



Published:

- with international search report
- 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.

METHODS AND DEVICES OF DATA TRANSMISSION USING CLOCK DOMAINS

The present invention relates to a transmitter.

The present invention further relates to a method of transmitting.

The present invention further relates to a receiver.

The present invention further relates to a method of receiving.

5 The present invention further relates to a data handling unit.

The present invention further relates to a network.

The present invention further relates to a mobile device.

The interconnect centric approach offers a powerful way to rapidly develop
10 new systems. In such an approach the system is developed as a plurality of nodes. The nodes,
also denoted as data handling units, comprise functional units e.g. storage units, dedicated
processors, general processors and data routing units such as routers and switches. The
functional units are arranged in a network formed by the data routing units. It is noted that
such a network may be a network on chip, a network coupling various integrating circuits, or
15 a network coupling various computers. It is a fact that the communication protocol of the
nodes tends to be standardized, and that a network like architecture may easily be expanded
with new nodes, facilitating design. For cost and power reasons, links between the nodes are
serial, use differential low-swing signaling, and run at high frequencies (1 GHz and above).
At these speeds it is not possible to run a multiple chip system at a single clock. For this
20 reason each chip has a local clock. Despite the fact that the clocks can mutually have the
same nominal frequency, variations within known tolerance will in practice occur. These
variations are caused by imperfections in the crystal oscillators and local temperature
differences. In other systems various nodes may have intentionally different clockrates. Data
transfer is still synchronous to a clock driven by the data producing node (transmitter). The
25 clock is either sent on second serial pair of wires (source synchronous data transmission), or
the clock is embedded in the data wires using for example an 8b10b encoding as in PCI
Express. The data is sampled in the destination node (receiver) using the clock of the
transmitter sent together with the data. Using one of the known clock-domain crossing
techniques, data is then transferred to the clock domain of the receiver. When implementing

systems providing guaranteed performance, one must control precisely the usage of the resources in the system. One cost-effective way of achieving this is to define frames of slots in which slots are reserved for guaranteed-throughput communication. This data for which a guaranteed throughput is required will also be denoted as primary data in the sequel. Data
5 used for control of various functions will be denoted as control data. Such a system requires that frames in all devices and switches are synchronized.

It is a purpose of the invention to synchronize the transmission of the frames between the nodes, despite the fact that the nodes have independent clocks.

10

According to the invention this purpose is achieved with a receiver as claimed in claim 1.

According to the invention this purpose is achieved with a method for receiving as claimed in claim 6.

15

According to the invention this purpose is achieved with a transmitter according to claim 7.

According to the invention this purpose is achieved with a method of transmitting according to claim 8.

20

In a practical embodiment a receiver according to claim 1 and a transmitter according to claim 7 are combined in a data handling unit as claimed in claim 12.

A plurality of data handling units may be combined in a network as claimed in claim 13.

25

According to the invention on the one hand the receiver communicates to the transmitter the number of slots with primary data that it has received, and converted to its own clock domain. The transmitter uses this information to determine whether the receiver is lagging or leading. If it determines that the receiver is lagging it interrupts its transmission of primary data and signals this to the receiver by transmission of a pause symbol, which is a particular form of control data. In this way the transmitter and the receiver are mutually synchronized.

30

In practice, a node will comprise both a receiver and a transmitter as is schematically shown in Figure 8. They share a counter: 'ns'. The transmitter sends data to all its neighbours to signal it slows down. The receivers are those part of a node that run being driven by the clocks in the neighbours. Using a standard clock domain crossing (e.g., 2 flip-flops, or a fifo), the received data is transferred in the node's clock domain which comprises

the transmitter. As a result, the transmitter doesn't signal PAUSEs to the receivers, but sends PAUSE messages to its neighbours. The goal is to synchronize the node and the neighbours. In a network of data handling units the mutual synchronization of each pair of nodes results in a global synchronization of network allowing for a global scheduling of primary data traffic.

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

Figure 1 schematically shows a data processing system in which the present invention may be applied,

Figure 2 schematically illustrates a scheme of data transfer between two nodes.

Figure 3 shows an example of a data packet in more detail.

Figure 3A shows an example of a datapacket in another embodiment of the invention,

Figure 3B shows an example of an escape symbol in said another embodiment,

Figure 4A schematically shows a method for receiving data, and

Figure 4B schematically shows a method for transmitting data,

Figure 5 schematically shows the interaction between a transmitter and a receiver as a function of time,

Figure 6 schematically shows an embodiment of a transmitter according to the invention,

Figure 7 schematically shows an embodiment of a receiver according to the invention.

Figure 8 schematically shows a pair of nodes both comprising a combination of a transmitter and a receiver according to the invention.

Figure 1 schematically shows a data processing system in which the present invention may be applied. The data processing system shown is a camera having various functional units, such as a modem 1, a communication accelerator 2, a first and a second general purpose processing engine 3, 6, a media accelerator 4, a camera 8, a display 9 and a mass storage units 5 and 10 as well as an auxiliary device 7. The functional units are coupled

in a network by switches S1, S2, S3 and S4. The various functional units and switches each operate at their own clock. Although the clocks may approximately have the same speed, an exact synchronization of the clocks cannot be provided. When transmitting guaranteed data, such as isochronous data it is essential that the transmission is globally synchronized in the network. The present invention provides a communication scheme that guarantees that this condition is fulfilled.

Figure 2 schematically illustrates a scheme of data transfer between two nodes. In the scheme shown, the available time for data transmission is subdivided in time-slots (SL), which are indicated as rectangles. Each time-slot is available for transfer of a packet of data. Part of the time slots is reserved for data requiring a guaranteed throughput, here denoted as primary data such as isochronous data. In this example these time slots are indicated by areas ISL. The other time slots are not reserved in advance, but can be granted at run-time for use by other data, also denoted as secondary data. Arbitration mechanisms known as such, e.g. round robin, priority scheduling may be used to select a data packet if two or more data sources want to transfer data along the same link. The remaining data can be transferred as bulk data BD, or as separate chunks of data. As can be seen in Figure 2, the slot reservations repeat after a fixed number of slots. This fixed number of slots is denoted here as a frame. In this case a frame comprises 128 slots, but any number could be applied. Here a first frame FR1 starts at $t=0$, a second frame starts at $t=128$, where the time unit is the duration of a slot.

In a practical embodiment a packet comprises for example 131 bytes and slot has a duration of approximately $1\mu\text{s}$. This corresponds to a data transmission rate of 1Gbps. In this example, where a frame comprises 128 slots, the frame repetition rate is 8kHz.

Figure 3 shows an example of a data packet in more detail. The data packet shown comprises a header H, a payload PL and a trailer T. In the embodiment shown the header H, payload PL and trailer T respectively comprise 2, 128 and 1 byte. As shown in more detail in the lower part of Figure 3, header H comprises the following information about the remainder of the packet:

A type indicator T1, T2. These two bits encode the following types:

Empty: An empty type of packet indicates that the link is active, but that the transmitted packet contains no data.

Isochronous, indicates a prescheduled package of a stream requiring a guaranteed throughput. This type of data is indicated as primary data.

Best_effort, or secondary data, the transmission of which is scheduled at run-time.

5 Escape: The escape type allows for a different format of the remainder of the header, which can be useful for various control functions, e.g. for activate a link via which the data is communicated, to deactivate the link, or to indicate an error.

 The flow control bits F1,, F5 serve to indicate to the receiver of the packet a number of credits. This number indicates the number of packets that can be accepted
10 until buffer overflow occurs.

 In case of an irrecoverable error a packet may be returned with the error flag E set. In response the device receiving the returned packet will execute a retransmission.

 The number of bytes used in a packet is indicated by the bits L1,, L7.

 The last packet in a sequence of BE packets is indicated by the EoP flag.

15 The trailer of the packet is preferably used for an error correction/detection code.

 An alternative data format is shown in Figure 3A. In this format a greater number of different types of packets is provided for so that for example flow-control and
20 error information is sent as separate messages, instead of sending them together with a payload.

 The greater number of types is encoded with type bits T1, T2, T3. Various types are for example

 ISOC isochronous data

25 BE control data, e.g. data for controlling a setting of the device, e.g. volume control, contrast control, which usually only comprises a single packet..

 BE bulk. Best effort data that comprises a plurality of packets.

* ESC symbols. The latter may be of normal or urgent type.

 The header may in addition comprise further data for indicating a source and a destination in
30 the data. For isochronous this information may be encoded in the slot table.

 Figure 3B shows in more detail the format of an escape symbol. Whereas the type bits T1, T2, T3 indicate that an escape symbol is present, the bits E1,, E5 identify the nature of the escape symbol: Escape symbols may be of type normal or urgent:

An example of a normal escape symbol is ESC_FC, which is used for flow control. In this case the payload P1,,P8, indicates the receiver the number of credits, i.e. the number of data units, which the transmitter of the escape symbol is ready to accept.

Escape symbols of type panic should be handled with urgency. These are for example

ESC_ERROR: to indicate that a received package has an irrecoverable error, and should be retransmitted.

In case of the ESC_SYNC, the payload comprises a slot number .

ESC_PAUSE is used to indicate that the transmitter temporarily stops transmitting primary data to allow the receiver to remain in pace with the transmitter.

As in the case of other data the ESC symbol includes a trailer for error checking and correction.

When handling packets, preferably the following priorities should be given:

- 1 (highest priority). Panic ESC symbols, such as ERROR, SYNC and PAUSE.
2. Isochronous data traffic, and
3. Normal ESC symbols: FLOW_CTRL
4. BE control
5. (lowest priority) BE bulk

Figure 4A schematically shows a method for receiving data, and Figure 4B schematically shows a method for transmitting data. The steps carried out by the transmitter are indicated by T1, ..., T8. The steps carried out by the receiver are indicated by R1, ..., R7

In step T1 of the transmitter after startup, a slot counter slot and a difference indicator ns are both initialized at 0.

In step T2 it is determined for which connected nodes data is available. Subsequently the links connecting those nodes are activated.

In step T3 those links connecting to nodes for which no data is available, are set into a sleep-mode.

Alternatively, in environments having no power constraints, e.g. apparatus supplied by the mains, all links may be kept active continuously.

In step T4 the available data is transmitted in the form of a packet to its destination. In this embodiment the packet comprises primary data as a payload, and control data in the form of a header and/or a trailer. The control data in the header indicates whether

it is followed by primary data in the form of a payload. In addition the control data may indicate the length of the payload.

In step T5 the slot counter is incremented, which is representative for the transmitted number of units of primary data. The step of incrementing the counter may be executed before the step of transmitting the primary data.

In step T6, it is verified whether the value of the difference indicator ns is greater than 0. The value of ns is the number of transmitted primary data units slot minus the number of primary data units which are received rcv_slot .

In step T7 a PAUSE symbol is transmitted. Subsequently the difference indicator ns is decreased by 1. The PAUSE symbol forms control data indicating to the receiving node (receiver) that there is no primary data. In this embodiment the PAUSE symbol replaces one packet of primary data. In other embodiments a different granularity may be selected, e.g. a PAUSE symbol replacing a single byte of primary data, or a PAUSE symbol replacing number of packets.

The operation of the receiver is now illustrated with reference to Figure 4A. In step R1 the receiver enters active power mode. Active power mode may be initiated by a special control word from the transmitter, or by power on of the data-processing system. In step R2 it receives a control word indicative for a slot number of the next unit of primary data that will be transmitted by the transmitter. It initializes a slot counter with this data. In step R3 the receiver receives a next unit of control data. In step R4 the receiver determines whether this control data unit indicates whether it is followed by primary data or whether it is a PAUSE symbol. In the latter case the receiver waits for the next data unit in step R3. In the case that the unit of control data indicates that it is followed by a payload of primary data, the receiver confirms receipt in step R5, increments its slot counter in step R6, and receives the primary data in the form of a payload and eventually further control data in the form of a trailer. The steps of confirming, incrementing and receiving may be executed in any order or executed in parallel. After steps R5, R6 and R7 the receiver continues with step R3.

The slot counter rcv_slot will also be incremented upon receipt of a packet of type EMPTY, or a packet containing secondary (best-effort) data. However, a bonus packet of secondary data, for which the rcv_slot is not incremented, may be transmitted immediately following the PAUSE symbol.

Alternatively a separate escape symbol may be used, e.g. indicated as ESC_STOP_SLOT for example, which, when preceeding an EMPTY or BE-packet suppresses an incrementation of the rcv_slot counter.

5 Upon receipt of the confirmation in step T8 the transmitter calculates the difference d between the number of transmitted slots of primary data slot and the number of received slots rcv_slot of primary data.

$$d = slot - rcv_slot$$

10 The counters slot and rcv_slot are wrap around counters, which can have a relatively low maximum value, in a practical embodiment for example 128. Due to the tight frame synchronization, the difference between the counters is limited to a small value, e.g. 1 or 2, so that aliasing is avoided.

15 In an alternative embodiment, does not store a counter for the number of received slots rcv_slot, but stores and updates this difference directly. In the alternative embodiment, when the link is activated, a word is received with the slot position of the neighbour (rcv_slot).

As a result, 'd' is initialized with (slot - rcv_slot). Following this, 'd' is incremented when the node's slot is incremented, and decremented when a data for a slot slot has been received from the neighbour.

20 In case of only a single receiver, the difference indicator ns is equal to this difference.

However, in case of a plurality of receivers, the transmitter has to adapt to the slowest one. In that case the difference indicator is calculated as:

$$ns = \max(ns, d)$$

25 Figure 5 schematically shows the interaction between a transmitter and a receiver as a function of time. In this example it is assumed that the receiver has a slower clock than the transmitter. For the purpose of illustrating the present invention this difference in clock speed is strongly exaggerated in the Figure. In practice the difference in clock speed is for example in the order of 0.1 %.

30 At t_0 the transmitter sends a first packet comprising a header, a payload and a trailer. The payload comprises of primary data and the header and trailer comprise control data. At t_1 the receiver has recognized the header and sends a message announce rcv_slot. After this message is received by the transmitter of the packet of data the counter rcv_slot is incremented at t_{1a} . Immediately the difference indicator ns is recalculated, and obtains the

value minus one. At t1b the transmitter has completed transmission of the packet and increases the slotcounter slot. In addition the difference-indicator ns is recalculated and obtains the value 0 again. Now the transmitter verifies the value of ns in step T6 and decides that a new packet can be transmitted. The receiving module recognizes the header at t2 and
5 sends a message announce rcv_slot. Upon receipt of this message at t2a, the transmitting module increments the counter rcv_slot and recalculates the value of ns, which obtains the value -1. At t2b the transmitter has completed transmission of the packet, increments counter slot, and recalculates the value of ns, which becomes 0 again. As this value is again 0, when the transmitter executes step T6, it decides to send a next packet. At t3 the receiver has
10 received and processed the header of this packet and transmits a message announce_rcv slot. At t3a the transmitter has received this message, and increments the counter rcv_slot. However, due to the higher clock speed of the transmitter, the latter has already finished transmitting its packet before this point in time, at t3b, and has increased its slot counter before t3a. Consequently, at the moment that the transmitter executes step T6 it finds that the
15 difference indicator is greater than 0. Consequently, instead of transmitting a packet, it now transmits a Pause symbol PS. After transmission of the Pause symbol it refrains from sending a payload and a trailer, so that the receiver has time to process the previous packet. At t3c, after the time slot is completed following the Pause symbol, the transmitter decrements the difference indicator ns with 1 to 0, and subsequently it transmits a next packet. During the
20 period after the Pause symbol, wherein the transmitter refrains from sending primary data, it may continue transmitting Pause symbols. Alternatively it may enter a low-power mode. In again another embodiment it may transmit secondary data, e.g. best effort data. In the embodiment shown here, the Pause symbol indicates that the receiver refrains from transmitting 1 packet of primary data. In another embodiment the transmitter may interrupt
25 transmission of primary data for a period longer than one packet. In that case the duration may have a predetermined duration e.g. the duration of a fixed number of packets. Alternatively the Pause symbol may include an indication for the length of the period during which transmission of primary data is interrupted.

When the receiver has received the Pause symbol it is 'aware' that the
30 transmitter refrains from transmitting the primary data. Consequently it refrains from communicating an 'announce-receive slot'

Due to the interrupted primary data stream, the receiver can now immediately start to process the next data packet transmitted by the transmitter at t3c, and send an 'announce-receive slot at t4. From that point in time the procedure repeats. At t6b the delay

of the receiver is again incremented to such an amount that the difference indicator is greater than 0, and the transmitter again transmits a Pause symbol.

Figure 6 schematically shows an embodiment of a transmitter TRM according to the invention. A controller CTRL controls a multiplexer M1 that selects one of a plurality of data sources to provide data for the output. In this case a data source HEAD is comprised, which provides a header for a data packet. In practice various headers may be used depending on the type of data, e.g. best effort data, or isochronous data, it may comprise information about the length of the payload. A second data source TRAIL provides a trailer, which may comprise for example an error correction code. A third data source provides a PAUSE symbol to indicate that the transmitter interrupts transmission of primary data. A fourth data source PRIMARY provides the primary data. A fifth data source SECONDARY provides secondary data. Various other data sources may be present for selection, e.g. to provide various control symbols, e.g. for activate a link via which the data is communicated, to deactivate the link, or to indicate an error.

The transmitter has an output TO for providing the selected data to a receiver. The transmitter further has an input TI for receiving the announced number of received slots.

In the process of transmitting primary data the controller observes the value of difference indicator ns. The difference indicator is decremented with signal DEC when a PAUSE symbol is transmitted and when the number of transmitted slots slot or the number of received slots, announced by the receiver, is updated. To that end the difference indicator is coupled to a subtractor S2 via a maximum function module MAX. The latter module has apart from a first input coupled to the subtractor S2 a second input coupled to difference indicator register. The subtractor calculates a difference between the actual number of transmitted primary data units (slot), and the number of primary data units rcv_slot that the receiver has announced it has received. This embodiment has the advantage that the transmitter can adapt to the receiver with the slowest clock.

Figure 7 schematically shows an embodiment of a receiver according to the invention. The receiver RCV has an input RI for receiving a stream of data from the transmitter TRM at a data rate corresponding to the clock of the transmitter. It further has a first comparator PRIM for determining whether the data it is receiving is the header of a packet of primary data. In that case it transmits a message announce_rcvslot to the transmitter. The receiver further has a second comparator PAUSE, for recognizing whether it

has received a PAUSE symbol. This comparator control a gate GT which couples a buffer BUF to the input I. If a PAUSE symbol is recognized, the gate is closed so that filling of the buffer is interrupted during a length of time corresponding to a data packet. Otherwise the gate GT is opened, so that the buffer can be filled at the speed of the clock CLT of the transmitter. In the embodiment shown, the clock CLT is provided via a separate connection. In another embodiment the clock is embedded in the data stream. The buffer is read out by a data processing unit DPU at a clock rate CLR of the receiver. Instead of interrupting the filling of the buffer when a PAUSE symbol is recognized, all data may be loaded in the buffer. In that case a read pointer indicative for the current position that is read from the buffer may be advanced with a number of positions corresponding to the size of a packet. Alternatively the PAUSE symbol may include information indicating the number of positions in the buffer that may be skipped.



Figure 8 schematically shows a pair of nodes. The first node N1 comprises a combination of a transmitter TR1 and a receiver RC1. The second node N2 comprises a combination of a transmitter TR2 and a receiver RC2.

It should be noted that the above-mentioned embodiments illustrate rather than limit the invention, and that those skilled in the art will be able to design many alternative embodiments without departing from the scope of the appended claims. In the claims, any reference signs placed between parentheses shall not be construed as limiting the claim. The word "comprising" does not exclude the presence of elements or steps other than those listed in a claim. The word "a" or "an" preceding an element does not exclude the presence of a plurality of such elements. In the device claim in numerating several means, several of these means can be embodied by one and the same item of hardware. The mere fact that certain measures are resided in mutually different dependent claims does not indicate that a combination of these measures cannot be used to advantage.

The following document is included as an integral part of this application:

2005-06-14-ups-unipro-paris.ppt

Furthermore, any reference signs in the claims shall not be constitute as limiting the scope of the claims.



UPS-x

Unified Protocol System

A packet-based protocol for chip-to-chip interconnects

Andrei Rădulescu, Peter van den Hamer, Ewa Hekstra-Nowacka
Philips Research
June 14th, 2005

Philips confidential

MIPI Alliance Confidential

Legal Disclaimer

The material contained herein is not a license, either expressly or impliedly, to any IPR owned or controlled by any of the authors or developers of this material or MIPI. The material contained herein is provided on an "AS IS" basis and to the maximum extent permitted by applicable law, this material is provided AS IS AND WITH ALL FAULTS, and the authors and developers of this material and MIPI hereby disclaim all other warranties and conditions, either express, implied or statutory, including, but not limited to, any (if any) implied warranties, duties or conditions of merchantability, of fitness for a particular purpose, of accuracy or completeness of responses, of results, of workmanlike effort, of lack of viruses, and of lack of negligence. ALSO, THERE IS NO WARRANTY OR CONDITION OF TITLE, QUIET ENJOYMENT, QUIET POSSESSION, CORRESPONDENCE TO DESCRIPTION OR NON-INFRINGEMENT WITH REGARD TO THIS MATERIAL.

IN NO EVENT WILL ANY AUTHOR OR DEVELOPER OF THIS MATERIAL OR MIPI BE LIABLE TO ANY OTHER PARTY FOR THE COST OF PROCURING SUBSTITUTE GOODS OR SERVICES, LOST PROFITS, LOSS OF USE, LOSS OF DATA, OR ANY INCIDENTAL, CONSEQUENTIAL, DIRECT, INDIRECT, OR SPECIAL DAMAGES WHETHER UNDER CONTRACT, TORT, WARRANTY, OR OTHERWISE, ARISING IN ANY WAY OUT OF THIS OR ANY OTHER AGREEMENT RELATING TO THIS MATERIAL, WHETHER OR NOT SUCH PARTY HAD ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES.



Contents

- **“Coppola” design goals**
- **UPS in a nutshell**
- **Packet concepts**
 - ISOC, BE, ESC
- **Zooming in**
 - ISOC: frame synchronization (key innovation)
 - error handling & flow control
 - BE grouping
- **“Coppola” graphs**
- **Summary**

3

MIPI Alliance Confidential



“Coppola” design goals

- **Low Cost**
 - M-phy
 - buffers for real-time guarantees
- **Low Power**
 - phy needs dynamic power management
 - grouping of data is the only low-power alternative for M-phy
 - trade off against latency / buffering
- **Low Latency**
 - most traffic is not latency critical
 - media & communication streaming (real time)
 - file transfers, browsing, emails (bulk)
 - still there is latency-critical traffic (e.g., interrupts)
 - given a high priority
- **Low Latency is good, but don't sacrifice Cost & Power**
 - low latency → architecture flexibility




Francis Ford Coppola

- **COst Price**
- **Power**
- **LAtency**

4

MIPI Alliance Confidential




UPS in a nutshell

Too busy!
Simplify?

- Satisfies all UniPro required and recommended features
- Main features
 - three classes of services
 - guaranteed throughput, low latency and best effort
 - packet based
 - network based
 - interfaces to the MIPI serial phy
 - supports the current D-phy
 - also supports the future M-phy
 - aggressive power management
 - switching off the link when no used
 - data grouping
 - supports links at different speeds, multiple of a nominal frequency
 - robust to link errors
- Uses simple, well-known technology
 - one key innovation to have the option of minimizing power

5

MIPI Alliance Confidential



Packet concepts

- ISOC packets
- Best-effort packets
 - low latency
 - bulk
- Escape symbols

6

MIPI Alliance Confidential




UPS packets

▪ Packet structure

- header
 - type (3b), length (7b)
 - BE only: destination ID (8b), source ID (8b)
- payload: maximum 128 Bytes
- trailer: CRC-16 (16b)

▪ Packet types

- ISOC
- BE control
- BE bulk
- ESC symbols
 - live only at the link level




▪ Packet priorities

0. Panic ESC symbols: ERROR, SYNC, PAUSE
1. ISOC
2. Normal ESC symbols: FLOW_CTRL
3. BE control
4. BE bulk

7

MIPI Alliance Confidential




ISOC packets


- Always aligned to the slot boundaries
- Can only be transferred in reserved slots

▪ Optional

- Destination ID (8b)
- Source ID (8b)
- can be encoded in the slot table



- A link can be put into low-power mode anytime there is no data to send





8

MIPI Alliance Confidential




BE packets (low latency & bulk)

- Lower priority than ISOC
 - preempted by ISOC
- Use all the space left over from ISOC
- No constraints relative to slots
 - multiple BE packets in a slot
 - a BE packet can span over multiple slots
 - BE packets can start/end anywhere in a slot





9

MIPI Alliance Confidential



ESC symbols

- Very short
- Two classes
 - normal:
 - ESC_FC
 - panic:
 - ESC_ERROR
 - ESC_SYNC, ESC_PAUSE
- Can be sent at any packet boundary



10

MIPI Alliance Confidential




Zooming in

- Frame synchronization
- Error handling & flow control
- BE grouping

11

MIPI Alliance Confidential



Frame synchronization

- Why synchronizing frames across multiple links?
 - guarantees no buffer overflow at low buffer sizes in switches
 - offers the option of low power
- Our solution: frame synchronization via coarse-grained handshaking
 - requirement: same nominal frequency (or a multiple of it)
 - nodes communicate their **slot number** at least once during an interval
 - e.g., 1000ppm, max drift 1% of frame → slot info exchange every 10 frames
 - a node **slows down** its slot speed if it is faster than its neighbors


12

MIPI Alliance Confidential




Synchronizing church clocks

- Each church clock runs at its own speed.
- Speed difference (e.g. <1 part in 2000) is unknown and drifts.
- For high-impact global synchronous prayer events we want to synchronize (lock) the effective clock speeds
- = need to guarantee max time difference.
- Can communicate events (and times) by ringing bells.

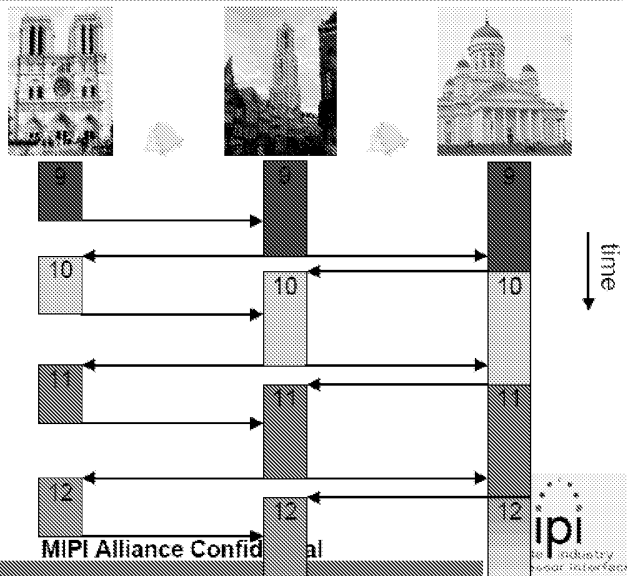


13

Simplest synchronization solution

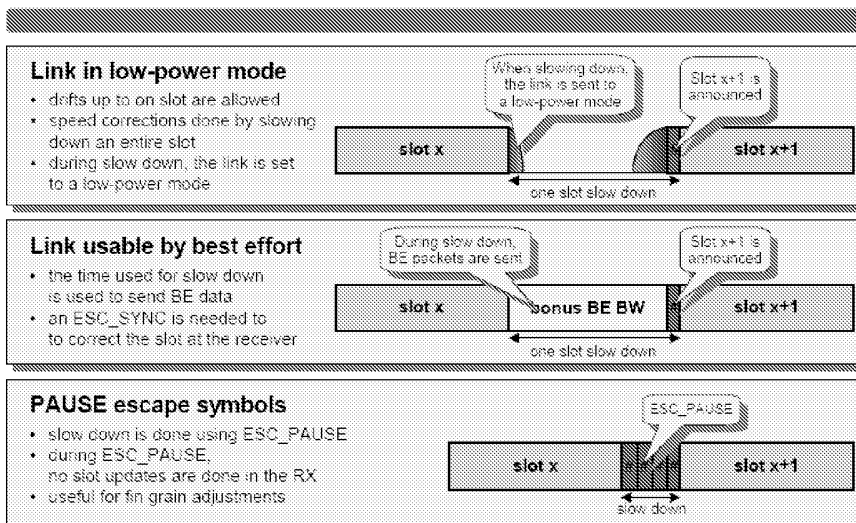


- Result: slowest locked clock determines rate.
- Trick also works for finite signal propagation speed.
- Typical time difference $\ll 1$ time unit.



14

Frame sync: slowing down a node



17

MIPI Alliance Confidential

mi
mobile industry
processor interface

Error handling & BE flow control


- Solved at the link level
- Error handling
 - uses error detection and retransmission
 - needed to recover from data and clock errors
 - required for BE
 - nice to have for ISOC
 - BE traffic
 - must be 100% robust to errors
 - worst-case timing is not important
 - ISOC traffic
 - must move forward (worst-case timing is essential)
 - repeated retransmissions conflict with the basic nature of ISOC
 - tradeoff between
 - reducing the ISOC error rate (cannot be 0), and
 - latency / buffer sizes
- Separate mechanisms for ISOC and BE
 - BE: sliding window, flow control included
 - ISOC: replay à la Rubee, no flow control

18

MIPI Alliance Confidential

mi
mobile industry
processor interface


Error handling: best effort



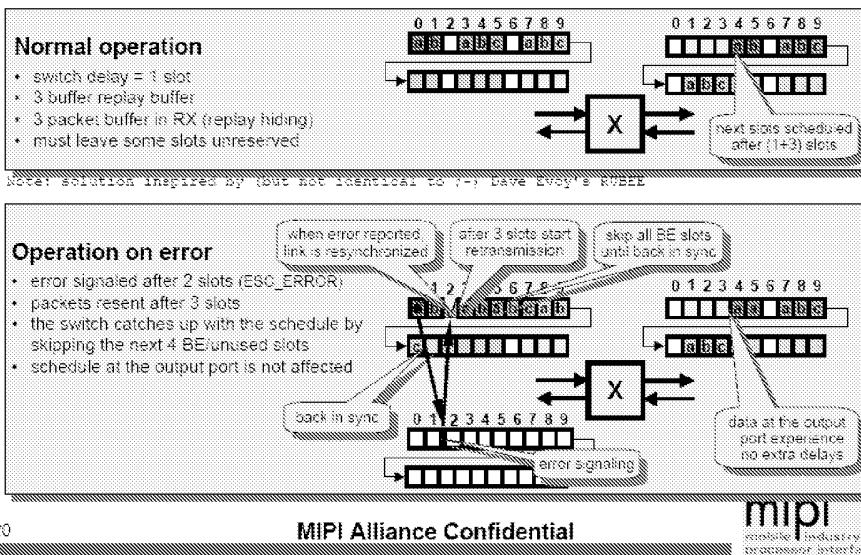
- Uses a sliding window protocol variation
 - 100% robust
 - combined with flow control
- **Producer**
 - cannot send data beyond *space*
 - *start(p)* to *last* serves as replay buffer
 - *end(p)* is the last valid data in the producer's buffer
- **Consumer**
 - data is removed → new *start(c)* is reported (ESC_FC) to update *start(p)*
 - on *error* → the last correct received slot is reported (ESC_ERROR)
 - *start(p)* is updated
 - all packets starting from the error are retransmitted as normal BE packets

19

MIPI Alliance Confidential



Error handling: ISOC



20


MIPI Alliance Confidential



Traffic grouping

Data threshold for BE traffic grouping


- link activated when threshold reached
- data in buffer sent in a continuous burst
- when no data can be sent, the link is sent to a low-power mode



- **Grouping data saves power (especially for the M-phy)**
 - ISOC
 - grouped slots
 - BE
 - data thresholds
 - send BE packets when link activated for ISOC
- **level of grouping is configurable via software**
- **tradeoff between low power and latency / buffering**

21

MIPI Alliance Confidential



“Coppola” graphs


- **Cost**
- **Power**
- **Latency**

22

MIPI Alliance Confidential



Cost




- M-phy has roughly half the cost of D-phy (for the same bandwidth)
- Grouping cost (power benefit 1-6x)
 - adaptor ports: ~30% extra cost
 - switch ports: same cost
 - overall interconnect: 15% extra cost

23

MIPI Alliance Confidential




Power




- Grouping saves significant power (especially for M-Phy)
- UPS allows any level of grouping
 - configurable via software
- UPS does not waste any unnecessary power
 - slots are filled up entirely (with ISOC and BE packets)
 - low-power communication is used for small urgent messages

24

MIPI Alliance Confidential




Latency




- More grouping → longer latencies
- UPS allows trade offs
 - low power vs. low latency / low buffering
 - for ISOC and BE

25

MIPI Alliance Confidential




UniPro high-end use case



- **Power**
 - spread: 53 mW (D), 82 mW (M), 58 mW (D+M)
 - grouped: 48 mW (D), 39 mW (M), 37 mW (D+M)
- **Cost**
 - spread:
 - grouped:

26

MIPI Alliance Confidential




Summary

- **UPS protocol features**
- **UPS benefits**

27

MIPI Alliance Confidential




Summary of protocol features

- **Packet-based**
- **Supports 3 traffic classes: isoc, messages & bulk**
- **Error handling for all traffic classes**
- **2 low-power states (deep/light sleep)**
- **Flow control for best-effort traffic**
- **Optimized for low buffer sizes**
- **Optimized for low latency**
- **No special PHY features required (no 8b10b, etc)**
- **Unused reserved bandwidth can be reused by BE**

28

MIPI Alliance Confidential





END

Changes from the Dallas version

- **Packet structure**
 - data is transported in ISOC and BE packets
 - the μ Packet concept has become obsolete
- **Packet alignment to the slot boundary**
 - BE packets can start anytime when there are no ISOC packets
 - 100% efficiency in link utilization
 - no partially filled slots
- **Escape symbols**
 - can occur anytime between packets
 - some of the previously piggybacked info moved to ESC symbols
 - e.g., flow control, error notification
- **Separate data recovery for ISOC and BE**
 - still at the link level
 - 100% robust for BE
 - in the case of multiple/long errors, ISOC packets are dropped
- **Frame synchronization**
 - during a node slowdown, BE packets are possible

CLAIMS:

1. Receiver (RCV) for receiving data from a transmitter, comprising
a first clock domain operating at a data rate synchronous with a clock of the
transmitter, and having an input (RI) for receiving data the data including primary data,
secondary data and control data,
5 a second clock domain operating at a clock rate independent from the
transmitter,
a clock-domain crossing unit for transferring data from the first to the second
clock domain,
a slot counter for counting a number of units of received data converted by the
10 clock-domain crossing unit into the second clock domain,,
a first identification unit (PRIM) for identifying control data indicative for the
presence of primary and secondary data,
a second identification unit (PAUSE) for identifying control data indicative
whether slot counter of the receiver will be updated or not,
15 an output (RO) for communicating an indicator indicative for the value of the
slot counter.
2. Receiver according to claim 1, wherein the clock-domain crossing unit
comprises a buffer, and the second identification unit (PAUSE) prevents data from entering
20 the buffer if it detects an absence of primary data
3. Receiver according to Claim 1, wherein the second identification unit
(PAUSE) prevents the slot information in the receiver to be updated.
- 25 4. Receiver according to claim 1, wherein the clock-domain crossing unit
comprises a buffer, all data being entered in the buffer, a data processing unit (DPU)
selecting primary data from the buffer.

5. Receiver according to Claim 1, wherein the receiver has a counter being initialized with a value transported by an escape symbol ESC_SYNC.

6. Receiver according to Claim 5, wherein the escape symbol ESC_SYNC for counter initialization is transferred at the link activation.

7. Receiver according to claim 1, wherein the receiver has a counter for counting a number of received primary and secondary data units and communicates the counted number to the transmitter.

10

8. Receiver according to claim 1, wherein the receiver communicates the occurrence of receiving a primary or secondary data unit (rcv_slot).

9. Method for receiving data from a transmitter comprising the following steps,

15 A. receiving a data word (R3),

B. identifying the data word (R4),

C. if the data word is a control word indicative for the absence of primary data continue with step A

20 D. otherwise receiving a packet of words (R7) at a clock determined by the neighbour's transmitter,

F. communicating an indicator for receipt of the packet to the transmitter (R5) when at least a part of the packet has been read,

G. continue with step A.

25 10. Transmitter (TRM) comprising
a unit (M1, HEAD, TRAIL, PAUSE, PRIMARY, SECONDARY) for
transmitting primary data and control data,

a counter (slot) for counting the transmitted amount of data of the primary type,

30 an input (TI) for obtaining an indication for a received amount of primary data,

a facility (S2) for calculating the difference between the transmitted amount of primary data and the received amount of primary data,

the unit for transmitting data being controlled by the outcome (ns) of this calculation,

11. Method of transmitting primary data from a transmitter to a receiver in a
5 network, comprising the steps of
transmitting units of primary data (T4),
counting the transmitted number of units of primary data (T5),
receiving an indication for the number of units of primary data received by the
receiving node (T8),
10 determining a difference between the transmitted number and the received
number (T8),
if the received number of primary data units is less than the transmitted number of primary
data units, transmitting control data indicative that primary data is absent (T7).
- 15 12. Method according to claim 9, wherein the transmitter continues transmitting
said control data until the difference between the transmitted number (slot) and the received
number (rcv_slot) is reduced to zero.
13. Method according to claim 9, wherein the transmitter temporarily stops
20 transmission after transmission of the control data indicative for the absence of primary data.
14. Method according to claim 9, wherein the transmitter transmits a unit of
secondary data after transmission of the control data indicative for the absence of primary
data.
25
15. Data handling unit comprising a combination of a receiver according to claim
1 and a transmitter according to claim 10.
16. Network comprising a plurality of data handling units as claimed in claim 15.
30
17. Mobile electronic device containing a network as claimed in Claim 16.




FIG.1




FIG. 2




FIG. 3




FIG. 3A




FIG. 3B

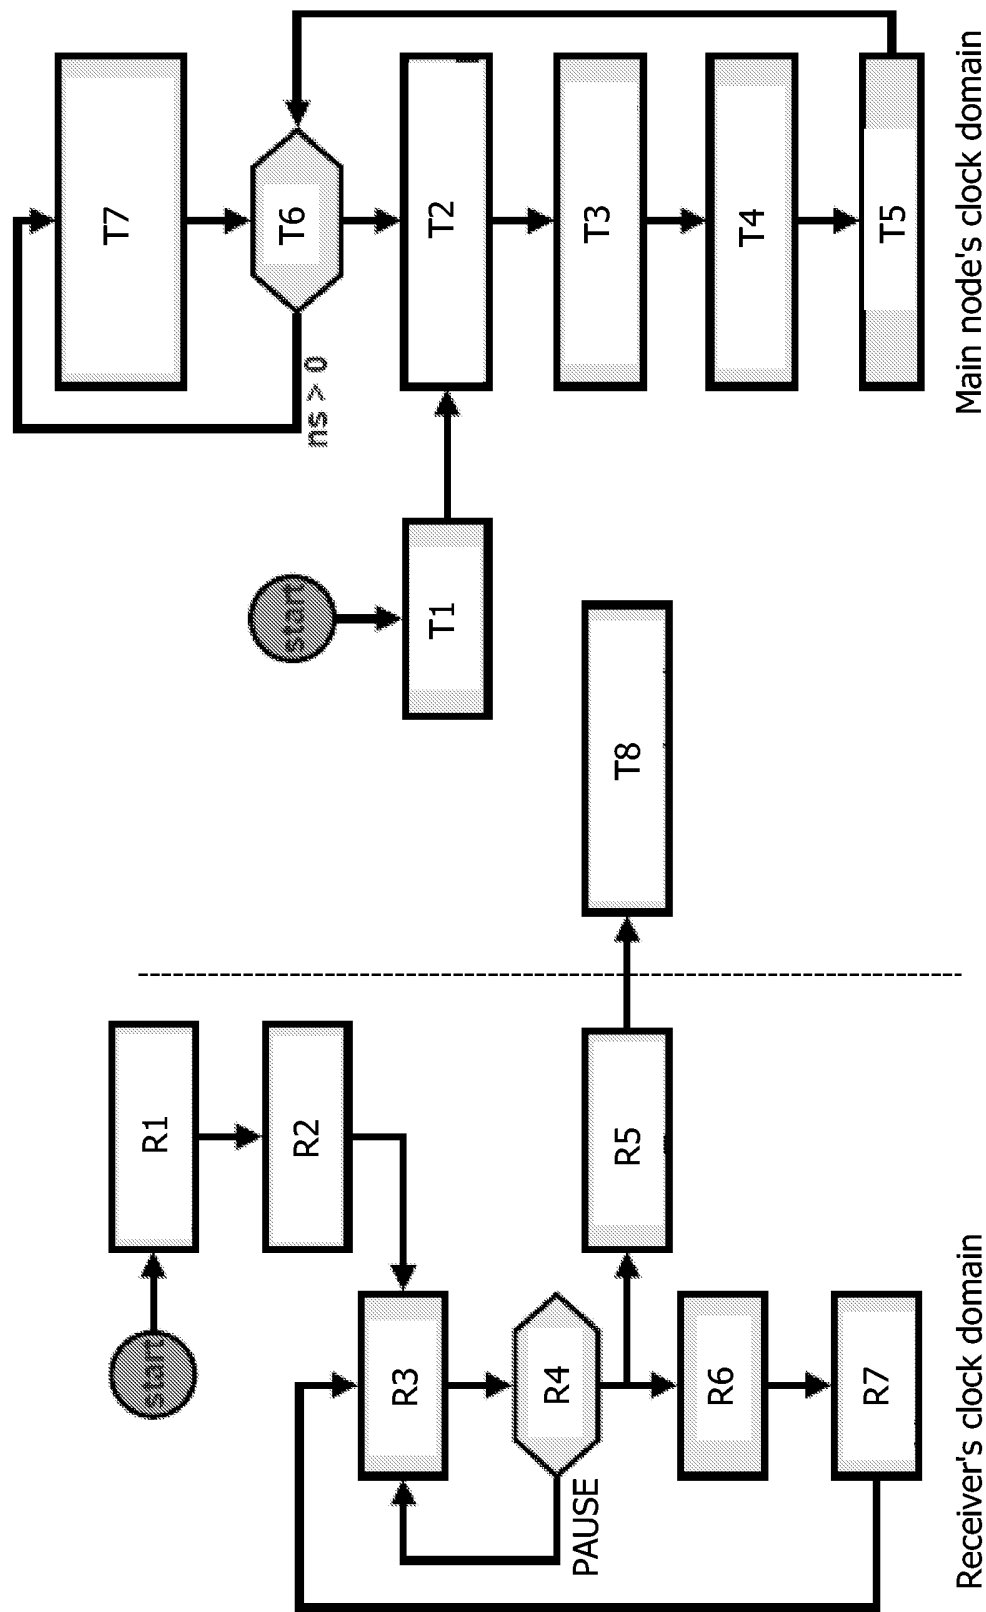


FIG.4B

FIG.4A




FIG.5

6/8




FIG.6

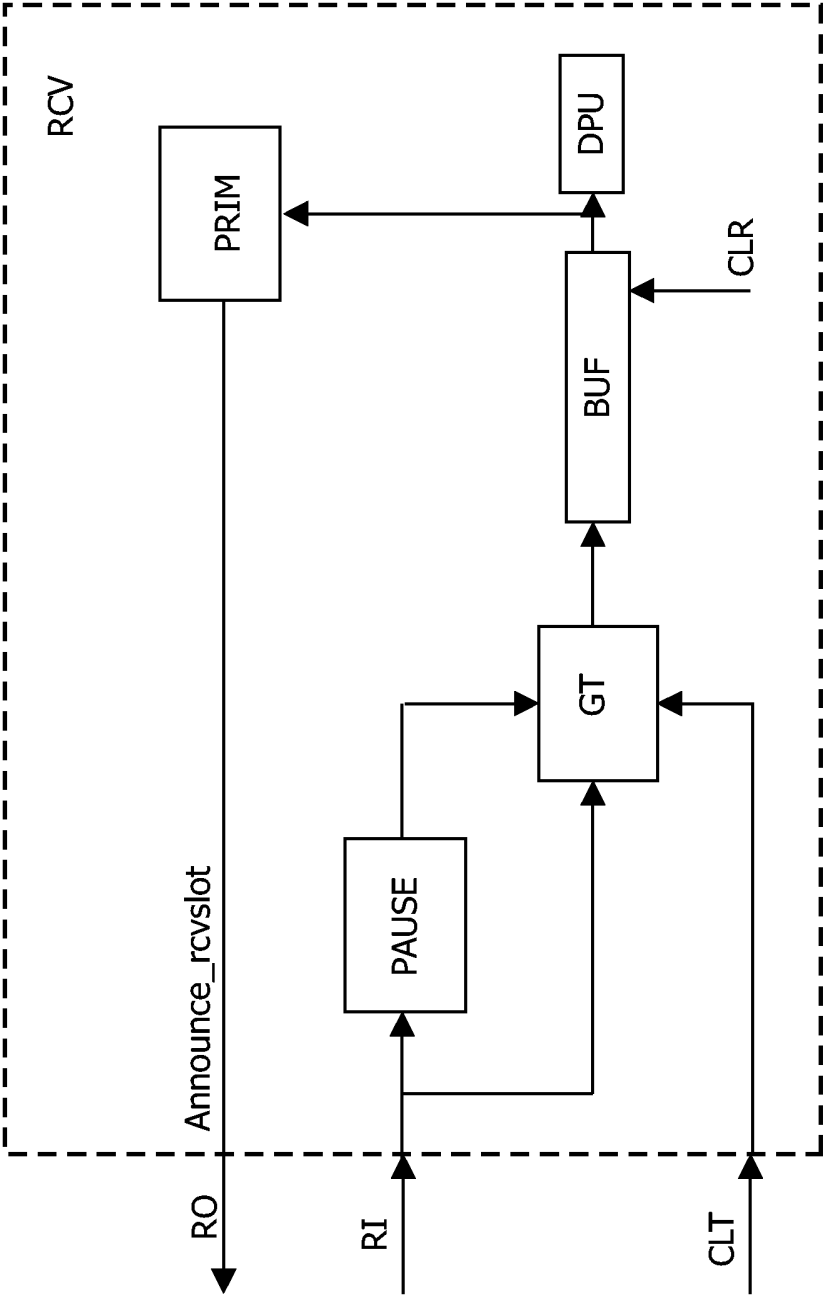


FIG.7




FIG.8

INTERNATIONAL SEARCH REPORT

International application No
PCT/IB2006/051856

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L7/02 G06F1/12 H04L12/56

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)
H04L G06F

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>DIELISSEN J ET AL: "Concepts and Implementation of the Philips Network-on-Chip" INTERNET CITATION, [Online] 13 November 2003 (2003-11-13), XP002330547 Retrieved from the Internet: URL: http://www.homepages.inf.ed.ac/kgoose/n/2003-ipsoc.pdf [retrieved on 2005-06-06] page 1, left-hand column, paragraph 3 - right-hand column, paragraph 2 page 2, right-hand column, paragraph 3 page 4, left-hand column, paragraph 2 - right-hand column, paragraph 1 figures 1,3,4,6</p> <p style="text-align: center;">----- -/-</p>	1-17

☒ Further documents are listed in the continuation of Box C.

☒ See patent family 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

17 October 2006

Date of mailing of the international search report

25/10/2006

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

Pieper, Thomas

INTERNATIONAL SEARCH REPORT

International application No
PCT/IB2006/051856

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 1 152 573 A2 (HEWLETT PACKARD CO [US]) 7 November 2001 (2001-11-07) paragraph [0003] - paragraph [0005]; figure 1 paragraph [0024] - paragraph [0026]; figure 3 paragraph [0036] -----	1-17
A	GB 2 362 777 A (3COM CORP [US]) 28 November 2001 (2001-11-28) page 3, line 7 - line 17 page 6, line 5 - page 7, line 4 page 8, line 1 - line 11; figures 1,2 -----	1-9
A	US 5 894 567 A (DODD JAMES M [US] ET AL) 13 April 1999 (1999-04-13) column 4, line 51 - line 65 column 5, line 43 - line 53 figure 2 -----	1-9

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/IB2006/051856

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
EP 1152573	A2	07-11-2001	US	6813275 B1		02-11-2004
GB 2362777	A	28-11-2001	US	6882661 B1		19-04-2005
US 5894567	A	13-04-1999	NONE			