RELAY UNIT BETWEEN A STATION AND COMMUNICATIONS CHANNEL, IN PARTICULAR FOR AN ETHERNET NETWORK

The invention concerns broadcasting communications networks, in particular those of the "ETHERNET" type. The general structure of the latter is defined by the standards IEEE 802.3 or ISO 8802-3 which are equivalent.

The network includes a communications medium or channel equipped with branches. In principle, with each branch there is associated a station. The latter should be able to emit messages on the network, and also to receive them by way of the network. A station may be constituted by any information technology system, whether standard or otherwise.

A branch comprises a transceiver on the channel side, and a "coupler" on the station side. The coupler/-transceiver exchanges, standardized, comprise bidirectional data signals, to which are added status signals, including an emission signal (TX), a reception signal (RX), and also a collision signal (COL). This notion of collision is particular to networks of the ETHERNET type, where emission is free in normal time. It corresponds (in simple terms) to the fact that two stations have sought to emit at the same moment. Finally, the coupler customarily provides the electrical supply to the transceiver.

Numerous applications have been developed according to the standard. It has also been proposed to integrate the coupler and/or the transceiver physically in the station, which then interfaces directly with the channel. The possibility of interconnecting the stations with one another is in principle still offered, since the coupler/-transceiver link is standardized.

But difficulties arise in certain cases:

- The standard provides for a collision resolution process termed "BEB", described hereinafter. This process does not

set any limit to the time which will be devoted to the resolution of a collision, which is very penalizing for certain applications. It is true that improvements to this protocol are known which can render it "determinist", in the sense that the time for resolution of a collision can still remain below a maximum fixed in advance (FR-A-2 597 686). But this is not compatible with the standard couplers conforming to the basic standard.

- It may be useful to render the local network redundant, by doubling the communications channel. But, in the present state of affairs, this entails the need to double the whole of the equipment, including the transceivers and the couplers, which throws the burden of managing this redundancy onto the user station. This also is not compatible with the current systems established according to the basic standard.

More generally, any modification applied with respect to the standard is likely to make it necessary to develop new equipment, specifically adapted to that modification and, at the same time, to render a great deal of existing equipment obsolete.

The present invention provides a solution to this problem.

One of the objects of the present invention is to provide a device which can operate by offering exactly the same interface to the user station, whatever the type of communications medium used, and its working mode. In other words, the station will not encounter any difference between the network proposed here and the known networks, conforming to the standard definition of the coupler.

Another object of the invention is to permit optimum use of the pass-band of the communications channel, in a manner which is "transparent" for the station.

The invention additionally aims to achieve these ends

independently of whether the communications medium is or is not split in order to function in redundant mode. It also aims to do this in a manner compatible with a "determinist" type of communications protocol.

Another object of the invention is to provide communications networks which can operate in real time with high security, including in demanding or difficult environments.

A further object of the invention is to permit the circulation of service messages, also in a manner which is "transparent" for the stations.

It is known to provide processing means between the coupler and the transceiver, in order to provide monitoring of the data and of the status signals exchanged between the latter. These processing means are distributed, in principle, between the coupler and the transceiver.

According to the invention, these processing means are installed in a relay unit, which is interposed between the station and the channel. The relay unit includes a memory, operated in a first in/first out mode. The processing means are arranged to use the memory in order to manage separately the relay/channel exchanges and the relay/-station exchanges, according to two modes which are at least partly different (the word "mode" here relates both to the physical characteristics of the network, especially of the channel, and the process or protocol according to which it is used).

Preferably, this management is effected in response to a condition bearing on the filling of the memory.

In one embodiment, the relay unit comprises:

- a connection to the station, which may comprise a station coder/-decoder,
- a connection to the communications channel, which may comprise a channel coder/decoder, and even incorporate the transceiver of the communications channel,
- on the channel side, the processing means co-operate with the transceiver of the channel in order to define channel

status signals, (RX, TX, COL), at the level of the communications channel. They respond to the fact that the channel status signals indicate reception, by temporarily storing the channel data in the memory with a view to submitting them immediately or subsequently to the station coder/decoder;

- on the station side, the processing means co-operate with the coupler of the station in order to define other status signals, termed station status signals (RX, TX, COL). They are capable of responding to a request for emission by the station by temporarily storing the emission data (station to channel) in the said memory, with a view to submitting them immediately or subsequently to the channel coder/-decoder.

In one variant, the connection to the station incorporates a "relay/station" transceiver, while the station comprises not only a coupler, but also a transceiver, and the station/relay link is then provided, between these two transceivers, by a means similar to a communications channel ("pseudo-network").

In another variant which can be combined with the preceding variant, the processing means are provided with a means for management of two channels in redundancy, on the channel side, the two channels being able to have dissymmetries of topology (difference in length, for example) and/or of nature (different materials, for example).

Other characteristics and advantages of the invention will be revealed by examination of the detailed description which follows and of the accompanying drawings, in which:

- Figure 1 is a general flow chart of an embodiment of a conventional coaxial cable network, of the ETHERNET type;
- Figure 2 is a diagram showing how a device according to the invention is inserted into such a network;
- Figure 3 is a diagram similar to Figure 2, but in the

case where a redundancy is used at the level of the coaxial cable;

- Figure 4 is a general flow chart of a coupler device of known type;
- Figure 5 is a general flow chart of a device according to the present invention;
- Figures 6 and 7 illustrate two variants of a mechanism operating the device according to the invention;
- Figure 8 illustrates a variant of the device according to the present invention;
- Figures 9 and 10 are two tables which are of use for the description of the management of the redundancy; and
- Figures 11, 12 and 13 respectively illustrate an emission mechanism, a reception mechanism and a divergence mechanism operating the device according to the invention during the processing of the redundancy.

The accompanying drawings are essentially definitive in character. They consequently form an integral part of the present description. They may therefore not only serve to make the latter more easily understood, but also contribute to the definition of the invention if need be.

The invention applies in particular to broadcasting communications networks (the messages are physically accessible to all the users, no matter for whom they are actually intended), and "contention" communications networks (the emission of messages is free, hence the possibility of emission collisions; in order to resolve the latter, some must abstain or restrain themselves from emitting, in accordance with a predefined protocol).

To this class belong the so-called CSMA/CD networks (Carrier Sense Multiple Access with Collision Detection), of which the so-called "Ethernet" networks form part. Broadcasting takes place on a communications medium or channel, which may be a coaxial link or "bus" (such as "10base2" or "10base5"), a twisted electrical pair ("10baseT"), or even an optical fibre (FOIRL"), for

example.

Figure 1 shows a part of a network, where the channel, of the bus type, is a coaxial cable.

The connection of a station to the channel comprises:

- on the channel side CC1, a transceiver (transceiver or "MAU", for Medium Attachment Unit) TR-i, charged with transforming the signals circulating on the channel into electrical signals of conventional format that can be used by electronic processing circuits, and vice versa.
- on the station side S-i, a coupler C-i, of which the link to the transceiver is standardized (level termed "AUI" for Attachment Unit Interface of the standards ISO 8802-3 or IEEE 802.3). In addition to the data signals, the coupler C-i and the transceiver TR-i exchange, through links in differential pairs, status signals, including an emission signal (termed TX or DO), a reception signal (RX or DI), and also a collision signal (COL or CO). Finally, the coupler customarily provides the electrical supply to the transceiver (lines termed VC, VP). Together these lines form a DC (Drop Cable) linking cable arrangement.

In certain cases, like those of the twisted pair or of the optical fibre, the link produced on the channel is of the "point-to-point" type. For point-to-point links, and for certain bus-type channels, transceivers are also provided, such as TR-R (Fig. 1), which act as amplifier or repeater (bidirectional) for the communications medium.

The structure of the messages obeys format rules, hence the frequent use of the word "frame" as synonymous with "message".

In an Ethernet network, the communications channel is, at a given moment and at a given point, in one of the following three states:

- channel empty (V), no emitter, no carrier and no collision;
- message (M) on the channel, 1 emitter, a carrier and no collision;

- collision (C), more than one emitter.

In the case of a collision, the status of the channel depends on different choices made when designing the network.

In a first type of network (coaxial link), the status of the channel reflects the existence of emissions, for example by the existence of a continuous voltage, which increases according to the number of emissions taking place. A station which wishes to emit can then immediately detect the fact that an emission is already taking place. Thus detecting a collision, it will reinforce it, for example by emitting a "jam" on the channel. This reinforcement lasts long enough for all the stations to be able to detect the collision.

In other networks (twisted pair, optical fibre), where the links are point-to-point, emission and reception coexist only in the sections of the channel where the emitting stations are located; in the other sections there is only reception. These signals are then allowed to spread freely on the channel, without intervening. The result is that a priori only the stations which have participated in the collision, by starting to emit, are aware of the existence of a collision.

In general, different means of making other stations aware of the collision may be envisaged, depending on the nature of the channel, and the design choices of the network:

- if, for example, the length of a minimum frame is at least equal to a threshold duration (51.2 microseconds according to the standard), the reception of a frame having a length which is less than this threshold duration then indicates the existence of a collision (this is the most current method for detecting a collision).
- The reception of undecodable information is in principle an indication, belated unfortunately, of a collision (this indication cannot be processed at the level of the

transceiver).

- conversely, a collision could be defined by the existence of a prolonged silence on the channel, especially in the case where systematic service emissions are carried out on this channel (information "apart from station data"), in the absence of emission by the stations themselves.

Let MPT be the maximum time taken by a frame to pass through the whole network (from one end to the other).

Time intervals are defined, being termed:

- "inter-frame spacing" or IFS
- "slot time" or SLT, which is at least equal to double the MPT time.

The inter-frame spacing IFS is linked to the emission/reception duality. When a station wishes to emit a frame, it first listens on the channel; if it is occupied, the station carries on listening; when the channel becomes free, the station must still wait for a time equal to the inter-frame spacing IFS (in order to take into account the time for passing from reception to emission at the level of the transceivers).

It should be noted that during its emission, a station could continue to listen, in order to compare what it hears with what it is emitting. Thus, it is still possible for a station which is emitting to detect a collision.

The slot time SLT is the critical time interval for collisions. At a point in the network, a silence at least equal to the MPT time signifies that no-one is emitting. But the most unfavourable collision occurs when a station A is emitting at one end of the network and a station B located at the other end starts to emit just before it has been able to detect the emission of A (therefore B is off-set by almost MPT with respect to A); at the moment (MPT lag) when it detects the emission of A, the station B becomes aware of the collision (and immediately emits the aforementioned "jam", if the structure of the network allows). But a time equal to 2 x MPT will pass before A

detects the emission of B (with or without jamming) and also detects the collision.

The basic Ethernet standard provides for the resolution of collisions by a standard procedure termed BEB (Binary Exponential Back-off). When a collision has occurred, all the stations retain the ability to emit, but not at any moment whatever. In any collision situation (the first and the following), each station which has, by emitting, participated in the collision taking place, will draw by lot a binary number BN in the range $[0, 2^k-1]$, where k is the minimum between 10 and the number i of consecutive collisions which have been perceived by the station under consideration. The station only has the right to emit, after the last collision perceived, at the end of a time equal to BN times the slot time SLT. To this is added a maximum accepted number of consecutive collisions (16 according to the standard).

The result is that it is impossible to guarantee a maximum time of waiting for emission, which raises a problem in certain applications, most particularly in "real time" networks. One imagines, in fact, that such a network cannot function without a guaranteed maximum waiting time.

The term "determinist" is used for a method of collision resolution (and, by extension, a network) which provides a guaranteed maximum waiting time. The solutions envisaged deviate from the BEB procedure, in the case of collision, and maintain the conventional CSMA/CD method for the remainder.

One of these solutions consists in assigning a number of its own to each station. After a collision (the end of the collision state), one first waits for an inter-frame spacing IFS; then the stations will count, incrementing each time that a characteristic time has passed, with the following rules:

- a station has the right to emit when its count is equal to its number,

- the characteristic time is SLT if no one emits, or, if someone emits, the duration of the frame emitted + IFS.

With SN being the number of stations, and TRMAX the maximum duration of the emission frame, the maximum time for resolution of a collision is the sum:
- of the maximum duration, restricted, of the first collision, and
- of the product SN * (TRMAX+IFS), in the case where the stations all wish to emit and are all in operation (it is possible to broadcast in the network, service information, including the number of stations in operation and the definition of their numbers).

In certain cases, in particular those of the real time type, it is of interest to apply this "determinist" solution and/or to provide a redundancy therein, that is to say, doubling (or more) of the communications channel.

But the existing or already designed networks are based on a standard interface, which is called a "coupler", standardized for dialogue with the transceiver which provides the physical link to the channel.

A priori, the application of the above-mentioned improvements to the existing or already designed networks assumes the complete remodelling of the couplers. Similarly, for future networks, it would be necessary to depart from the standardized couplers, which adversely affects the general applicability of networks thus designed.

The Applicant has discovered how to eliminate this difficulty.

One of the basic ideas of the invention (Figure 2) consists in providing, between the coupler C-i and the transceiver TR-i, a relay device ID-i, equipped with a memory. A distinction is made between the channel status signals RX1, TX1 and COL1 and the station status signals RX2, TX2 and COL2. If the channel is doubled for the

redundancy (Figure 3), there are then two sets of channel status signals, that is to say:

- RX11, TX11 and COL11 for the channel CC1, and
- RX12, TX12 and COL12 for the channel CC2.

This is made possible by the temporary storage of the channel data and/or the station data in the memory.

It will already be observed that this device makes it possible:

- on the one hand to simulate an emission for the station while data is being received from the channel (reception has priority); one will then transmit to it the data received while emitting one's own.
- On the other hand, it makes it possible to work with a station operating in the conventional BEB collision resolution mode, while in reality exploiting the channel according to a determinist collision resolution mode.
- Moreover, in the case of redundancy of the channel, it makes it possible to manage the redundancy in a manner which is "transparent" for the station.

An interesting variant of Figures 2 and 3 consists in integrating the transceiver (TR-i) or the transceivers (TR1-i and TR2-i) actually within the unit ID-i. Without adversely affecting the desired aim, this makes it possible to use, for the transceivers, circuits which are better integrated with the remainder of the block ID-i, and consequently to reduce the electrical consumption thereof. Preferably, the assembly is then fed from the coupler ID-i, in particular in the case of Figure 3, where this integra-tion is illustrated by the framework IDT.

Figure 4 illustrates an example of known structure of a coupler C-i, in the particular case of an industrial application in real time.

In Figure 4 the elements TR-(i-1), TR-i and TR-R are to be seen mounted on the coaxial cable CC1, as in the preceding figures.

According to the prior art, to the transceiver TR-i

there is directly connected the output of the coupler C-i, in this case defined by a MANCHESTER DO9 coder/decoder, which may be based on the circuit 82C501 sold by the INTEL CORPORATION.

In the upper part of Figure 4, an application bus AB9 is defined in order to operate according to the protocol TCP/IP (IP for "Internet Protocol"; TCP for "Transmission Control Protocol"). A person skilled in the art knows that the protocol IP makes it possible to convert all the services supplied by the network into a service termed "datagram", having a particular pack format. In its turn, the protocol TCP provides temporary safeguarding of the data, both in the reception mode and in the emission mode, as well as monitoring of their flow, and optionally multiplexing allowing a plurality of processes of the same host machine to use the same coupler.

Below the application bus AB9 appears a linking bus LB9, with, between these buses, an exchange memory ME9, which can be managed by a microprocessor control unit MP9.

The signals circulating on the bus LB9 are of a format directly usable by a block RE9, forming what is called a MAC ("Medium Access Control") half-layer.

This block RE9 may be based on the integrated circuit 82590 (or 82592, which differs from 82590 only by the absence of direct access memory or DMA - these circuits are sold by the INTEL CORPORATION). Such a circuit provides the functions of controller CSMA/CD, of collision resolution according to the aforesaid BEB protocol, of serialisation/deserialisation (passage series/parallel), of encapsulation/disencapsulation, as well as an error check of the CRC type ("Cyclic Redundancy Check"). The notion of encapsulation means that the actual messages are framed in monitoring information useful for their transmission on the network.

It is, for example, in such a context that the relay unit device ID-i according to the invention (Figure 5) will

be inserted between the coupler C-i and the transceiver TR-i, instead of these two elements being directly connected, as according to the prior art, by respective coder/decoders (not shown).

Figure 5 is a general flow chart of the relay unit ID-i according to the invention, in the case of a redundancy, where two communications channels are provided (see Figure 3). Here, the unit ID-i (hereinafter abridged to ID) can be broken down into coding/decoding units IDN1 and IDN2, respectively on the channel side and the station side, with, between the two, and from top to bottom, a data management unit IDD, and a channel control unit IDC.

Consequently (Figure 5, lower part), the unit IDN1 of the device ID includes two coder/decoders DO11 and DO12, respectively intended to be connected to those of the transceivers TR1-i and TR2-i in Figure 3. Similarly, (Figure 5, upper part), the unit IDN2 of the device ID-i includes a coder/decoder DO2. In one embodiment, these are coder/decoders using the code called MANCHESTER, which coder/decoders may be similar to the integrated circuit 82C501.

A description will now be given of the remainder of Figure 5, from top to bottom. After the unit IDN2, the series inputs/outputs of the coder/decoder DO2 are applied to the unit IDD. The latter starts with a bidirectional series/parallel interface ISP2 (series side at the top, parallel side at the bottom). Then, the unit IDD comprises a memory, which is here constituted by two memory blocks, one, FIFO_e, serving as a buffer memory for emission, that is to say, from the station to the channel, and the other, FIFO_r, serving as a buffer memory for reception, that is to say, from the channel to the station. As a variant, a single memory divided into two zones can be provided. At the bottom, the memories FIFO_e and FIFO_r are connected to another bidirectional series/parallel interface, denoted by ISP1 (parallel side at the top, series side at the bottom).

A microprogrammed monitoring unit UC1 manages the memories FIFO_e and FIFO_r, especially their addressing. It can also access the data themselves which these memories exchange (the data paths are illustrated in the drawing by thick lines).

The processing on the station side is carried out under the control of the upper level control unit UC2, which operates the data transfers executed by the emission buffer memory FIFO_e and reception buffer memory FIFO_r, and also by the two bidirectional series/parallel interfaces ISP1 and ISP2, and of course the monitoring unit UC1.

The exchanges between the unit IDD and the unit IDN1 are provided through a channel side exchange unit IDC, which comprises, from top to bottom:

- two collision detection modules MDC11 and MDC12, respectively connected to the two coder/decoders DO11 and DO12 of the unit IDN1, and
- a bidirectional internal multiplexer MI1, capable of connecting the two modules MDC11 and MDC12 to the series/parallel interface ISP1 of the unit IDD.

As a variant, the series/parallel interface ISP1 could be moved to the other side of the multiplexer MI, provided it was doubled. Moreover, when a single communications channel is used, the elements MDC12 and DO12 are disabled or omitted.

The functions of the block MI1 (bidirectional multiplexing) are to effect the transformation of the signals issuing from the two channels into a single group of signals, which also involves, in the embodiment described, managing the redundancy, which will be examined hereinafter.

Reference will now be made to the resolution of collisions.

This may be defined functionally as follows.

- the assembly IDN2+IDD (in detail, UC1, UC2, ISP2, DO2, FIFO and ISP1) carries out overall the same functions as the integrated circuit 82590 in Figure 4, except that the protocol BEB of the latter would be inhibited, so as to permit, within the unit ID, the immediate re-emission of the station data to the circuit MI1, save in the case where the latter prohibits it (the reception of messages from the channel has absolute priority in principle). At the level of the FIFO memories, the unit UC1 will then prohibit emission from the memory FIFO_e, while favouring reception to the memory FIFO_r.

- in its turn, the assembly IDC+IDN1 (in detail, MI1, MDC11, MDC12, DO11, DO12) manages the channel or channels according to any desired protocol, preferably with a determinist mode for the resolution of collisions.

Functionally, a distinction may therefore be made between channel side processing means (data management) MGC and station side processing means (data management) MGS. The memories FIFO_e, FIFO_r, their control or monitoring units UC1 and UC2, and also the series/parallel interfaces ISP1 and ISP2 straddle the two, as is shown by the broken line separating in Figure 5 the upward connections and the downward connections of the memories FIFO_e and FIFO_r.

Instead of being separated as shown, the memory blocks FIFO_e and FIFO_r may be combined in a single memory unit, managed on a first in, first out basis, but selectively for emission and reception, with priority to access in reception, that is to say, for the channel data. It is also of advantage for the memory or memories to be of the double port type, that is to say, they simultaneously permit reading access and writing access. This simultaneity may be physical or simply apparent, that is to say, the writing and reading operations are multiplexed sufficiently rapidly in time as to appear simultaneous at

the desired working speed.

The assembly MI1, MDC11 (optionally MDC12), DO11 (optionally DO12) may be combined in a single specialized integrated circuit (ASIC), to which will be added, as a variant, the transceiver or transceivers (TR; TR1 and TR2). The Applicant now manufactures such a circuit.

The channel side processing means MGC generally have priority for access to the memory FIFO_r for reception from the channel. Moreover, the control units UC1 and UC2 have full control over the sending of the contents of the emission memory FIFO_e to the channel. Consequently, the channel side processing means MGC can function in a generally conventional manner, except for:

- the fact that these means preferentially use a determinist collision resolution protocol, instead of the standard BEB collision resolution detection protocol, and
- the management of the redundancy.

It is therefore considered as accessible to a person skilled in the art to produce these channel side processing means MGC, with the hypothesis that they would use a non-redundant channel with a conventional protocol, of the BEB type.

It is also considered as accessible to a person skilled in the art to produce these same channel side processing means MGC, with the hypothesis that they would use a non-redundant channel with a determinist protocol, these latter being already known, and usable for example on the basis of the integrated circuits INTEL 82592 and equivalents, and/or of the above-mentioned French Patent.

The management of the redundancy will be dealt with hereinafter.

On their side, the station side processing means MGS are necessarily adapted according to the invention, since they do not have priority as a general rule. The mechanism of the exchanges between the station and the FIFO memory or memories should therefore:

- take into account precisely the priority in principle attributed to the channel side reception, and
- at the same time effect the exchanges with the station by simulating for the latter the normal behaviour of a transceiver, or, more generally, of a standardized station/-channel link.

A description will now be given of two examples of this mechanism, of which it is essential to understand that it concerns only the exchanges between the station and the memories FIFO_e and FIFO_r. Consequently, for example, "emission" means "emission from the station to the relay unit".

The mechanism preferentially used in connection with the embodiment in Figure 5 is illustrated in Figure 6.

It begins with a starting phase 100, which comprises steps of initialization of the device, with, at 102, the zeroing of a noted magnitude or variable Nbcoll (number of collisions).

The following step 110 consists in determining whether the reception memory FIFO_r is empty, that is to say, that it contains no data received from the channel and awaiting re-transmission to the station (alternatively, very few such data).

If the response to the test 110 is "no reception waiting", a step 120 detects whether a start of emission by the station is present, that is to say, whether the station is beginning to emit a frame to the unit ID (believing it to be the channel: this is therefore a pseudo-emission). In the absence of an emission request, one returns to the output of step 102. The steps 110 and 120 therefore define a kind of loop awaiting an occurrence, either of reception (very precisely, re-transmission of the memory FIFO_r to the station of data already received from the channel and stored in the memory FIFO_r), or of emission from the station to the memory FIFO_e of the unit ID.

In the case where the response to the test 110

indicated that data received from the channel are waiting, the step 112 produces a start of "reception" to the station ("ETHERNET subscriber"). The word "reception" is used, since the station believes itself to be receiving data from the channel, while it is a question of the re-transmission of information received from the memory to the station, therefore pseudo-reception. In practice, this start of reception consists in sending the customary preamble of the data, which one also terms encapsulation, but which does not comprise the useful data themselves.

If, during this sending of the preamble, the test 114 indicates that on its side the station has started to emit data to the unit ID, then the step 140 stops "reception".

Since only the preamble has been sent, these data have remained in memory, and thus remain available for a later attempt at re-transmission to the station.

In the absence of an attempt at emission by the station at the level of the test 114, the step 116 continues the reception, that is to say, a complete frame is re-transmitted from the memory FIFO_r to the station, through the devices ISP2 and DO2 in Figure 4, and monitored by the control unit UC2, while the memory is itself operated by the unit UC1.

After that, one passes to the final step 190 which is the customary wait for the time IFS, in order to return to the output of step 102.

If, in the absence of reception waiting at the test 110, the test 120 indicates a start of emission, then the step 122 inquires whether the level of filling of the emission memory FIFO__e is below a threshold, marked "threshold_e". If so, then the emission memory is regarded as being able to accept new data, and the step 124 continues with the emission by the station to the unit ID, that is to say, the entry into memory FIFO_e of the data coming from the station. At the same time, the units UC1 and UC2 co-operate in order to prohibit any action on the

memory FIFO_r, but only for emptying of the latter to the station.

There can then be provided a step 126 re-setting the magnitude Nbcoll to zero, after which one passes back to step 190 as before.

If the response to step 122 is that the memory FIFO_e is already well filled, there is in principle a "collision" on the station side (it is not a collision on the channel, but a pseudo-collision, simulated for the station, and simply due to the fact that the unit ID refuses to receive new data being emitted from the station).

As a preliminary, a test 130 determines whether the number of collisions counted earlier is equal to 15. If so, one is at the sixteenth collision, and, by absolute exception to the known BEB collision resolution protocol, one will then accept nonetheless the data emitted by the station, in order to store them in the memory FIFO_e, at the step 124, without indicating a new collision to the station (which, believing itself to be emitting on a channel operating in BEB mode, would not repeat its attempt to emit once again).

If, on the contrary, the number of collisions has not yet reached the value 15 (that is to say, the collision occurring is not the sixteenth), then the step 132 will simulate a collision to the station or ETHERNET subscriber, including, if need be, the sending of the "JAM". One therefore refuses the data coming from the station, in accordance with the BEB protocol. After that, the step 134 increments by one unit the number of collisions Nbcoll, and one passes back as before through the step 190 in order to recommence the general loop.

It is now possible to describe what happens after step 140.

After the stopping of reception, that is to say, the interruption of the attempt to re-emit data received from the channel to the station (data which, one recalls, have

remained in the memory FIFO_r up to this stage), one passes to step 122, in order to execute the station emission requested in step 114, under the same conditions as for the other emissions, as just described.

This mechanism has several particular features which are important, and, to a certain degree, independent:

- When the test of step 122 determines that the memory FIFO_e has not reached its filling threshold, there is, for the exchanges between the station and the memory, priority to the subscriber, that is to say, to the station, and not to reception (it is recalled that, on the channel side, for the exchanges between the channel and the memory, there is priority to the channel).

- If, on the contrary, the level of filling of the memory FIFO_e is above the threshold, then there is, for example, a simulated collision sent back to the station, in order to cause the latter to wait. An exception is provided, however, if an emission has already been refused fifteen times: one will then force its acceptance at the sixteenth time (at the sixteenth simulated collision), since it is the last occasion on which to use this emission, in default of which it would be irretrievably lost.

In fact (it is one of the advantages obtained by means of the device ID according to the invention), with suitable fixing of the filling threshold of the memory FIFO_e, one may allow oneself thus to accept supplementary emissions which one would otherwise have had every chance of losing.

For example, in a particular embodiment of the determinist collision resolution protocol, the relay unit according to the invention has a chance of being able to emit a frame for the benefit of the station approximately every 150 milliseconds (at the most) in the collision situation, and if all the frames effectively emitted are of maximum length. In its turn, as applied here between the

station and the relay unit ID, the BEB (modified) collision resolution protocol will create a waiting time varying, according to estimates made by the Applicant, from 0 to 350 milliseconds, with a mean value at approximately 175 milliseconds. The result of these two findings is that in principle no frame can be lost.

In other words, the relay unit/channel exchange occurs in the determinist collision resolution mode. In the latter mode, one is sure of being able to emit the emission frame, in the "slot" which is assigned to the relay unit, at the level of the determinist protocol. The fact that the memory FIFO_r is empty or practically empty means that there is no reception on the channel side; this in turn means:

- either that the channel is free, in which case the station can of course emit,
- or that the channel is in a collision resolution period, what is often called "era", in which case there is a time slot which is allocated to the station for emitting (the relay unit knows it but the station does not know it).

Insofar as the memory FIFO_r is empty or sufficiently empty, the Applicant proposes, preferably, to give priority to emission, in order to offer the station the possibility of making the best use of the pass-band, that is to say, the transmission capacities of the channel.

Although it has been indicated above that at the step 110 one could cause the test to bear upon the fact that the memory FIFO_r is almost empty, it is preferable, in the majority of cases, to determine simply whether the memory FIFO_r is strictly empty, so that the general rule remains that the relay unit ID according to the present invention acts in principle to re-emit immediately to the station the data received from the channel.

It is indicated above that the conventional BEB protocol is modified according to the present invention. This protocol is not in fact modified in itself, since the

station will apply it strictly. It is simply modified on the part of the dialogue between the relay unit ID and the station, since this relay unit ID almost always gives priority to the station, which is not true in the general application of the BEB protocol.

A description will now be given of a variant of the mechanism in Figure 6, with reference to Figure 7. This variant has two particular features.

The first particular feature is that the gangway between the emission and reception, defined in Figure 6 by the step 140, is omitted.

The second particular feature consists in that the operations termed "reception" are modified starting from step 112. The test 144 determines whether the reception started at 112 has been able to terminate correctly. If so, one returns to the final step 190. If not, one inquires at 146 whether there is a station/relay unit collision (this time, a true collision on the pseudo-network). In the absence of a collision, step 148 makes possible a return to the reception terminated test 144. If, on the contrary, step 146 indicates a collision, then at 150 one sends, if need be, a jam to the station. The output of step 150, and the output of step 144 indicating that the reception is terminated both lead to the final step 190.

The variant in Figure 7 may replace the mechanism of Figure 6. This variant is necessary in the case (Figure 8) where the link between the device ID-i and the coupler C-i is not produced directly, but through two transceivers TRK1 and TRK2, which effect the exchange between the two blocks ID-i and C-i in the manner of an ETHERNET network (or as a "pseudo-network"), on a "pseudo-channel" PSC.

The functions of management of the redundancy are described hereinafter with reference to Figures 9 to 13.

It is recalled that the word redundancy is used in the sense of doubling of the number of paths for emission,

reception and collision detection. According to the invention, the two paths used (also called medium) may have dissymmetries of topology (difference in length, for example) and/or in nature (different materials, for example). Consequently, it is no longer necessary to symmetrize the two paths, and that may allow the rejection of a certain number of common modes.

Up to now, the management of the redundancy was effected in general by means of a special redundancy pilot, external to the coupler.

The protocol proposed here by the Applicant is internal, and obeys the following rules:

- emission takes place systematically in an identical fashion on the two channels (or paths);
- for reception, one makes use of what one of the channels gives (from the point of view of the chronograms of the signals), while checking its coherence with what the other channel gives; a divergence is signalled as network status information (service information), at the end of a slot time;
- there is a collision status if a collision is observed on at least one of the channels;
- finally, a degraded mode is possible, with a single channel in service.

Save for this degraded mode, in terms of states (Figure 9) one has:

- an empty channel V if the two channels (path 0 and path 1) are empty,

- a message state M if the two paths have a message, or one has a message and the other is empty,

- a collision state C, in the other cases, where at least one of the paths is in collision.

This notion of states is to be completed by the consideration of the time intervals to which they apply. TM0, TC0 and TV0 are noted, time intervals during which the path 0 is respectively in the state M, C and V. The same is observed for the path 1. The final time interval, after management of the redundancy, is also similarly defined, without the final figure.

In that case (Figure 9):

- a message time TM is the minimum of TM0 and TM1,

- in the case of collision on one path, with message or collision on the other, the collision time TC is the combination of the time intervals TC0 and TC1, or TM0 and TC1 (for example),

- a collision TC1 or a message TM1 on one path while the other is empty manifest themselves through a collision time TC=TC1 or TM=TM1, respectively,

- if the two paths are empty during TV0 and TV1, there is an empty channel for a time TV equal to the combination of the time intervals TV0 and TV1.

It is considered as accessible to a person skilled in the art to produce the logic functions which are described above, preferably in microprogrammed form. The latter are incorporated in the multiplexer MI (in which one could distinguish a pure multiplexing function, and a multiplexing control function, incorporating the said microprogrammed logic).

Asymmetrical situations occur above on collision, in particular for couples such as TM0+TC1 (message on one channel and collision on the other), or TM1+TV0 (message on one channel and the other channel empty). These situations may be regarded as errors, the processing of which will

depend on the choices made during the designing of the network. It will still remain possible to pass into degraded mode, by using the channel which functions best (in theory, that which has no collision, in the two cases under consideration).

The mechanism of management of the redundancy according to the invention will now be described in more detail. Reference is made first to Figure 11 in order to describe the emission mechanism.

Whether one is in emission mode or reception mode, the internal multiplexer MI is in a mode waiting for a start of emission or reception. It is therefore (step 200) in a watching state. As long as it does not detect any start of emission, it remains waiting. On the other hand, when the multiplexer detects a start of emission from the half-layer MAC (Medium Access Control), it carries out a first test (step 201) in order to determine whether the two paths 0 and 1 are free (in other words, a check is made on the one hand that there is no modulation on the network, and on the other hand that the network is not in an IFS period). As long as the result of the first test 201 is negative, the said test (201) is re-started.

When this result is positive, MI gives the authorization to emit (step 202) to the half-layer MAC. Then (step 203), MI authorizes emission on the path or both paths according to their validity status. Thus, if the two paths are valid, the emission is authorized on each of them, while in the case of validity of only one of the two paths, the authorization to emit is given only for the validated path.

After a predetermined time interval, MI carries out a second test (step 204) in order to determine whether the MAC emission is terminated. As long as the result of the second test 204 is negative, the said test (204) is re-started. When this result is positive, then MI considers that the emission is terminated, and it reassumes the state

in which it was at the start of the request (before step 200), which brings the emission mechanism to an end.

Reference is now made to Figure 12 in order to describe the reception mechanism.

As in the preceding mechanism, the multiplexer is in a watching state at the start. As long as it does not detect any start of reception, it remains waiting. On the other hand, when the internal multiplexer MI detects a start of reception of signals from the network, it carries out a first test (step 250) in order to determine whether there is activity on one of the two paths, for example the path 0.

If the result of this first test 250 is positive, then reception takes place on the path 0. MI then proceeds to a second test (step 251) which consists in determining whether the reception on the path 0 is validated. In the case of a positive response, the path 0 is regarded by MI as selected (step 251) and one passes to the following step 260, to which one will return later. In the case of a negative response, one passes to the step 253, to which one will return later.

On the other hand, if the result of the first test 250 or of the second test 251 is negative, MI carries out a second test (step 253) identical to the first test 250 in order to determine whether there is activity on the path 1.

If the result of this third test 253 is negative, then MI reassumes its initial state (watching state).

On the other hand, if the result of this third test 253 is positive, then reception takes place on path 1. MI then proceeds to a fourth test (step 254) which consists in determining whether the reception on path 1 is validated. In the case of a positive response, path 1 is regarded by MI as selected (step 255) and one passes to the following step 260, to which one will return later. In the case of a negative response to the test 254, then MI reassumes its initial state in order to re-start the tests for activity

at step 250.

When one of the two paths has been selected (step 252 or 255), MI delivers (step 260) to the half-layer MAC, on the one hand, the data to be found on the selected path (0 or 1), and on the other hand the state of redundancy of the network according to the truth table in Figure 9.

Then MI carries out a fifth test (step 261) in order to determine whether the activity on the selected path (0 or 1) is terminated. As long as the result of the fifth test 261 is negative, the said test (261) is re-started.

When this result is positive, MI carries out a sixth test (step 262) in order to determine whether the activity on the path which had not been selected (1 or 0) is terminated.

If the result of this sixth test 262 is positive, then MI reassumes its initial state in order to re-start the tests for activity at step 250.

On the other hand, if the result of this sixth test 262 is negative, then MI proceeds with a seventh test (step 263) in order to determine whether there is new activity on the path initially selected (0 or 1). If the result of this seventh test is negative, then MI re-starts the fifth test by returning to step 261. On the other hand, if the result of this seventh test is positive, MI carries out an eighth test (step 264) in order to determine whether the reception on the path initially selected (0 or 1) is authorized. If the result of this eighth test is positive, then MI returns to step 260 in order to deliver to MAC the signals which have just arrived on the selected path. If the result of this eighth test is negative, then MI reassumes its initial state in order to re-start the tests for activity at step 250, which brings the reception mechanism to an end.

Reference is now made to Figure 13 in order to describe the divergence mechanism.

The internal multiplexer MI carries out a divergence

test starting from the moment when it has detected activity on at least one of the paths. It is in fact a matter of tests of coherence between the two paths which are carried out as soon as a frame is received. In order to do this, it determines in a first test (step 270) whether there is activity on one of the two paths, for example path 0, then, if this is not the case, it interrogates path 1.

If the result is negative on the two paths, then MI returns to its initial state before step 270.

On the other hand, if the result is positive (activity on at least one of the two paths), MI passes into a state waiting for diagnostic (step 271). After reception of the diagnostic (step 272), MI carries out a second test (step 273) in order to determine whether there is a divergence of status between the two paths 0 and 1. If the result of this second test 273 is positive, then MI produces a divergence signal (step 274), the consequence of which, for example, is to inhibit a path if a blocking anomaly (of the permanent collision type) is detected on the path concerned, then MI returns to the starting step 270. If the result of the second test 273 is negative, then MI proceeds to a third test (step 275) in order to determine whether the activity on the first path detected is finished.

As long as the result of this third test 275 is negative, then MI re-starts step 275. On the other hand, if the result of the third test 275 is positive, MI carries out a fourth test (step 276) in order to determine whether the activity on the path other than the first detected is terminated.

If the result of this fourth test 276 is negative, then MI proceeds to a fifth test (step 277) in order to determine whether there is new activity on the first path detected. In the case of a positive result, MI returns to the start of step 271 in order to await a new diagnostic. On the other hand, if the result of the fifth test 277 is negative, MI returns to the start of step 276 in order to

test the activity of the other path again.

If the result of the fourth test is positive, then the divergence mechanism comes to an end and MI returns to the start of the said mechanism, at step 270.

The Applicant has preferred to describe complete embodiments above, seeking to draw the maximum from the ETHERNET networks.

The relay unit ID according to the present invention may of course be used either solely for the passage from a conventional BEB mode of resolution between subscriber and relay to a determinist mode of resolution between relay and channel, or solely for the passage from a network mode with a single channel between station and relay unit to a mode with two redundant channels between the relay unit and these channels. More generally, the device according to the invention could be used in particular for any application which would require on the one hand management which was rigorously according to the standard, for the network protocol between the station and the relay unit, and on the other hand management which is modified with respect to the standard, between the relay unit and the channel or channels.

As already indicated, it is possible to broadcast, in the network, service information, including the number of stations in operation and the definition of their numbers. To that are added statistical data on the load of the network, for example, or other frames "apart from station data", emitted at a selected speed in order to check the status of the network. One of the advantages of the invention is that the relay units proposed (or some of them) can intercept these frames, which they use themselves, without transmitting them to their station.

C l a i m s

1. Electronic device forming a relay unit, intended to be inserted into the branch between a station and the communications channel of a broadcasting network, in particular of the Ethernet type, the device being characterized in that it comprises a connection (IDN2) to the station, a connection (IDN1) to the channel, a memory, having a channel zone (FIFO_r) and a station zone (FIFO_e), independently accessible in a first in/first out mode, and processing means (UC1, UC2; MI), arranged to use the two zones of the memory in order to manage separately the unit/channel exchanges and the unit/station exchanges, according to two at least partly different modes.

2. Device according to Claim 1, characterized in that the channel zone (FIFO_r) of the memory is sized to accept any reception, up to the vicinity of the maximum load of the communications channel.

3. Device according to one of the preceding claims, characterized in that the memory (FIFO) is a double port memory.

4. Device according to one of Claims 1 to 3, characterized in that the processing means (UC1, UC2) are provided with a determinist mode for the resolution of collisions on the channel side, and with the standard mode (BEB) for the resolution of collisions on the station side.

5. Device according to one of Claims 1 to 4, characterized in that the connection to the station (IDN2) comprises a station coder/decoder (DO2), and in that the connection to the channel (IDN1) comprises at least one channel coder/decoder (DO11, DO12).

6. Device according to Claim 5, characterized in that the connection to the station comprises a bidirectional series/parallel interface (ISP2), followed by the station coder/decoder, and in that the connection to the channel comprises at least one bidirectional

series/parallel interface (ISP1) followed by the channel coder/decoder.

7. Device according to one of Claims 1 to 6, characterized in that the processing means are provided with a means for management of two channels in redundancy (MI), on the channel side.

8. Device according to Claim 7, characterized in that the connection to the channel comprises two channel coder/decoders (DO11, DO12), respectively connected to two transceivers (TR1-i, TR2-i) capable of being mounted on two communications channels used in redundant mode, and fed directly in the relay unit.

9. Device according to one of Claims 7 and 8, characterized in that the two channels have dissymmetries of topology and/or nature.

10. Device according to one of the preceding claims, characterized in that the processing means comprise:

- channel side processing means (MGC), capable of
  - . co-operating with the channel connection in order to deal with channel status signals, representing reception (RX1), emission (TX1) or a collision (COL1), at the level of the communications channel,
  - . responding to the fact that the channel status signals indicate reception, by temporarily storing the reception data (channel to station) in the channel zone of the memory,
  - . attempting the emission of data contained in the station zone of the memory,
  - . processing the unit/channel collisions according to a channel mode,
- station side processing means (MGS), capable of
  - . co-operating with the station connection in order to deal with station status signals, representing reception (RX2), emission (TX2), or a collision (COL2), at the level of the unit/station exchange,
  - . responding to the fact that the station status

* C L M S *

signals indicate an attempt at emission by the station, by temporarily storing the emission data (station to channel) in the station zone of the memory,

. attempting the transmission to the station of data contained in the channel zone of the memory,

. processing the unit/station collisions according to a station mode.

11. Device according to Claim 10, characterized in that the station data processing means (MGS) are arranged to respond to a request for emission by the station by storing the emission data, or by refusing them with the production of a collision station status, according to whether the station zone of the memory is or is not below a selected filling condition.

12. Device according to Claim 11, characterized in that, the station side collision resolution mode being the standard mode (BEB), the station data processing means (MGS) are arranged to accept nonetheless a request for emission by the station at the end of a fixed number of station side collisions.

13. Device according to one of Claims 11 and 12, characterized in that, the connection to the station being produced by a station coder/decoder, capable of being connected to a coupler which the station comprises, the collision resolution mode of the station data processing means (MGS) is modified to tend to give priority to emission by the station, when the channel zone of the memory is below a selected filling condition.

14. Device according to one of Claims 11 to 13, characterized in that, the connection to the station being produced by a station coder/decoder, capable of being connected to a coupler which the station comprises, the station side processing means (MGS) attempt reception on the station side in the absence of a request for emission by the station, but interrupt this attempt, without

* C L M S *

emitting a collision, if a request for emission by the station occurs in the meantime.

15. Device according to one of Claims 11 to 14, characterized in that, the channel side collision resolution mode being of a type which systematically allocates a time slot to each station, the station data management means (MGS) accept the emission of the station data, in a channel side collision resolution situation, for their despatch into the said slot by the channel side processing means.

* C L M S *