications link. For example, in PCI Express, links
10 between devices can be set to power down based on some threshold criteria (e.g., predetermined period of inactivity of the link). When the link is signaled to power up in order to receive/transmit data, the link must resynchronize. Necessarily, there is a period of latency between the
15 indication that the link should be powered up and the completion of resynchronization.

[0027] Various embodiments take advantage of this inherent latency period by combining packets prior to and/or during resynchronization of the link. In other embodiments,
20 resynchronization is intentionally delayed. The additional latency caused by an intentional delay advantageously allows additional time for the combining of packets.

[0028] In this regard, **Figure 3** is a flow chart of one embodiment of a method for intentionally delaying
25 resynchronization. At block 28, available packets are combined into a single packet. At decision block 30, it is determined whether a pre-determined fixed amount of time has elapsed and/or a minimum combined packet size threshold has been satisfied. If the required condition or conditions
30 have not been satisfied, the combining of packets continues at block 28. If, however, the necessary condition or conditions have been met, resynchronization of the link

proceeds, at block 32. Once the link is resynchronized, the single packet is sent, at block 34.

[0029] **Figure 4** shows a simplified diagram of two devices that can be used to perform some of the various methods described above. Specifically, first PCI Express device 36 is coupled to second PCI Express device 38. As stated above, each device can be one of the following: a PCI Express root device, a PCI Express switch device, and a PCI Express peripheral device.

10 [0030] In one embodiment, device 36 (and possibly device 38) is capable of combining a plurality of packets that are compliant with PCI Express into a single packet that is compliant with PCI Express and contains identification information for each of the combined packets. Device 36 is also capable of sending the single packet over a PCI Express link to device 38.

[0031] Device 38 is capable of routing the single packet, received from device 36, to another PCI Express device; extracting the combined packets from the single packet; and sending the extracted packets to an appropriate destination. In various embodiments, the appropriate destination is determined by the header of the single packet and/or the header for each of the combined packets within the single packet.

25 [0032] Device 36 and device 38 each include port 35, which is a group of related transmitters and receivers. In various embodiments, ports 35 should have matching capabilities (e.g., both ports 35 should be able to combine and partition packets in both an upstream and a downstream direction, as necessary). In the embodiment shown, there are four lanes 37 between device 36 and device 38. A "lane" is a set of differential transmit and receive pairs. The

term "link" (e.g., PCI Express link) as used herein refers to two (or more) ports and their interconnecting lanes.

[0033] **Figure 5** shows a system that includes graphics controller 46, a first PCI Express device (e.g., root 40) coupled to graphics controller 46, and a second PCI Express device (e.g., switch 48). The system also includes central processing unit ("CPU") 42 coupled to root 40, memory 44 coupled to root 40, legacy endpoint 50 coupled to switch 48, and endpoint 52 coupled to switch 48. Although root 40 is shown as a single element in **Figure 5**, root 40 may comprise a memory controller hub (e.g., north bridge coupled to CPU 42, graphics controller 46, and memory 44) and an IO controller hub (e.g., south bridge coupled to the north bridge and to switch 48).

[0034] Root 40 and switch 48 are capable of implementing the methods described above (e.g., combining, partitioning, and routing packets). Likewise, endpoint 52 is also capable of implementing the previously described methods. However, legacy endpoint 50 is not capable of combining and partitioning packets.

[0035] **Figure 6** shows different data streams based on the enablement of the combining function of the devices in the system of **Figure 5**. For example, if combining is disabled on root 40, the packet stream on line 47 would appear as Stream 0 in **Figure 6** (e.g., each packet sent separately). Most likely, the separate transmission of each packet would necessitate bit stuffing, described above, to meet the minimum packet size for transmission.

[0036] However, if packet combining is enabled on root 40, switch 48, and endpoint 52, the packet stream on line 47 would appear, for example, as Stream 1 of **Figure 6**, with multiple packets appropriately combined based on size,

destination, and the combining capability (and/or enablement) of the destination device. Switch 48 routes the received packets based on their appropriate destination. In the case with packets appearing as in Stream 1, non-combined packets C and D are destined for legacy endpoint 50, and the packets combined to contain AB and EFG are destined for endpoint 52. Thus, the packet stream on line 49 would appear as Stream 2 (e.g., no combined packets) of **Figure 6**, and the packet stream on line 51 would appear as Stream 3 of **Figure 6**.

[0037] By utilizing the methods and devices disclosed herein, the efficiency of certain PCI systems can be increased between 5-10% over conventional systems that do not combine packets. Among other advantages, when the various embodiments disclosed herein are used in a chipset south bridge, port utilization is greatly increased, which permits more bandwidth to be used for high traffic peripherals in the same device.

[0038] It is to be understood that even though numerous characteristics and advantages of various embodiments have been set forth in the foregoing description, together with details of structure and function of the various embodiments, this disclosure is illustrative only. Changes may be made in detail, especially matters of structure and management of parts, without departing from the scope of the various embodiments as expressed by the broad general meaning of the terms of the appended claims.

CLAIMS

I claim:

1. A method comprising:
combining a plurality of packets that are compliant
5 with a protocol into a single packet that is compliant with
the same protocol;
preserving identification information for each of the
combined packets within the single packet; and
adding an indication to the single packet to indicate
10 that the combined packets are contained within the single
packet.
2. The method of Claim 1, wherein adding comprises:
attaching a header to the single packet.
3. The method of Claim 1, further comprising:
15 sending the single packet to a device capable of
combining and partitioning packets.
4. The method of Claim 3, further comprising:
receiving the single packet; and
sending the single packet to another device.
- 20 5. The method of Claim 3, further comprising:
receiving the single packet; and
extracting the combined packets.
6. The method of Claim 5, further comprising:
sending each extracted packet to an appropriate
25 destination.
7. The method of Claim 1, wherein combining is
performed during at least one of a time (i) prior to
resynchronization of a Peripheral Component Interconnect
("PCI") Express link, and (ii) during resynchronization of a
30 PCI Express link.

8. The method of Claim 1, further comprising:
delaying resynchronization of a Peripheral Component Interconnect Express link based on at least one of a fixed amount of time and a minimum combined packet size threshold.

5 9. A method comprising:
receiving a single packet that is compliant with a protocol, the single packet containing a plurality of combined packets and a header to indicate that the combined packets are contained within the single packet, each
10 combined packet being compliant with the same protocol and having identification information associated therewith.

10. The method of Claim 9, further comprising:
routing the single packet to a device coupled to a bus.

15 11. The method of Claim 9, further comprising:
extracting the combined packets from the single packet.

12. The method of Claim 11, further comprising:
sending the extracted packets to an appropriate destination.

20 13. An apparatus comprising:
a first Peripheral Component Interconnect ("PCI") Express device capable of
combining a plurality of packets that are compliant with PCI Express into a single packet that is compliant with PCI Express, the single packet to contain
25 identification information for each of the combined packets within the single packet and a header to indicate that the combined packets are contained within the single packet, and
sending the single packet over a PCI Express link;
and
30 a second PCI Express device to receive the single packet.

14. The apparatus of Claim 13, wherein the first PCI Express device comprises one of a PCI Express root device, a PCI Express switch device, and a PCI Express peripheral device, and the second PCI Express device comprises one of a
5 PCI Express root device, a PCI Express switch device, and a PCI Express peripheral device.

15. The apparatus of Claim 13, wherein the second PCI Express device is capable of at least one of (i) routing the single packet to another PCI Express device, (ii) extracting
10 the combined packets from the single packet, and (iii) sending the extracted packets to an appropriate destination.

16. A system comprising:
a graphics controller;
a first Peripheral Component Interconnect ("PCI")
15 Express device coupled to the graphics controller, the first PCI Express device capable of
combining a plurality of packets that are compliant with PCI Express into a single packet that is compliant with PCI Express, the single packet to contain
20 identification information for each of the combined packets within the single packet and a header to indicate that the combined packets are contained within the single packet, and
sending the single packet over a PCI Express link;
and
25 a second PCI Express device to receive the single packet.

17. The system of Claim 16, wherein the first PCI Express device comprises one of a PCI Express root device, a PCI Express switch device, and a PCI Express peripheral
30 device, and the second PCI Express device comprises one of a PCI Express root device, a PCI Express switch device, and a PCI Express peripheral device.

18. The system of Claim 16, wherein the second PCI Express device is capable of at least one of (i) routing the single packet to another PCI Express device, (ii) extracting the combined packets from the single packet, and (iii) sending the extracted packets to an appropriate destination.

19. A machine-readable medium containing instructions that, when executed by a machine, cause the machine to perform operations comprising:

combining a plurality of packets that are compliant with a protocol into a single packet that is compliant with the same protocol;

preserving identification information for each of the combined packets within the single packet; and

adding an indication to the single packet to indicate that the combined packets are contained within the single packet.

20. The machine-readable medium of Claim 19, wherein adding an indication comprises:

attaching a header to the single packet.

21. The machine-readable medium of Claim 19, wherein the operations further comprise:

sending the single packet to a device capable of combining and partitioning packets.

22. The machine-readable medium of Claim 21, wherein the operations further comprise:

receiving the single packet; and

sending the single packet to another device.

23. The machine-readable medium of Claim 21, wherein the operations further comprise:

receiving the single packet; and

extracting the combined packets based on the identification information.

24. The machine-readable medium of Claim 23, wherein the operations further comprise:

 sending each extracted packet to an appropriate destination.

5 25. The machine-readable medium of Claim 19, wherein combining is performed during at least one of a time (i) prior to resynchronization of a Peripheral Component Interconnect ("PCI") Express link, and (ii) during resynchronization of a PCI Express link.

10 26. The machine-readable medium of Claim 19, wherein the operations further comprise:

 delaying resynchronization of a Peripheral Component Interconnect Express link based on at least one of a fixed amount of time and a minimum combined packet size threshold.

15




FIG. 1




FIG. 2

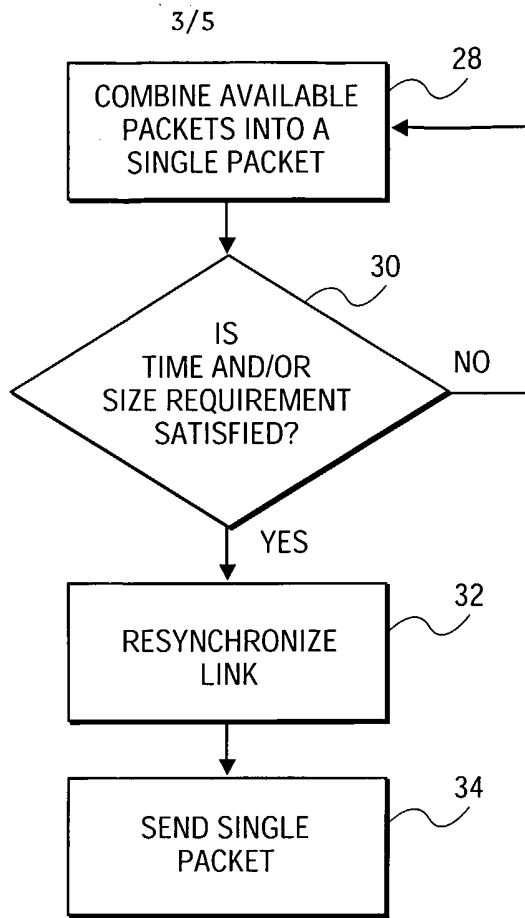


FIG. 3




FIG. 4




FIG. 5




FIG. 6