(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
30 May 2002 (30.05.2002)

PCT

(10) International Publication Number
WO 02/42917 A2

(51) International Patent Classification[7]: G06F 13/00

(21) International Application Number: PCT/GB01/05138

(22) International Filing Date:
21 November 2001 (21.11.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
GB 0028353.1 21 November 2000 (21.11.2000) GB

(71) Applicant *(for all designated States except US)*: **ASPEX TECHNOLOGY LIMITED** [GB/GB]; Sanderson House, Oxford Road, Denham, Buckinghamshire UB9 4DX (GB).

(72) Inventor; and

(75) Inventor/Applicant *(for US only)*: **WHITAKER, Martin** [GB/GB]; Aspex Technology Limited, Sanderson House, Oxford Road, Denham, Buckinghamshire UB9 4DX (GB).

(74) Agents: **AHMAD, Sheikh, Shakeel** et al.; David Keltie Associates, 12 New Fetter Lane, London EC4A 1AG (GB).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent

[Continued on next page]

(54) Title: BROADCASTING DATA ACROSS A BUS

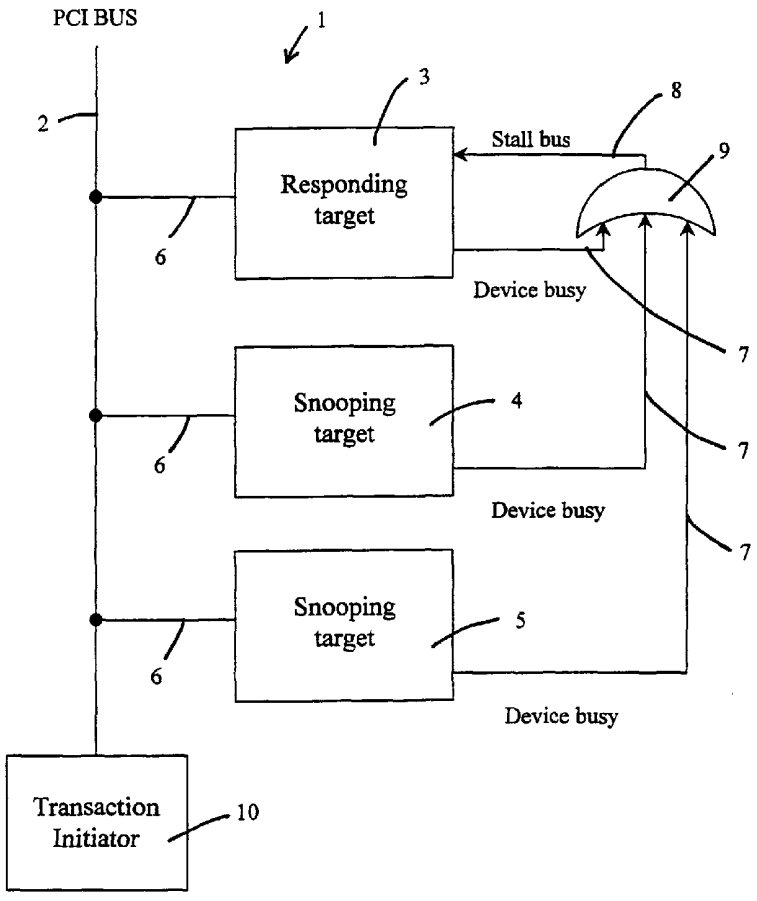

(57) Abstract: A method of broadcasting data to multiple targets across a systembus (2), such as the peripheral component interconnect. (PCI) bus, that does not normally support broadcast transfers wherein one target (3) respnds to the bus transaction (20, 24, 28) and the remaining targets (4, 5) listen in (snoop) on the bus transaction to receive data from the system bus (2), the method comprising theresponding target (3) stalling (38) or slowing down the bus transaction, or forcing the re-sending of already transmitted data, when any of the listening (snooping) targets (4, 5) communicate (32) to the responding target (3) that they are temporarily unable to accept the data on the bus.

WO 02/42917 A2

(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.

Broadcasting data across a bus

This invention relates generally to the transfer of data across a bus within an electronic system, and more specifically to a method of broadcasting data to multiple targets on a bus, such as the peripheral component interconnect (PCI) bus, that does not normally support broadcast transfers.

Modern electronic systems make use of connecting many components to a central bus, commonly known as a databus, so that data can be efficiently moved between components such as the processor, memory and peripherals. At their conception, databuses were envisaged as being used to move data between one transaction initiator and one responding target at a time. Initially, with computerised systems consisting of only one processing unit and few peripheral components, this proved to be an adequate solution to the problem of allowing the processing unit and peripheral components to transfer data and communicate.

With the advent of modern electronic components running at much higher speeds than when the original specifications for many databuses were drawn up, and the development of many more types of components able to connect to these databuses, the efficient use of the limited transmission bandwidth along databuses is essential in achieving good overall system performance in computerised systems. Examples of high-performance plug-in components which are compatible with the PCI databus standard include video capture cards, Double Data Rate (DDR) memory technologies, IEEE 1394 communication cards, cable and xDSL modems and LAN cards based on optical or Fast Ethernet protocols. Equally compatible are surface mounted components present on motherboards and daughterboards of computer systems. Many of these types of PCI databus compatible components require heavy bandwidth usage.

A common limiting factor for overall system performance of computerised

systems are bottlenecks caused when databus bandwidth is not enough to accommodate all the data that needs transporting simultaneously. For example, when large amounts of information are frequently transferred to and from memory, the memory databus bandwidth can quickly become saturated leading to a cap on performance. This problem is exacerbated in high performance computer systems, particularly within systems that contain more than one processing unit wherein a transaction initiator may be transmitting data to one or more responding targets simultaneously.

In these situations data broadcasting (i.e. the simultaneous transfer of data from a single source location to multiple destinations) is often required. If the system databuses used in these high-performance systems do not support broadcast transfers then the data has to be individually and sequentially transferred to each of the required responding targets, unnecessarily consuming bus bandwidth and increasing both the time taken and the number of processor cycles needed to complete the task. These are detrimental factors to the performance of the system and need to be minimised.

Unfortunately, many of the standard databuses used within computerised systems do not support broadcast data transfers. Some of these databuses do provide a simple message broadcasting scheme, but this has been found in practice to be an impractical method for transferring large blocks of data. For example, the databus most commonly and widely used in computerised systems is the PCI bus, and the specification for this databus provides for a 'Special Cycle' transaction which allows a message to be broadcast to all components on the bus, but this message can only consist of a single data item, with no associated addressing information. An example of when a 'Special Cycle' transaction may be used is in sending power management directives to all the components on the bus. This type of transaction transmits very low bandwidth information and does not allow burst transfers, hence it is unsuitable for transmitting usable data blocks in real-time applications. Also, because no

handshaking is performed on this transaction, there is no way to stall the data transaction if a target is temporarily unable to accept the data, for example, if the receiving data buffer is full.

The only other way in which the PCI databus specification allows more than one component to listen in on data not being directly transmitted to that component is via 'VGA palette snooping'. This is a means of having a passive listener on the PCI bus listen in, or 'snoop', on the data being transmitted across the PCI bus and is used when special VGA video display cards and other high-end display-related components, such as MPEG decoder cards, need to be able to look at the video card's VGA palette to determine what colours are currently in use. This feature is only rarely used and is also unsuitable for transmitting large blocks of data to multiple components at speed as only a restrictive address range (only of interest to display-related components) is passed along. Again, no handshaking is present in "VGA palette snooping" and a further drawback is present in that an assumption is made that the passive listener is always at least as fast as the intended responding target in being able to receive the transmitted data, otherwise the passive listener may miss some information.

In many cases the length of the circuit path between a transaction initiator and the intended multiple responding targets are different, perhaps the data having to pass through bridges and other connectors before reaching certain responding targets, and this can cause timing delays between different responding targets receiving the same broadcast data. The broadcasting system has to be flexible enough to accommodate any lags in the responding targets receiving data by varying or even pausing the data transmission.

Examples of prior art that relates to PCI bus broadcasting transactions are provided by US Patent Nos. 5507002, 5634138, 5983024 and 6230225.

US 5507002 to AT&T Global Information Solutions Company describes a method of allowing sideband signalling to be made available to standard PCI compliant components, and also describes part of the PCI specification relating to "Special Cycle" transactions, and describes the limitations of using this type of transaction.

US 5634138 to Emulex Corporation concerns a method of burst broadcasting data from a transaction initiator to multiple responding targets over a PCI bus, but the teaching in this specification is to modify both the transaction initiator and the responding target which can be costly and impractical. This method requires extra PCI bus transactions to be performed to initiate and terminate broadcast operations, and responding targets cannot respond that they are not ready to accept data.

US 5983024 to Honeywell, Inc. relates to a method of robust data broadcasting on a PCI bus and it requires that both the transaction initiator and the responding targets are hardware-modified in accordance with the teachings of the specification. The responding targets have no way of signalling to the transaction initiator that they are not ready to accept data, and this method only works if the snooping targets are capable of accepting the data transacted across the PCI bus at whatever speed the transaction initiator and intended responding target are communicating at.

US 6230225 to Compaq Computer Corp. teaches how broadcasting can be achieved with components connected to a PCI bus through the introduction of a separate multicast bus dedicated to informing components of when to listen in on PCI bus transactions. This method requires changes to the hardware of all components and the introduction of a separate dedicated multicast bus. In addition, only the intended responding target can signal that it is not ready for receiving data.

From the prior art it is demonstrated that for a databus broadcasting system to work effectively there must be provided a way for the responding target or snooping targets, to be able to at least suspend or even slow down or force resends on any of the data transactions being carried out on the databus. It is also evident that any system which requires significant changes to be made to existing hardware in order to be implemented will not be readily adopted when there is such a large existing user base of unmodified components.

It is with a view to solving these problems that there is provided a method of broadcasting data to multiple targets across a system bus, such as the peripheral component interconnect (PCI) bus, that does not normally support broadcast transfers, in which one target responds to the bus transaction and the remaining targets listen in on the bus transaction to receive data from the system bus, the method comprising the responding target stalling the bus transaction when any of the listening targets communicate to the responding target that they are temporarily unable to accept the data on the bus.

The present invention allows use of the existing handshaking communications established between the transaction initiator and the responding target to signal stall bus commands from any one or more of the snooping targets or the responding target itself. As a consequence, the integrity of the data received by every target on the PCI bus is not compromised through a target not being ready to receive the data transmitted through the PCI bus. The benefit of using existing standard hardware specifications to provide the conduit for signaling is that robust data broadcasting according to the present invention can be carried out cheaply and quickly, as no modifications to the PCI bus connectors or associated and supporting components, except for the targets themselves, are needed.

Put in another way, what is proposed is a method of broadcasting data to multiple targets across a data bus, such as the peripheral component interconnect

(PCI) bus, that does not normally support broadcast transfers, in which one target responds to the bus transaction and another target snoops on the bus transaction, the method comprising the snooping target transmitting a device busy signal to the responding target when it is temporarily unable to snoop on the data carried by the data bus and the responding target responding to the device busy signal by stalling the data transfer on the data bus.

Preferably, the transmitting step comprises any one of a plurality of snooping targets generating and transmitting a device busy signal to the responding target for stalling the data bus.

In a further preferred feature the method comprises the step of Oring together all of the signals from the plurality of snooping targets to provide a single bus stall trigger signal. This provides a simple low-cost way of logically combining all of the device busy signals into a single bus stall trigger signal which only requires a single logic gate.

In order to react to the responding target or any of the snooping targets not being able to receive data, the method preferably further comprises the responding target generating a temporary device busy signal and using the same to trigger stalling of the data bus.

So that this method may be adopted with minimal changes to the existing hardware and software in use already, it is advantageous that the method as previously described comprises any of the snooping targets communicating with the responding target via an unused control channel of the data bus. This again advantageously minimises cost and complication as existing connections between the target and snooping targets are used.

Once the responding target and/or snooping targets are able to receive data

again, it is advantageous that the method comprises a responding step that stalls the system bus until such time as the device busy signal is negated.

In an alternative aspect, the invention also resides in a data transfer apparatus for broadcasting data to multiple targets across a system bus, such as the peripheral component interconnect (PCI) bus, that does not normally support broadcast transfers, the apparatus comprising a responding target connected to the system bus for responding to a data transfer and at least one snooping target connected to the system bus for listening in on data transfers on the system bus, and processing means provided at the responding target for receiving device busy signals from the at least one snooping target, the processing means being arranged to stall the system bus in response to receipt of a device busy signal.

Phrased differently, the invention could also be said to be a method of broadcasting data simultaneously from a transaction initiator to a responding target and to at least one other target across a databus, the method comprising: snooping on the data being transmitted across the databus by the one other target; generating a busy signal if the one other target is unable to receive the data being transmitted across the databus; and the transaction initiator responding to the generation of a busy signal by stalling the transmission of data across the databus until the busy signal is removed.

The present invention can also be considered to be a system for broadcasting data comprising a transaction initiator, a responding target and at least one other target all on the same databus wherein: the transaction initiator is arranged to broadcast data across the databus and the at least one other target comprises means for snooping on the data; each of the at least one other target comprises means for generating a busy signal if the target is unable to receive the data being sent from the transaction initiator; and the transaction initiator comprises means for responding to

the generation of a busy signal by stalling the transmission of data across the databus until the busy signal is cleared.

In a preferred feature, the transaction initiator further comprises means to vary the data transmission rate across the databus down to zero. By being able to vary the speed of the data transmission to at least one speed between the maximum or normal speed and zero, any target that falls behind in processing data received from the data bus may have a chance to catch up without having to totally halt the transmission process and causing further delay.

It is also envisaged that for such a system it is advantageous to allow for the data receiving rate to be different for each target so that robust data broadcasting according to the present invention may be implemented with systems containing PCI bus compatible targets of differing specifications, allowing for maximum design flexibility.

The present invention is now described with reference to the accompanying figures wherein:

Figure 1 is a block and circuit diagram of a communications apparatus embodying the present invention; and

Figure 2 is a flow diagram showing the steps implied in implementing a method of broadcasting data using the apparatus of Figure 1 according to the embodiment of the present invention.

In Figure 1 there is illustrated a functional block and circuit diagram of a communications apparatus, shown generally at 1, for effecting broadcast of data across a databus. In this embodiment a databus complying with the PCI specification is represented at 2, although this system is equally applicable to any databus suitable

for use in computerised systems. Data signals are sent along the bus 2 from a transaction initiator 10 addressed to an intended responding target 3, the responding target being in the form of a plug-in PCI board. Also connected to the same bus 2 are other snooping targets 4 and 5, both also in the form of further PCI plug-in boards. Two snooping targets are shown in Figure 1 although the number of snooping targets may be as many as required. The responding target and each snooping target are connected to the PCI bus through standard PCI plug-in board connectors 6, although this method is equally applicable to components connected to the PCI bus through traces on a circuit board. For example, surface mounted components may be connected to the PCI bus in this method whilst still complying with the inventive concept.

Provided with the responding target are two connectors, one for connecting a DEVICE BUSY line 7 and one for connecting a STALL BUS line 8. Each snooping target also possesses a connector for connecting a DEVICE BUSY line 7. Each DEVICE BUSY line 7 feeds from the responding and snooping targets into a logical OR gate 9, and the output of logical OR gate 9 feeds into the STALL BUS connector on the responding target through STALL BUS line 8.

Referring now to Figure 2, a method of using the above-described communications apparatus is now described. When data is to be transmitted across the PCI bus 2, before the main data itself is sent at step 20, there is first sent addressing data which signals the responding target 3 to be ready to accept the following main data. At the same time, every other snooping target 4, 5 also responds at step 22 by getting ready to accept the following data as well. When handshaking has been established at step 24 between the responding target 3 and the transaction initiator 10, the main data is then transmitted at step 28 and is received at step 30 by the responding target and all the snooping targets.

At any time during transmission of the main data should any one or more of the responding 3 or snooping 4,5 targets not be able at step 30 to receive data at the rate at which it is being sent, an output is sent at step 32 along the DEVICE BUSY line 7 of the responding 3 or snooping target 4, 5 which has the problem, and this causes the STALL BUS line 8 state to change at step 34 due to the logical OR gate 9 receiving an input. On detecting at step 36 the change in the status of the STALL BUS line, the responding target 3 signals at step 38 to the transaction initiator 10 to pause the data transmission until all of the responding and snooping target(s) which have difficulties are able to continue receiving data again. The responding target 3 signals the transaction initiator 10 through the normal control lines of the PCI bus 2, and because of this, there is no need to modify the transaction initiator 10 either in terms of hardware or software, nor is there any need to modify the BIOS settings of the main board upon which the PCI bus 2 is situated. On receipt of the signal from the responding target 3, the transaction initiator 10 stalls at step 40 the PCI bus 2.

In this embodiment, the STALL BUS signal prompts a pause in the data transmission, but it is envisaged in alternative embodiments (not shown) that the data transmission may also be re-sent or slowed down according to different types of signals sent from the responding target. One way in which these alternative embodiments may be put into effect is to have separate outputs on the responding and snooping targets, similar to the DEVICE BUSY outputs, for indicating that the responding or snooping target requires data to be re-sent or the transmission rate to be slowed down. This can be achieved using components and techniques that are well known to the skilled addressee. The responding target 3 would then pass on the required type of signal to the transaction initiator 10, which would then act accordingly to re-send or slow down the data transmission.

In Figure 1 the STALL BUS line 8 and DEVICE BUSY line 7 are both separate hardware components on the PCI plug-in boards. An alternative embodiment (not shown) uses unused pins (to unused lines of the PCI bus) in the PCI

bus connector specification to act as the STALL BUS line and the DEVICE BUSY line, thus minimising the changes needed to implement the present invention on existing systems.

The embodiment of Figure 1 shows one responding target and one STALL BUS line. In a further alternative embodiment (not shown), any or all of the snooping targets may also act as a responding target. In this case, all of the snooping targets 4 shown in Figure 1 are provided with their own STALL BUS line inputs from the logical OR gate, so that should they in turn be the responding target they will also be able to signal to the transaction initiator to pause, re-send or slow down the data transmission.

Of course, in any one data transaction along the PCI databus 2 there is only one responding target and if the responding target is different to the previous responding target, the previous responding target may change to being a snooping target.

In a yet further embodiment (not shown), any one out of the responding or snooping targets may become the transaction initiator, responding to signals to pause, re-send or slow the data transmission from the other responding or snooping targets. The PCI bus specification also allows the same component to act as both the transaction initiator as well as a responding target at the same time.

It will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, although the present invention has been illustrated in the context of a PCI bus, the same concept works with any bus that lacks a dedicated burst broadcast mode but does permit two-way communication between the transaction initiator and at least one responding target. Accordingly, it is to be understood that the invention is not to be limited to the specific illustrated embodiment, but only by the scope of the appended claims.

Claims

1. A method of broadcasting data to multiple targets across a system bus, such as the peripheral component interconnect (PCI) bus, that does not normally support broadcast transfers, in which one target responds to the bus transaction and the remaining targets listen in on the bus transaction to receive data from the system bus, the method comprising the responding target stalling the bus transaction when any of the listening targets communicate to the responding target that they are temporarily unable to accept the data on the bus.

2. A method of broadcasting data to multiple targets across a data bus, such as the peripheral component interconnect (PCI) bus, that does not normally support broadcast transfers, in which one target responds to the bus transaction and another target snoops on the bus transaction, the method comprising the snooping target transmitting a device busy signal to the responding target when it is temporarily unable to snoop on the data carried by the data bus and the responding target responding to the device busy signal by stalling the data transfer on the data bus.

3. A method according to Claim 2, wherein the transmitting step comprises any one of a plurality of snooping targets generating and transmitting a device busy signal to the responding target for stalling the data bus.

4. A method according to Claim 3, further comprising Oring together all of the signals from the plurality of snooping targets to provide a single bus stall trigger signal.

5. A method according to any of Claims 2 to 4, further comprising the responding target generating a temporary device busy signal and using the same to trigger stalling of the data bus.

6. A method according to any of Claims 2 to 5, further comprising any of the

snooping targets communicating to the responding target via an unused control channel of the data bus.

7. A method according to any of Claims 2 to 6, wherein the responding step comprises stalling the system bus until such time as the device busy signal is negated.

8. A data transfer apparatus for broadcasting data to multiple targets across a system bus, such as the peripheral component interconnect (PCI) bus, that does not normally support broadcast transfers, the apparatus comprising a responding target connected to the system bus for responding to a data transfer and at least one snooping target connected to the system bus for listening in on data transfers on the system bus, and processing means provided at the responding target for receiving device busy signals from the at least one snooping target, the processing means being arranged to stall the system bus in response to receipt of a device busy signal.

9. A method of broadcasting data simultaneously from a transaction initiator to a responding target and to at least one other target across a databus, the method comprising:

snooping on the data being transmitted across the databus by the one other target;

generating a busy signal if the one other target is unable to receive the data being transmitted across the databus;

and the transaction initiator responding to the generation of a busy signal by stalling the transmission of data across the databus until the busy signal is removed.

10. A system for broadcasting data comprising a transaction initiator, a responding target and at least one other target all on the same databus wherein:

the transmitting is arranged to broadcast data across the databus and the at least one other target comprises means for snooping on the data;

each of the at least one other target comprises means for generating a busy signal if the target is unable to receive the data being sent from the transaction initiator;

and the transaction initiator comprises means for responding to the generation of a busy signal by stalling the transmission of data across the databus until the busy signal is cleared.

11. A system according to Claim 10 wherein the transaction initiator further comprises means to vary the data transmission rate across the databus down to zero.

12. A system according to Claim 10 wherein the data receiving rate is different for each target.

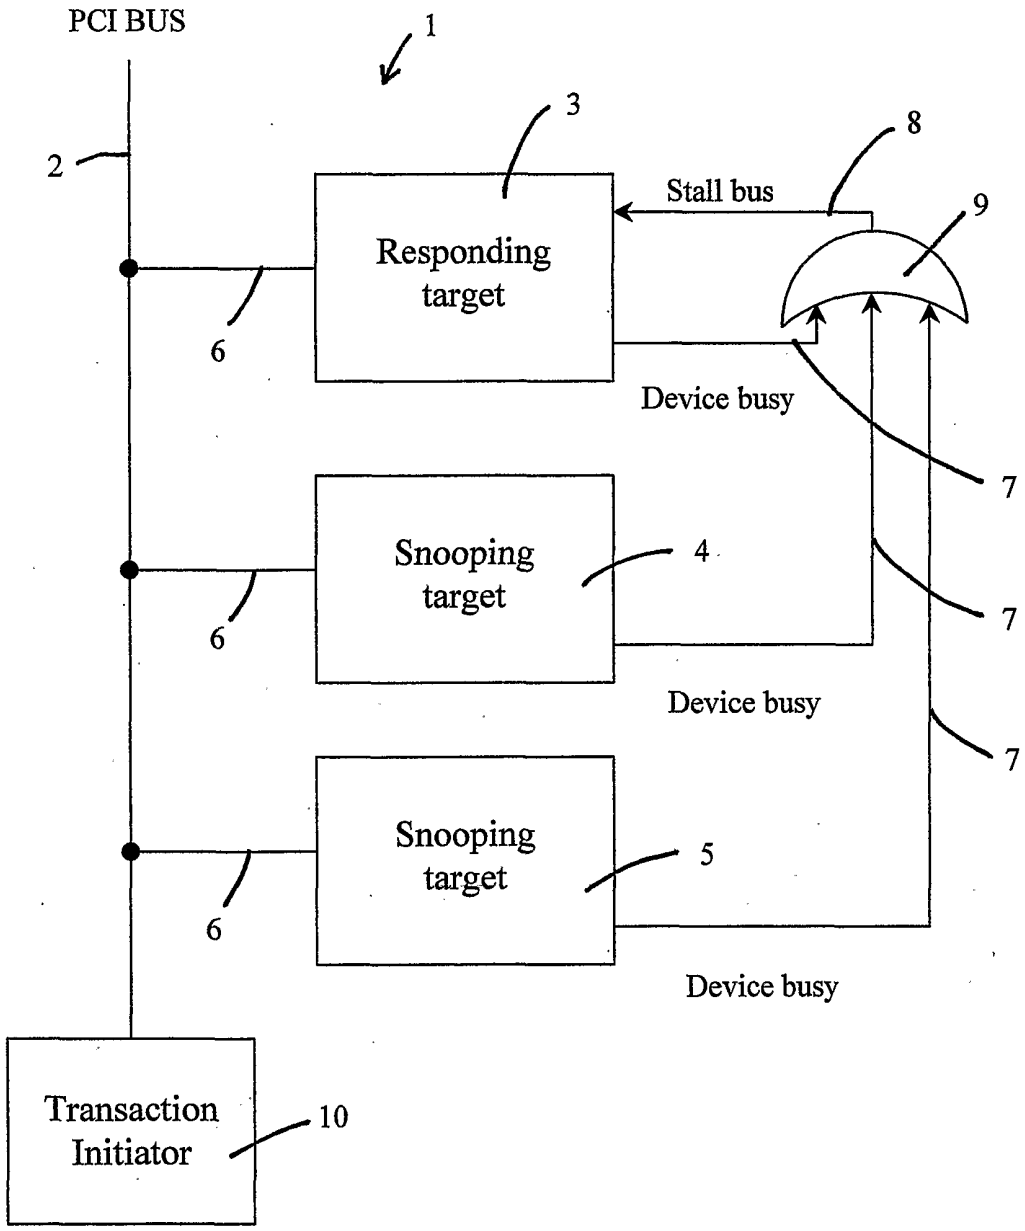
PCI BUS
1
2
3
8
9
Stall bus
Responding target
6
Device busy
7
Snooping target
4
6
7
Device busy
7
Snooping target
5
6
Device busy
Transaction Initiator
10

Figure 1

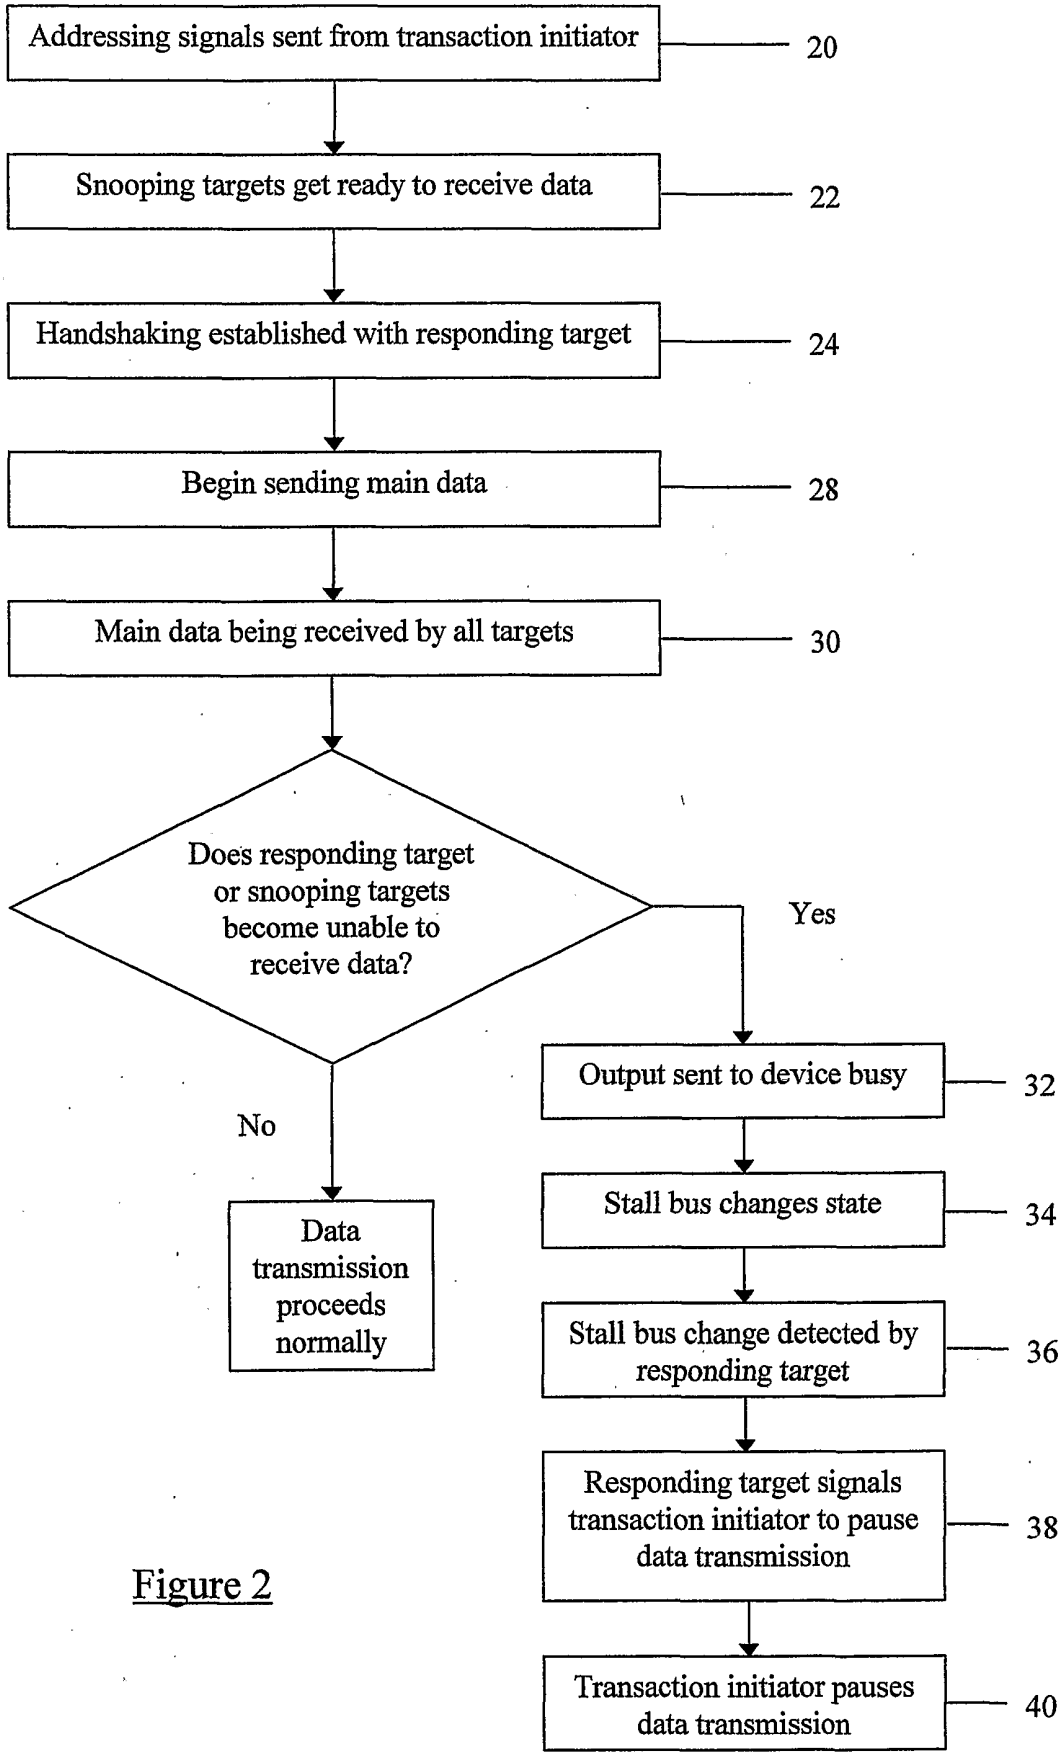
Addressing signals sent from transaction initiator
20
Snooping targets get ready to receive data
22
Handshaking established with responding target
24
Begin sending main data
28
Main data being received by all targets
30
Does responding target or snooping targets become unable to receive data?
Yes
No
Data transmission proceeds normally
Output sent to device busy
32
Stall bus changes state
34
Stall bus change detected by responding target
36
Responding target signals transaction initiator to pause data transmission
38
Transaction initiator pauses data transmission
40

Figure 2