

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
18 November 2004 (18.11.2004)

PCT

(10) International Publication Number
WO 2004/100005 A1

(51) International Patent Classification⁷: **G06F 13/36**,
13/42

(21) International Application Number:
PCT/IB2004/050580

(22) International Filing Date: 4 May 2004 (04.05.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
03101266.9 8 May 2003 (08.05.2003) EP

(71) Applicant (for all designated States except US): **KONINKLIJKE PHILIPS ELECTRONICS N.V.** [NL/NL];
Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **RADULESCU, Andrei** [RO/NL]; c/o Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). **GOOSSENS, Kees, G., W.** [NL/NL];
c/o Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).

(74) Agent: **DE JONG, Durk, J.**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).

(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).

Declaration under Rule 4.17:


— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations 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, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (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 patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

Published:

— with international search report

[Continued on next page]

(54) Title: PROCESSING SYSTEM AND METHOD FOR COMMUNICATING DATA



(57) Abstract: A method for communicating data between an initiator unit (INIT) which initiates the communication and a target unit (TRGT) is described. Therein the initiator unit (INIT) indicates a request (TID) to initiate a communication. In response the target unit (TRGT) provides information (READY, ACCEPTC) indicating whether one of the following situation exist, the initiator unit (INIT) has to maintain the request, the request of the initiator (INIT) is accepted, the request of the initiator (INIT) is rejected. In addition a processing system is described.



-
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*
- 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.*

Processing system and method for communicating data

The invention relates to a processing system.

The invention further relates to a method for communicating data.

Modern on-chip interconnect protocols such as VCI and OCP introduce communication-thread identifiers to allow multiple unrelated transactions to evolve independently. On each communication thread all the transactions are ordered (i.e., requests are delivered at the slave, responses are issued by the slaves, and responses are delivered to the master in the same order the requests were issued by the master). Transactions with different communication threads do not have any ordering constraints. This allows unrelated communications to proceed independently of each other.

These protocols are defined as split protocols in that transactions are split in a request and a response. After a transmission of a request is completed from a source functional unit to the first intermediate functional unit in a communication path the source functional unit can proceed with a next transmission, instead of having to wait for a response to that request from the destination functional unit. The destination or any intermediate functional unit will start a separate arbitration procedure if necessary to give a response. A split bus protocol is more efficient when a response generation at the slave takes time (because it is pipelined, for example) or if the delivery of the request from the master to the slave takes time (as is often the case in networks). The protocols further allow pipelining, i.e. allow a master to have multiple outstanding requests (i.e., requests waiting for a response). All transactions within the same communication thread are ordered. This implies that requests are executed by the slave in the same order as the requests for those responses were issued by a master, and responses are delivered in the same order as the requests for those responses were issued by a master. Transactions with different communication threads do not have any ordering constraints.

25

US 6,182,183 provides a link level protocol for exchanging the message between two subsequent functional units in the communication path from the source functional unit to the destination functional unit. According to the known protocol a master

functional unit produces information, e.g. a command (Cmd), an address (Addr), or data (DataReq) and at the same time provides an identification of the thread (ReqThreadID) to which the information belongs. Likewise the slave functional unit may provide information (DataResp), and indicate the communication thread to which it belongs by an identification
5 RespThreadID.

A functional unit can be any unit involved in a data stream for example a unit which performs operations on data, such as a CPU, a DSP or a VLIW, or a unit for storing data such as a memory, or a unit for transmitting data such as a router or an interface.

In the sequel the following denotations will be used. An initiator functional
10 unit or initiator is a functional unit which initiates a link-level transaction with another functional unit referred to as target functional unit. According to the OSI definition at the link level, data packets are encoded and decoded into bits. It furnishes transmission protocol knowledge and management and handles errors in the physical layer, flow control and frame synchronization. A source functional unit is denoted as a functional unit which at the network
15 level sends a message, e.g. data or a command to another functional unit, denoted as the destination functional unit. According to the OSI-definition the network layer provides switching and routing technologies, creating logical paths, known as virtual circuits, for transmitting data from node to node. Routing and forwarding are functions of this layer, as well as addressing, internetworking, error handling, congestion control and packet
20 sequencing. A functional unit may operate as a source functional unit or as a destination functional unit depending on the circumstances.

It is a purpose of the invention to improve the known protocol. According
25 therewith a method as defined in claim 1 and a processing system as defined in claim 6 are provided.

According to the invention the target functional unit receiving the request from the initiator functional unit to perform an action relating to a particular thread has the options to signal the initiator functional unit that

- 30
1. the initiator unit (INIT) has to maintain the request,
 2. the request of the initiator (INIT) is accepted,
 3. the request of the initiator (INIT) is rejected.

When the target functional unit rejects the request according to option 3, the initiator is enabled to issue a new request for another or the same communication thread. If

this is accepted, a data transmission can take place which otherwise would have been delayed. In this way overall throughput is improved.

The invention is particularly suitable for application in a processor network for the following reasons. Systems on silicon show a continuous increase in complexity due to the ever increasing need for implementing new features and improvements of existing functions. This is enabled by the increasing density with which components can be integrated on an integrated circuit. At the same time the clock speed at which circuits are operated tends to increase too. The higher clock speed in combination with the increased density of components has reduced the area which can operate synchronously within the same clock domain. This has created the need for a modular approach. According to such an approach the processing system comprises a plurality of relatively independent, complex modules. In conventional processing systems the modules usually communicate to each other via a bus. As the number of modules increases however, this way of communication is no longer practical for the following reasons. On the one hand the large number of modules forms a too high bus load. On the other hand the bus forms a communication bottleneck as it enables only one device to send data to the bus. A communication network forms an effective way to overcome these disadvantages. The communication network comprises a plurality of partly connected functional units. Messages from a functional unit are redirected by the functional units to one or more other functional units. Between two directly connected functional units messages relating to different communication threads may be transmitted originating from plural sources and going to plural destinations. According to the method and processing system of the invention the target functional unit has an efficient way to inform the initiator functional unit that a request for a certain thread can not be handled for some time. This enables the initiator functional unit to select another communication thread for which data transmission otherwise should have been unnecessarily delayed.

A message sent by a source functional unit may comprise a command or a packet of data. It is forwarded via one or more intermediate functional units until it arrives at the destination functional unit. The destination functional unit may on its turn send a message to the source functional unit.

30

These and other aspects of the invention are described in more detail with reference to the drawing. Therein

Figure 1 shows a processing system comprising plurality of functional units which are coupled in a network,

Figure 2 shows a typical way of coupling two functional units in a network according to the prior art,

5 Figure 3 shows a first and a second mutually coupled functional unit in a first embodiment of the invention,

Figure 4 shows a first and a second mutually coupled functional unit in a second embodiment of the invention.

10

Figure 1 schematically shows a data processing system, which comprises a network connecting a plurality of functional units. The processing system is arranged to transmit data and a communication thread identifier for said data according to a split protocol along a communication path (indicated by arrows) from a source functional unit SFU to a
15 destination functional unit DFU via one or more intermediate functional units IFU1, ..., IFU5. By transmitting the data together with a communication thread identifier, multiple unrelated transactions, having mutually different communication thread identifiers can evolve independently.

Figure 2 shows a typical way of coupling two functional units in a network
20 according to the prior art. Therein a data producing functional unit PFU issues a request to a data consuming functional unit CFU by activating a signal TID. This signal is indicative for the communication thread for which it attempts to transfer data. The data to be transferred is offered at a connection DATA. Validity of the indicated communication thread and the related data is indicated with a further signal VALID. As soon as the data consuming
25 functional unit has handled the request it indicates this with a signal ACCEPTC. Until the data producing functional unit has received this signal no further transmission can take place between the two functional units. It may occur however that a request of the data producing functional unit PFU cannot be handled in time by the data consuming functional unit CFU, for example if the data consuming functional unit CFU has no buffer space left for said
30 thread. This has the consequence however that data transmission between those functional units for all other communication threads are also delayed. In the embodiment shown it is supposed that the initiator is only capable to send data for a particular communication thread. The initiator can indicate this with the TID signal. In practice the initiator may be capable of issuing a plurality of commands, e.g. write, read, conditional write, etc. One or more

command signals may be added to enable the initiator to do so, or alternatively the command may be transmitted time multiplexed via another connection, e.g. via the connection used by the TID signal. It is also possible to send only commands.

Figure 3 shows a first and a second mutually coupled functional units in a first embodiment of a network according to the invention. A data transmission is initiated in the same way as was illustrated with reference to Figure 2, i.e. the data producing functional unit PFU issues a request to the data consuming functional unit CFU by indicating a communication thread with a signal TID. The data to be transferred is offered at a connection DATA. Validity of the indicated communication thread and the related data is indicated with a further signal VALID. Contrary to the situation illustrated in Figure 2 however the target functional unit, here the data consuming functional unit, has the following options:

1. Indicating the initiator unit INIT that it has to maintain the request.
2. Indicating that the request of the initiator INIT is accepted.
3. Indicating that the request of the initiator INIT is rejected.

In practice the target functional unit needs time to interpret the request and to determine whether it can be handled at that moment or not. During that time the target functional unit provides an indication according to option 1. This causes the target functional unit to hold the signals TID, DATA and VALID. As soon as the target functional unit has interpreted the request and can handle it, it indicates according to option 2. However if the target functional unit determines that the request can not be handled at that time, for example because a buffer overflow would occur for said communication thread, it signals according to option 3. This allows the initiator to withdraw the request for said communication thread and initiate a transaction for another communication thread which otherwise would have to wait until the earlier request could be handled. In this way the overall communication in the processing system is improved.

In the embodiment of Figure 3 there is one additional signal that encodes that a link-level data exchange can proceed or it is delayed. With only one additional signal, there is only a proceed/delay feedback possible. This can be generalized to encode more elaborate feedback on why a transaction cannot proceed on a communication thread (when there are multiple causes possible: e.g., empty/full buffers, a process not expecting data on a communication thread is running on a CPU, etc), or how long transactions can proceed on a communication thread (e.g., there is enough buffering/data to proceed with at least N transactions on a communication thread).

It is noted that a request for a particular communication thread may involve a request for sending commands, addresses and/or other data. Alternatively it may involve a request for receiving commands, addresses and/or other data.

Figure 4 shows a first and a second mutually coupled functional units in a second embodiment of a network according to the invention. Contrary to the embodiment shown in Figure 3 now the data consuming functional unit CFU serves as the initiator INIT.

The data consuming functional unit CFU as an initiator INIT issues a request to the target, here the data producing functional unit PFU by indicating a communication thread with a signal TID. Validity of the indicated communication thread and the related data is indicated with the signal VALID. As in the embodiment shown in Figure 3 the target functional unit has the options to

1. Indicating the initiator unit INIT that it has to maintain the request.
2. Indicating that the request of the initiator INIT is accepted.
3. Indicating that the request of the initiator INIT is rejected.

During the time necessary for determining whether the request can be handled the target functional unit responds according to option 1. If the target functional unit determines that it can handle the request on a short term it provides an indication according to option 2. However if it is determined that this impossible, for example because the required data is not expected to become available soon, the target functional unit responds according to option 3. Again this allows the initiator to withdraw the request for said communication thread and initiate a transaction for another communication thread which otherwise would have to wait until the earlier request could be handled. In this way the overall communication in the processing system is improved.

The indications according to the above-mention three options can be provided as follows using the signals READY and ACCEPTP.

Option	READY	ACCEPTP ACCEPTC	Meaning
1	0	0 (*)	The initiator functional unit has to continue indicating the communication thread identifier TID.
2	1	0	The initiator functional unit is requested to indicate an other communication thread identifier.
3	1	1	The indicated communication thread identifier is accepted.

As long as the target functional unit is not ready, indicated by READY = 0, the value of the signal ACCEPTP in the embodiment shown Figure 3 or ACCEPTC in the embodiment of Figure 4 is irrelevant. However, preferably the signal ACCEPTP or ACCEPTC is kept at a value indicating rejection as long as the target functional unit is not ready to decide between rejection and acceptance. This makes it possible to connect the target functional unit to a conventional initiator functional unit, having no input for receiving the signal READY. In this way the conventional initiator functional unit can normally operate, i.e. wait until the request associated with the indicated communication thread TID is accepted, regardless whether the waiting is caused by the target functional unit evaluating the request, or because the target functional unit cannot serve the request at that moment. For clarity the logical signal values 1 and 0 were used to indicate a negative and a positive outcome respectively, e.g. ready and not ready, or acceptance or non acceptance. It is clear however that these values could also be used the other way around. Many other alternatives are possible to provide an indication according to one of these options as described in the table. For example by providing the signals READY and ACCEPTC, ACCEPTP serially along one physical connection.

It is remarked that the scope of protection of the invention is not restricted to the embodiments described herein. It is noted that information, e.g. data for a communication thread, information about a communication thread, can be exchanged between the processing units in several ways, e.g. serial, parallel or in a combination of ways.

Neither is the scope of protection of the invention restricted by the reference numerals in the claims. The word 'comprising' does not exclude other parts than those

mentioned in a claim. The word 'a(n)' preceding an element does not exclude a plurality of those elements. Means forming part of the invention may both be implemented in the form of dedicated hardware or in the form of a programmed general purpose processor. The invention resides in each new feature or combination of features.

CLAIMS:

1. Method for communicating data between an initiator unit (INIT) which initiates the communication and a target unit (TRGT), wherein the initiator unit (INIT) indicates a request for a communication thread having a communication thread identifier (TID) to initiate a communication,
5 the target unit (TRGT) provides information (READY, ACCEPTC) indicating whether one of the following situations exist,
 - the initiator unit (INIT) has to maintain the request,
 - the request of the initiator (INIT) is accepted,
 - the request of the initiator (INIT) is rejected.
10
2. Method according to claim 1, wherein the initiator unit (INIT) indicates a reference (TID) to a communication thread for which it indicates the request, and wherein the initiator attempts to initiate a communication for another communication thread if the request is rejected.
15
3. Method according to claim 1, wherein the target unit (TRGT) provides a first and a second signal, the first signal (READY) indicating with a first state (0) that the request should be maintained, and indicating with a second state (1) that a second signal is valid, the second signal (ACCEPTC; ACCEPTP) having a first state (1) indicative that the target unit
20 accepts the selected communication thread, the second state (0) indicating that the initiator (INIT) should abandon the request.
4. Method according to claim 1, wherein a data producing unit (PFU) functions as the initiator (INIT).
25
5. Method according to claim 1, wherein a data consuming unit (CFU) functions as the initiator (INIT).

6. Processing system comprising an first functional unit (IFU5) and a second functional unit (DFU), the first functional unit being capable of initiating a communication with the second functional unit using a facility for exchanging signals between the functional units, the signals comprising a signal (TID) by the first functional unit (IFU5) that it desires
5 to initiate a communication,
signals (ACCEPTC, READY) by the second functional unit (DFU) indicative for one of the following situations
- the initiator unit (IFU5) has to maintain the request,
 - the request of the initiator is accepted,
 - 10 - the request of the initiator is rejected.

7. Processing system according to claim 6, wherein the first functional unit (IFU5) indicates a reference to a communication thread (TID) for which it indicates the request, and wherein the first functional unit (IFU5) attempts to initiate a communication for
15 another communication thread if the request is rejected.

8. Processing system according to claim 6 or 7, wherein the first (IFU5) and the second functional unit (DFU) form part of a network.

1/2




FIG.1




FIG.2

2/2

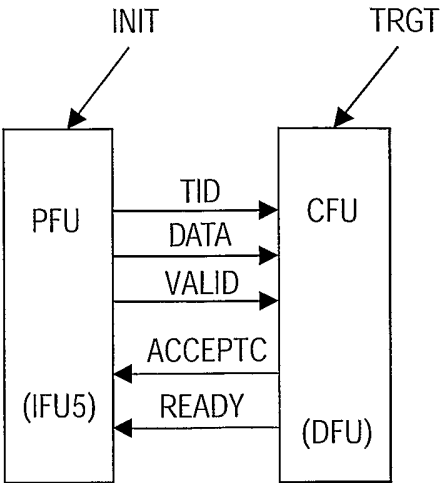


FIG.3




FIG.4

INTERNATIONAL SEARCH REPORT

International Application No

IB2004/050580

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 G06F13/36 G06F13/42

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC 7 G06F H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 922 063 A (OLNOWICH ET AL) 13 July 1999 (1999-07-13) column 7, line 26 - column 9, line 20 column 14, line 16 - line 52 claim 1; figures 1-3 -----	1-8
X	US 6 101 194 A (ANNAPAREDDY ET AL) 8 August 2000 (2000-08-08) column 3, line 43 - column 4, line 31 column 5, line 44 - column 7, line 52 figures 1,5 -----	1,2,4-8
A	US 5 175 733 A (NUGENT) 29 December 1992 (1992-12-29) column 3, line 13 - column 4, line 46 column 9, line 4 - column 11, line 23 -----	1-8



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

* & * document member of the same patent family

Date of the actual completion of the international search

6 October 2004

Date of mailing of the international search report

13/10/2004

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

McDonagh, F

INTERNATIONAL SEARCH REPORT

Information on patent family members

ional Application No

IB2004/050580

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5922063	A	13-07-1999	JP 6214965 A	05-08-1994
US 6101194	A	08-08-2000	NONE	
US 5175733	A	29-12-1992	GB 2251356 A , B	01-07-1992
			HK 56495 A	21-04-1995
			JP 5153163 A	18-06-1993
			SG 10695 G	01-09-1995