



- (51) International Patent Classification:
H04L 12/46 (2006.01)
- (21) International Application Number:
PCT/GB2012/051321
- (22) International Filing Date:
12 June 2012 (12.06.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
1110009.6 15 June 2011 (15.06.2011) GB
11275096.3 15 June 2011 (15.06.2011) EP
- (71) Applicant (for all designated States except US): **BAE SYSTEMS plc** [GB/GB]; 6 Carlton Gardens, London SW1Y 5AD (GB).
- (72) Inventors; and
- (73) Inventors/Applicants (for US only): **FOUNTAIN, James Graham** [GB/GB]; Bae Systems, Warton Aerodrome, Preston, Lancashire PR4 1AX (GB). **EISSA, Rania Hamdi** [GB/GB]; Bae Systems, Warton Aerodrome, Preston, Lancashire PR4 1AX (GB). **HUDSON, Peter Noble** [GB/GB]; Bae Systems, Warton Aerodrome, Preston, Lancashire PR4 1AX (GB).
- (74) Agent: **BAE SYSTEMS plc**; Group IP Dept, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hampshire GU14 6YU (GB).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).


Published:

— with international search report (Art. 21(3))

[Continued on next page]

(54) Title: DATA TRANSFER

Fig. 1



(57) Abstract: A data transfer method and system includes at least one switch device (102) configured, in use, to transfer data, directly or indirectly, between a plurality of computing devices (104). The at least one switch device is configured to provide a plurality of Virtual Local Area Networks (VLAN), each said VLAN being configured as a Virtual Channel (VC) by receiving data from said computing device designated as a sole source of data over the VC. The VC data is transmitted from the sole source computing device as Ethernet frames (200) tagged with a VLAN tag (202). The VCs are routed through the at least one switch device according to a VLAN Id field of the VLAN tag of a said frame, with fixed and pre-determined routing of the frame being determined by configuration of the at least one switch device.



-
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

- 1 -

Data Transfer

The present invention relates to data transfer.

Packet switched networks have long been seen as a way of providing systems, including avionic systems, with very large bandwidths and connectivity, avoiding the bottlenecks and limitations associated with time division multiplexing (TDM) protocols like MIL-STD-1553B and STANAG 3910. These TDM protocols share network resources by scheduling transfers consecutively, which requires a significant level of co-ordination between the disparate sources of the data. This co-ordination of the transmissions onto the network is not generally a problem where the network is based on a shared media. In this case, a centralized control function, like a Bus Controller, can be used to control all transmissions and ensure that there are no contentions. However, this co-ordination can be a very significant problem in packet switched networks. This is because, in packet switched networks, such a centralized control function is difficult to implement (and would be extremely restrictive if it were), and distributing such a function requires close synchronization of the entire network and thus severely restricts the scheduling mechanisms that may be used in such a system.

An option usually used in commercial packet switched network protocols, such as Ethernet, is to allow all users to transmit at any time. FIFO queues are then used to buffer traffic at the shared resources whenever the loads are temporarily greater than the resource can cope with. However, these buffers will always have a limited ability to cope with the peaks in the traffic, due to their finite size. Hence, if they are overloaded, i.e. the peaks in the traffic are too

- 2 -

great or last too long, the resource becomes congested, the buffer overflows, and data is consequently lost. Even where the buffer does not overflow, it may become so full that the delay through it is excessive. Both of which are generally unacceptable in real-time systems

5 One known method for preventing this congestion is to limit the bandwidths of the traffic sharing the resources and predict the maximum load or buffer usage this will cause. The predictions can then be compared with the actual sizes of the buffers to show whether overflow is or is not possible and the delay caused by this usage can be used to show whether transfers can be or
10 cannot be excessively delayed.

 The buffering does introduce an additional delay that does not occur in TDM protocols; however, this delay is very small in comparison with the delays that can occur at the sources when resources are shared by TDM, e.g. when data waits for its allotted transmission time slot, which can be up to the
15 transmission interval of the data (e.g., 20 ms for 50 Hz data transfers). Since, in commercial packet switched networks, the users can transmit as soon the data is ready, these networks suffer no such transmission delays.

 The known ATM protocol is in many ways well suited to the transport of data in real-time systems, as it has all the necessary properties to route traffic
20 deterministically and limit bandwidths. This allows the loads on shared resources and consequent delays to be predicted. ATM does this routing using Virtual Channels (VCs), which are connection oriented "pipes" that carry the data along deterministic routes. The mechanism used to limit the bandwidths in ATM is hardware based traffic shaping at the source and traffic policing in the

- 3 -

network: Traffic policing in the network effectively fixes the size of the pipe that is a VC, and traffic shaping at the source ensures that the source does not try to pump too much data down this pipe.

While ATM can provide deterministic data transfer, it has failed to spread
5 significantly beyond its niche market in telecommunications infrastructures, partly because the complexity of the hardware in the network interfaces made them significantly more expensive than the equivalent Fast/Gigabit Ethernet interfaces. Whereas, Ethernet is a much more widely used, virtually ubiquitous, packet switched network protocol in the commercial sector. Hence, it is
10 desirable to take advantage of the breadth of components, such as processor boards, etc, that already incorporate interfaces for the Ethernet protocol. This can reduce initial costs of avionic systems and reduce the costs of ownership by reducing obsolescence issues and easing performance upgrades, i.e. providing higher network bandwidths.

15 Ethernet networks that will need to carry this type of deterministic data will, in some cases, already exist and will already be carrying uncontrolled or best efforts traffic when the requirements to support deterministic data transfers emerge. As a result, the protocols that will allow such a network to meet deterministic data requirements will have to be applied after its inception; the
20 only other alternative would be to provision for it, e.g. with modified/bespoke interfaces, initially, in the expectation of these requirements arising at some later data. It is therefore important that the process of upgrading an existing network should have the minimum possible impact on the system and on the pre-existing best efforts traffic, i.e. there should be no need to modify

- 4 -

subsystems that do not need to take part in these new transfers. Additionally, in some instances, the security of the data will have to be proved, i.e. that it is authentic and resistant to unauthorized access or modification.

There can be significant advantages if systems and subsystems are able
5 to use more advanced scheduling mechanisms than the cyclic executive, such as priority based pre-emptive scheduling. There are also advantages in systems supporting software re-use, if data transfers can be made at the natural rate of the application, which may not be harmonically related to other application rates. For example, adding an application that has a natural rate of 60 Hz to a
10 system where most applications run at 50 Hz and the sub-harmonics thereof, which is a difficult task where cyclic scheduling is used.

It is therefore desirable that any deterministic data transfer should use standard Ethernet interfaces; allow for data transfers with requirements for security, reliability, and or timeliness; be able to be applied to existing networks
15 with the minimum impact, i.e. allow uncontrolled, legacy traffic to continue to use the network and support the maximum possible range of scheduling mechanisms, and allow it to be proved in advance that these deterministic requirements will be met. Optional requirements are for resistance to system faults, i.e. over-transmission by subsystems, and network faults, i.e. failures of
20 the network components – switches, network interfaces, and cables.

US2006/002370 describes a control plane/transport plane architecture for controlling a carrier network formed by Ethernet switches. Carrier edge switches may be logically separated into single Provider Edge (PE) Core and one or more PE-Edge functions. The PE-Core function can be employed to

- 5 -

encapsulate incoming Ethernet traffic from the customer using MAC type encapsulation and to forward the encapsulated traffic across the carrier network. VLAN tags are used to provide customer separation at the logical PE-Core with each different customer site connected to each edge switch having a
5 unique VLAN tag. The logical PE-Core of an edge switch encapsulates each Ethernet frame in a further Ethernet frame using the MAC address of edge switch as the source address and the MAC address of the appropriate egress point (the destination address). At the PE-Core of the edge switch, the original frames are stripped of their encapsulation and sent over a communications link
10 via the PE-Edge of edge switch to a customer switch. Every switch, receiving an Ethernet frame on one port, broadcasts the Ethernet frame out on every other port. The process repeats as the frame is received by other switches. Thus, the frame is broadcast across the entire network. MAC address auto-learning functionality is provided to improve configuration
15 efficiency in switched Ethernet networks. Ethernet frames have source and destination MAC addresses corresponding to their source and destination Ethernet switches. When an Ethernet frame sent out by a source switch is received by a receiving intermediate or destination Ethernet switch, the receiving switch observes the port on which the frame was received and the
20 source address of the frame. The receiving switch then builds up a forwarding table for use in future frame switching. The VLAN tag defined in IEEE 802.1q is applied such that the Ethernet switches of carrier network are 802.1q VLAN-aware but arranged to use a combination of destination address and VLAN tag to forward data traffic. This is preferentially achieved by reusing the existing
25 capabilities in each Ethernet switch to store separate forwarding tables for each

- 6 -

VLAN tag configured, the VLAN tag acting as a mapping (or indexing) to forwarding tables.

WO2006/070197 discloses a communications scheme for configuring a network comprising a plurality of connected switching apparatus, the scheme comprising: determining in a control plane index header field values to identify connectionless traffic received at switching apparatus for which a connection is to be established between a source node and a destination node; providing each switching apparatus necessary to implement the connection with information from the control plane, the information enabling the data forwarding tables of the switching to be populated with said index header field values in association with egress ports of the switching apparatus; and disabling all other functionality on said switching apparatus capable of populating the data forwarding tables with index information associated with said egress ports of the switching apparatus necessary to establish said connection. Forwarding tables of the switching apparatus are directly provided with address information associated with outgoing ports of the switching apparatus. In some embodiments the control plane populates the entries in the forwarding table with at least one other Ethernet header field in addition to the destination address field and further associates a VLAN with an outgoing, or egress port of the switch. This VLAN-Id is used to distinguish between multiple paths across the communications network.

US2005/157737 describes a network including an edge node configured to define per hop behaviours using a set of bits in an Ethernet header of a frame

- 7 -

and a core node configured to receive the frame and to forward the frame according to the per-hop-behaviours.

US 7 369 495 relates to a method and a network device for sharing bandwidth among a group of classes of traffic for an interface. Bandwidth may
5 be allocated to at least one traffic class of a first priority for the interface. At least some unused bandwidth of the at least one traffic class may be allocated to at least one other traffic class of a second priority for the interface.

In US2003/0154259 a VPN service is provided through a shared network infrastructure comprising a plurality of interconnected provider edge devices
10 having customer edge interfaces. Some of the CE interfaces are allocated to a VPN supporting a plurality of virtual local area networks (VLANs) and are arranged for exchanging traffic data units with CE devices. Each traffic data unit includes a VLAN identifier. At least one virtual connection is established in the shared network infrastructure between two CE interfaces of different PE
15 devices, allocated to said VPN, for forwarding traffic data units including a VLAN identifier. Said virtual connection (VC) has an identifier determined from said VLAN identifier and an identifier of said VPN.

US2004/0213254 discloses a method for servicing subscriber ends by utilizing virtual LAN on ATU-R (ADSL Termination Unit – Remote) of ADSL,
20 which utilizes VLAN to generate a plurality of virtual connections in an ADSL connection, couples equipment of each subscriber end to a plurality of subscriber input/output (I/O) ports in ATU-R, adds a switching hub in the ATU-R for identifying tagged data in the virtual connections, connects each subscriber end in the virtual connections with the switching hub via Ethernet, and assigns

- 8 -

each subscriber I/O port in the switching hub to one of the virtual connections having a distinct tag.

None of the above patent publications specifically suggest designating one computing device as the sole source of data over the VC in a network in the same way as the present invention and so do not solve the technical problems
5 relating to security and congestion as described herein.

Embodiments of the present invention are intended to address at least some of the problems discussed above and provide for the transfer of data with requirements for reliability, timeliness, security, and or fault tolerance, across a
10 packet switched Ethernet network, in parallel with ordinary, best efforts data, and using standard Ethernet network interfaces. The embodiments can use VCs for reliable, secure and/or timely carriage of data over a full duplex switched gigabit Ethernet network using the Virtual Local Area Network (VLAN) protocol to identify the VCs and the Class of Service (CoS) protocol to separate
15 the data on these VCs from other data that do not have such requirements. The embodiments can be used to prove to a customer or certifying authority that data transfers between standard Ethernet interfaces in a switched Ethernet network will always meet their requirements for the security (the sender can rely on the data going only to where it should, and the receiver can rely on it having
20 come only from where it should), reliability (that the data will not be lost due to the effects of other traffic, i.e. there can be no congestion, in the network) and/or timeliness (that the data arrives within a specified deadline, relative to its transmission time).

- 9 -

According to a first aspect of the present invention there is provided a data transfer system including:

at least one switch device configured, in use, to transfer data, directly or indirectly, between a plurality of computing devices,

5 wherein the at least one switch device is configured to provide a plurality of VLANs, each said VLAN being configured as a VC by receiving data from a said computing device designated as a sole source of data over the VC, the VC data being transmitted from the sole source computing device as Ethernet frames tagged with a VLAN tag,

10 the VCs being routed through the at least one switch device according to a VLAN Id field of the VLAN tag of a said frame, with fixed and pre-determined routing of the frame being determined by configuration of the at least one switch device.

In embodiments comprising a plurality of the switch devices, each said
15 switch device may use VLAN trunking to allow for transmission on, and reception from, a plurality of the VLANs, thereby allowing the plurality of switch devices to be connected together. In this case the switch devices can be configured not to insert or strip the VLAN tag of a said frame.

A said VLAN tag may use any available Tag Protocol ID (TPID) value
20 and may designate at least one said computing device other than the sole source computing device as at least one destination for the VC data. Alternatively, where the VC data have requirements for security, the VC may be limited to a said computing device designated as a sole destination for the VC

- 10 -

data (to avoid destinations receiving spurious data from other, faulty or malicious, destinations, as may occur on conventional VLANs).

At outputs of the at least one switch device, the VCs can be assigned to and routed through a VC buffer. The routing through the VC buffer may be by
5 means of the CoS protocol described in IEEE 802.1d, &co. A Priority Code Point (PCP) field in the VLAN tag may be used to identify the VC data for routing through the VC buffer. The buffer shall be given priority over all other buffers associated with the output of a said switch device, such that no other (e.g. non-critical) traffic is transmitted from the output when there is any
10 Ethernet data frame held in the VC buffer.

Each said VC shall be allocated a maximum bandwidth. The at least one switch device may employ traffic control measures that limit each said VC individually to its said allocated maximum bandwidth, e.g. by delaying and/or discarding any traffic that is in excess of this limit (determined, e.g., by arriving
15 too close to the preceding frame of the same VC, based on the bandwidth and tolerated jitter for that VC). This may be done, e.g., using the known token bucket or leaky bucket algorithms.

A maximum fill level for each of the said VC buffers may then be calculated and the maximum fill level may be compared with actual sizes of the
20 buffers to avoid, or reduce a probability of, congestion within or between the buffers assigned to the VCs. The end-to-end delays for the VCs can also be calculated, using the fill levels for the buffers through which it is routed, and compared with any deadline requirements for the VCs.

- 11 -

The at least one computing device may comprise an, e.g. Avionic Line Replaceable Units (LRUs).

The VLANs will normally comprise IEEE 802.1q VLANs. The plurality of switches will normally be configured to provide an IEEE 802.3 full duplex
5 switched Gigabit Ethernet network.

According to a further aspect of the present invention there is provided a communications network including at least one switch device substantially as described herein and a plurality of said computing devices.

According yet another aspect of the present invention there is provided a
10 switch device configured substantially as described herein.

According to another aspect of the present invention there is provided a method of providing data transfer between computing devices, the method including:

configuring at least one switch device, in use, to transfer data, directly or
15 indirectly between a plurality of computing devices,

wherein the at least one switch device is configured to provide a plurality of VLANs, each said VLAN being configured as a VC by receiving data from a said computing device designated as a sole source of data over the VC, the VC data being transmitted from the sole source computing device as Ethernet
20 frames tagged with a VLAN tag,

the said VCs being routed through the at least one switch device according to a VLAN Id field of the VLAN tag of a said frame, with fixed and pre-

- 12 -

determined routing of the frame being determined by configuration of the at least one switch device.

According to another aspect of the present invention there is provided a computer program element comprising: computer code means to make the
5 computer execute a method substantially as described herein. The element may comprise a computer program product.

Whilst the invention has been described above, it extends to any inventive combination of features set out above or in the following description. Although illustrative embodiments of the invention are described in detail herein
10 with reference to the accompanying drawings, it is to be understood that the invention is not limited to these precise embodiments. As such, many modifications and variations will be apparent to practitioners skilled in the art. Furthermore, it is contemplated that a particular feature described either individually or as part of an embodiment can be combined with other individually
15 described features, or parts of other embodiments, even if the other features and embodiments make no mention of the particular feature. Thus, the invention extends to such specific combinations not already described.

The invention may be performed in various ways, and, by way of example only, embodiments thereof will now be described, reference being
20 made to the accompanying drawings in which:

Figure 1 schematically illustrates a communications network including a plurality of switch devices configured to transfer data between a plurality of computing devices, and

- 13 -

Figure 2 illustrates the format of an Ethernet frame.

Referring to Figure 1, there is shown a communications system including a set of switch devices 102A – 102C. Each switch device is connected to at least one computing device 104A – 104F, which may be a conventional
5 computing device, e.g. an Avionic Line Replaceable Units (LRU), and may be configurable to act as a networked host. The system components are used to provide an IEEE 802.3 full duplex switched Gigabit Ethernet network in which a plurality of VLANs, as described in IEEE 802.1q, &co, are used as VCs for the transfer of data. The hosts in this system may, through the use of VLAN
10 trunking, be enabled for the transmission on and reception from a plurality of VLANs.

A VLAN is configured as a VC by having one source of critical data from amongst the computing devices 104. This data is transmitted in the form of Ethernet frames tagged with a VLAN tag, which may use any available TPID
15 value. The VC may have a plurality of destinations for this data. However, where the data have requirements for security, the VC may be limited to a single destination from amongst the computing devices 104 to avoid destinations receiving spurious data from other, faulty or malicious sources as may occur on VLANs. These VCs/VLANs, being routed through the switch
20 device(s) 102 on their VLAN Id (which is contained within the VLAN tag, as discussed below) and have fixed and pre-determined routes. The predetermination of the routes is contained within the configuration of the switch device(s).

- 14 -

Critical data can be transferred over these VCs only from the authorized source and only to the authorized destinations. However, this will apply only where the configuration of the network switches is suitably secure. Specifically, the configuration of the switch(es) through which the VCs are routed must be
5 limited to avoid VLAN hopping attacks, e.g., double tagging and switch spoofing, and can only be modified by authorized means. The limits on the configurations are likely to include such requirements that no switch output through which one or more VCs are routed may strip VLAN tags. The limits on means of configuration are likely to include that, e.g., the Dynamic Trunking
10 Protocol or Multiple VLAN Registration Protocol cannot be used to modify the VLANs being used as VCs (which prevents the critical data VCs being subject to switch spoofing attacks), and the methods of configuring the switches for the VLANs being so used must be suitably secure. There may also be a requirement that the interfaces of hosts that receive these critical data over the
15 VLAN VCs shall only receive Ethernet frames that are tagged with VLAN tags, including any non-critical data they may be required to receive, where the VLANs used for the non-critical data will be of lower priority (have a lower PCP value) than the critical data (which avoids the switch outputs connected to these hosts from needing to strip default VLAN tags from non-critical data, and thus
20 prevents the hosts and critical data VCs being subject to double tagging attacks).

For those VCs that are configured with a single source and single destination, the constraints on how the VCs are routed through the switch(es), i.e. using the VLAN id, will ensure that data received by the destination must

- 15 -

have come from the expected source: frames that contain the correct VLAN id that are transmitted by other, erroneous or malicious, hosts will not be allowed to join this VC, by virtue of there being no routing data supplied to the switch for the inputs to which said erroneous or malicious hosts are connected. These
5 VCs are thus, excepting in circumstances involving failures of the network switches and or the data protection scheme incorporated into the Ethernet standard, immune to data spoofing attacks.

The switch devices 102 may be core switches and subsystem switches, which meet the basic requirements of support for IEEE 802.1q (including VLAN
10 trunking), and for IEEE 802.1p, allowing for at least two levels of priority with separate buffers for each and with the high priority traffic buffer having simple priority over the lower. The specifics of IEEE 802.1q support are that the switches have to be configurable so that frames arriving at a port from the connected host tagged with a specific VLAN Id are routed only to a subset of
15 the other switch ports. Ideally, the set of ports that a given VLAN Id is routed to should be specific to each input, i.e. the port id augments the VLAN Id. However, in common with normal usage of VLANs in COTS (commercial off the shelf) networks, it is acceptable that the association is made in configuring these other ports to have these frames routed to them. The difference being
20 that the former method would allow limited re-use of VLAN Ids within a switch.

The switch devices 102 also have to be configurable so that untagged frames arriving at a port from the connected host can be tagged internally with a default/native VLAN Id specific to that port, and so that frames internally routed

- 16 -

to it from other ports that are tagged with a default/native VLAN Id, specific to that port, can have their tags stripped before retransmission.

The switch devices 102 also provide wire speed operation, i.e. they provide a core switching bandwidth that is at least equal to the sum of the input
5 bandwidths. They use shared, centralized buffering or individual output buffers, and use store and forward switching. The wire speed operation will ensure that there is no input buffering in the switches that could lead to head of line blocking. The use of shared/centralized or output buffers allows the buffer usages to be calculated and compared with the available buffer sizes. The
10 requirement for store and forward switching ensures that frames are not routed on corrupted VLAN Ids, etc.

The switch devices 102 may be configured off platform, or out of band by a single subsystem on platform. They may also be configured in band by any appropriately secure method, e.g. the known Simple Network Management
15 Protocol. In all cases, other unused means of switch configuration, specifically the autonomous configuration of VLANs and traffic policing, including DTP and MVRP, will be disabled, i.e. the VCs implemented by the methods described herein shall be equivalent to the so called Permanent VCs (PVCs) used in ATM networks rather than the so called Switched VCs (SVCs).

20 The switch devices 102 are configured to provide a full duplex packet switched network implementations of Ethernet. It will be appreciated that the system shown in the Figure is exemplary only and the number and arrangement of switching and computing devices can vary. For instance, the simplest possible architecture comprises a star, with each networked host 104

- 17 -

connected to one of the ports of a single switch device (this can be duplicated, etc, for redundant networks). Where there are more hosts to interconnect than may be conveniently connected to a single switch, multiple switches may be used. These switches can then be connected together, to allow a host
5 connected to one to communicate with hosts connected to another. These inter-switch links can then restrict transfers, if there is a need for a significant number of connections between hosts on one and hosts on others. This may be eased by using higher bandwidth connections for these inter-switch links, i.e. use 10 Gigabit Ethernet in a generally Gigabit Ethernet network.

10 If more than two switch devices 102 are interconnected then this can create rings: e.g. A connects to B connects to C connects to A, etc. Such rings in a single broadcast domain can cause broadcast traffic to be repeatedly regenerated and generally flood the network. A known solution is to use the Rapid Spanning Tree Protocol to detect such loops and shut down one of the
15 links in the loop, thus breaking it and prevention the build-up of broadcast data. For the VLAN VCs described herein, it is assumed that either these will be configured without loops or the Multiple Spanning Tree Protocol will be used to break any loops that are formed. This may also be a potential method of implementing network redundancy, except that it can take at least 3 seconds to
20 re-activate a link should there be a physical layer failure.

The illustrated system comprises existing components of the Ethernet standard protocol: use of standard device drivers and COTS operating system, etc. The present inventors investigated the Ethernet standards and found what was apparently a means of implementing VCs already present in IEEE 802.1q:

- 18 -

Virtual (bridged) Local Area Networks (VLANs). Initially, this was seen as a way of implementing a fully COTS equivalent of the known AFDX protocol by using the VLAN Id instead of the VLink Id. However, after realizing that this would require existing networks to be upgraded so that all traffic used VLANs, the
5 inventors realized that IEEE 802.1p Class of Service (CoS) could be used to allow the traffic on the VLANs to be separated from the best effort traffic as well as one another, if this best effort traffic is assigned or defaults to a lower priority.

Figure 2 illustrates the format of a VLAN tagged frame 200, including a 4 byte VLAN tag 202. Ethernet frames transmitted on a VLAN are identified by a
10 specific 12-bit VLAN Id value 204 in the tag 202, which is inserted into the Ethernet frame in front of the Ethertype field. Another 3-bit PCP field 206 indicates one of 8 classes of service. The single bit Canonical Format Indicator (CFI) 208 is not used in this embodiment. Some small quantity of supporting data, e.g. responses to the Address Resolution Protocol or Neighbor Discovery
15 Protocol requests, will/may pass from the destination(s) to the source and which will also be tagged with the ID for the VLAN associated to the VC.

Conventional VLANs were developed to interconnect groups of hosts as though they, and only they, are attached to the same broadcast domain, regardless of their physical location. Hence, a VLAN has the same attributes as
20 a physical LAN, but it allows for hosts to be logically grouped together even if they are not located on the same network switch. Several of these VLANs can then share a physical domain, e.g. be connected to the same physical switch, each being isolated from the others. Switch devices that support IEEE 802.1q VLANs do so through a set of virtual switches, one for each VLAN, with the

- 19 -

ports of the physical switch being associated with a number of these VLAN virtual switches.

Normally, switches that support VLANs convert untagged traffic to a default VLAN on its arrival, i.e. add a tag indicating VLAN Id 1, and, generally, strip this tag off on transmission. They can also be configured to add a tag with a different VLAN Id value at a port, only route VLANs with this tag to that port and strip these on transmission. This allows a host connected to such a port to be on a VLAN, but unaware of this, with the switch doing all the work. This is known as port-based VLANs and port-based VLAN switching. Thus, the host connected to this port is on the VLAN associated with the port, but unaware of it, and does not need to comply with IEEE 802.1q, etc., as it neither adds VLAN tags to its transmissions nor receives any tagged frames. This, however, is not how VLANs are used in the present embodiment.

The present inventors considered each VLAN being isolated from all the others in port-based VLAN switching to be a limitation where there is a need to connect a host to multiple VLANs, e.g. where this host is a server, etc, or where multiple switches need to be connected together, each supporting a fraction of one or more VLANs. Hence, they decided to use another known mode of operation, called VLAN trunking. In this mode, switch ports can be associated with multiple VLANs, and can leave them tagged when transmitting these frames. A host connected to a trunking port then has to add and strip the VLAN tags itself, and can concurrently transmit on and receive from multiple VLANs. This is how VLANs are used for deterministic traffic in the present embodiment.

- 20 -

In the present embodiment, the switching of VLANs carrying deterministic data is always assumed to be using VLAN trunking for the deterministic data, and VLAN tags identifying these VLANs should never be inserted or stripped by the switch. As a result of this virtual switching, a
5 (trunked) VLAN is used to carry only one data type, where only one host is used as the nominal source of this data type and all the other connected hosts are then destinations of this data type, becomes, in all significant regards, a VC.

As the VLANs are routed through a virtual switch, rather than a virtual hub, the routing of traffic identified with a VLAN is also modified by the MAC
10 destination address. Hence, a frame identified with a VLAN and containing a specific MAC address will only go to the port associated with the VLAN and connected with the specified MAC address. However, a broadcast MAC address can be used, when the frame will be transmitted to all ports associated with the VLAN (except the one it is received on).

15 As stated above, the VLAN Id field 202 is only 12 bits in size, which limits the number of separately identified VCs to 4094 (a couple of values are reserved). If the identification of the VCs is required to remain consistent across the whole system, and this is the only identification, then this means only 4094 VCs in the entire system. However, this would be an artificial limit because
20 VLAN Id values can be reused. With COTS VLAN compliant switches, the real limitation is that there can be no more than 4094 VCs routed through any switch. Therefore, in larger networks, different VCs that are routed through different switches can reuse the same Id. Requirements for unique identification of data could then be met using, e.g., a message type Id in the payload.

- 21 -

Switch devices that comply with IEEE 802.1p CoS can be configured to transmit IEEE 802.1q tagged traffic through separate output buffers. This can be done by configuring the source to set the PCP value in the tag to a high value (e.g. 7) and configuring the switch to associate this service class with the high priority buffer. The switches can also be configured to service these high priority buffers with strict priority over the buffer used for the best efforts traffic, which can be defaulted to a lower priority. Hence, when a tagged, prioritized frame arrives at an output, it will be transmitted with a delay caused only by the other high priority traffic and at most one low priority frame (if there was one just started when the first high priority frame arrives).

At the output of each switch device 102, the VCs are assigned to and routed through a buffer, e.g. using the Class of Service (CoS) protocol described in IEEE 802.1d, &co. This buffer is given simple priority over all other buffers associated with the same switch output, such that no other, e.g. non-critical, traffic is transmitted from that output when there is any Ethernet data frame held in the VC buffer.

Each VC can be allocated a maximum bandwidth that it may use. The network switches may employ traffic control measures that limit each VC individually to its allocated bandwidth, e.g. delaying and or discarding any traffic that is in excess of this limit, determined, e.g., by arriving too close to the preceding frame of the same VC, based on the bandwidth and tolerated jitter for that VC. This may be done, e.g., using the known token bucket or leaky bucket algorithms, or an equivalent, i.e. an algorithm that strictly limits the nominal data rate and the maximum jitter of the traffic on an individual VLAN VC, separately

- 22 -

between each input and the buffers. The specific requirements for the range and resolution of the bandwidth (nominal interval) and jitter tolerance can vary. However, the range on the nominal interval (which is equivalent to AFDX's Bandwidth Allocation Gap) should encompass intervals down to at least 12 μ s
5 and up to one second. Jitter tolerance should encompass a range down to zero and up to, at least, several milliseconds (which is assumed to be sufficient for multi level switch hierarchies up to at least 4 layers deep).

Simplistically, the token bucket and leaky bucket algorithms measure the average rate at which data arrives on a VC and the variation around this,
10 essentially how much earlier they arrive than when they should, based on the nominal bandwidth. Basically, frames should arrive at the nominal rate, but some can arrive early, due to jitter. However, if it arrives more than the expected jitter before an "ideal arrival time" (derived from the nominal rate), it has to have been transmitted too soon, i.e. at too high a bandwidth, or jittered
15 more than allowed for, and it therefore it is dropped by the traffic policing function to protect the buffer.

This method proposed for system fault tolerance (traffic policing in the switch inputs), however, is not supported by current Ethernet protocols, at least not in relation to VLANs and CoS. However, the functions that provide this
20 feature in the present embodiments are contained entirely within the core network switch or switches. Hence, this feature can be implemented while still using only standard features of Ethernet in the subsystems. It does mean modifying the core network switch(es) where system fault tolerance is required,

- 23 -

but as there will be significantly fewer of these than subsystems, this should be acceptable in most cases.

To ensure that these traffic policing functions do not discard legitimate traffic, the bandwidth limit to which a VC is policed can be set slightly higher
5 than the allocated bandwidth for the VC, i.e. the bandwidth the source is allowed to use on the VC. Similarly, the jitter limit can be set slightly higher than the sum of the variations in the scheduling at the source of the VC and the delays between the point where the transmission is scheduled and the policing function, which has to include any variation in the delay between the switch
10 input and the policing function. However, this delay may be difficult to measure in isolation, either as part of characterization or validation against a specification. Hence, it may be necessary to use an overestimated value, e.g. the variation in the delay between the input and output of the switch, which will necessarily include delays that precede the policing function and delays that
15 follow it.

The delays through the switches will comprise two separate types: one type which is dependent on the traffic flowing through the switch (and thus specific to the system and network configuration), and one type which, while potentially variable in nature, is not dependent on the traffic flowing through the
20 switch and may, therefore, be characterized in isolation from a configuration. By appropriate specification of the switch devices, e.g. having wire-speed operation and either output buffers or a centralized and shared buffer that may be modelled as a set of virtual output buffers each associated with a specific switch output, it is possible to ensure that the buffering in the switch associated with

- 24 -

the switch outputs is the only source of delays that are dependent on the traffic flows and, moreover, ensure that these delays depend only on the traffic routed through the output with which the buffer is associated.

As a consequence of the predetermination of the routes and the
5 constraints applied to the bandwidths of the VCs, the fill levels for and consequent delays caused by the switch buffers in the network may be predicted. The maximum fill levels, either absolute values or ones with a given probability of occurrence, may then be compared with the actual sizes of the buffers to ensure that there can be no congestion within or between the buffers
10 assigned to the VCs or that congestion has a only given probability of occurring. Hence, the transport of these VLAN tagged frames over a VC will, excepting losses due to noise and component failures within their routes or hosts exceeding the allocated bandwidths of other VCs, be entirely reliable or, where predicted fill levels are probabilistic, have a given probability of success.

15 Since the bandwidth of each VC can be limited to a known level, it is possible, using known methods, e.g. network calculus or an equivalent method, to predict in advance what the maximum fill level of each of the buffers in the network might be for a given set of VCs. These values can then be compared to the available buffer sizes to check that there can be no congestion, i.e. the
20 buffers can never overflow, or it is a sufficiently low probability that they will. This can be used to show the reliability of the transfers. The delays these buffers might ever cause can also be calculated at the same time, and combined with the other delays, which can be characterized in isolation from the set of VCs sharing the network resources, to give the maximum end to end

- 25 -

delays, or values that it is a sufficiently low probability will be exceeded. These can then be compared with any deadline requirements for timely data transfers, to show these will be met.

The maximum delay caused by the buffers in the route of a VC can than
5 be added to the non-traffic dependent delays for the same route to give a maximum end-to-end delay in transportation through the network. Similarly, the variations in these delays may be summed for the route of a VC to give a maximum jitter in the delivery of the Ethernet frames carrying the critical data. These VCs thus provide real-time performance.

10 The use of traffic control in the network switches to limit transmissions on the VCs will ensure that, excepting in circumstances involving failures of the network switches and or the data protection scheme incorporated into the Ethernet standard, the predicted maximum fill levels and delays caused by the buffers in the switches will not be exceeded even in situations where the hosts
15 transmitting on other VCs shall exceed their allocated bandwidths. The predictions that there will be no congestion and for the maximum end-to-end delays will, excepting the circumstances just listed, remain valid even in such situations. These VCs are thus, excepting the circumstances just listed, immune to denial of service attacks causing congestion and consequent data loss in the
20 buffer's switch output.

As a result of using bandwidth and jitter in traffic control, rather than the absolute transmission times for the data, it is possible to use a wide variety of methods for scheduling the transmission of critical data, which should be an enabler to using advanced scheduling techniques for the applications that are

- 26 -

the sources and sinks of this data. Specifically, the commonplace Cyclic Executive method can be supported, as well as more advanced techniques such as pre-emptive priority based scheduling, as long as it is possible to specify the minimum interval between transmissions on a given VC. Also, because the system described herein is based only on the use of frequency domain techniques in the control and prediction of critical performances, transmissions may use non-harmonically related update rates. This is in contrast to methods that use temporal domain techniques, i.e. synchronize the transmission times over the entire network to avoid contentions at the shared resources (switch buffers). In these time domain controlled approaches, transmission rates must, essentially, be subdivisions of a base rate and all transmission rates are thus necessarily harmonically related. For example, the highest rate may be set at 50 Hz; in which case, the only lower transmission rates that are available will be 25, 12.5, and 6.25 Hz, etc. If this is not so, then contentions will occur at some rate determined by their different frequencies. In a network using the scheme described herein, however, one VC may use a transmission rate of 50 Hz while another, which may have the same source and or destinations, can use 60 Hz.

Network redundancy for providing network fault tolerance is not supported by the current Ethernet protocols. However, again, it is possible to implement this feature, with some costs, but without diverging from the use of COTS and standard components in the subsystem. One possible way of implementing this form of network redundancy would be in the application layer. One such solution is to employ multiple VCs, one through each of the identical

- 27 -

networks. These multiple VCs can extend from the applications in the source to the applications in the destination. The destination applications can then take data from only one of these, e.g., which ever delivers first. In the context of IMS (integrated modular systems) and the TLS (three layer stack), it may be seen
5 that this makes the application layer network dependent, as it would not need to do this were the network itself fault tolerant, etc. However, if the requirements for this fault tolerance are limited to a fraction of possible users, it can be the case that supporting network redundancy in the MSL/OSL makes these layers platform dependent, or at least adds complexity to them that is not a general
10 requirement. In which case, using multiple parallel VCs, despite the additional functionality in the application layer, may be an acceptable solution.

The use of existing components of the Ethernet standard allows the system described herein to be fully compatible with both COTS hardware and COTS software. For example, both Wind River's VxWorks™ (Network
15 Equipment Edition) and Green Hill's Integrity support™ IEEE 802.1q and IEEE 802.1p. Also, the IMS three layer stack (TLS) currently uses VxWorks™ device drivers, and so should also provide the required support. Greenhills Integrity™ and some versions of Linux also provide support for VLANs and CoS. The embodiments described herein can therefore be used to add deterministic data
20 transfers to new network implementations, i.e. new systems and systems upgrades; and can also be applicable to systems using existing switched Ethernet networks (and a compatible OS) with a minimum of impact on the existing traffic and operation of those systems. They can also allow new systems or new networks to be implemented knowing that such transfers can

- 28 -

be added at a later date without the need to add new interfaces, when and if the requirements for these transfers emerge. Device drivers, etc, for the network interfaces that are not intended to send or receive deterministic traffic will have no additional special requirements in relation for use in embodiments of the
5 system and so any IEEE 802.3 compliant driver should be sufficient.

In addition to providing means to prove that criticality requirements will always be met, the embodiments described herein also allow existing Ethernet networks may be upgraded to support critical data transfers. This is because the embodiment only uses the IEEE 802.1q and IEEE 802.1d protocols to
10 identify and route the VCs and to keep the critical data separate from the non-critical, ordinary Ethernet traffic. This system of VCs may, therefore, be implemented using only Ethernet interfaces and device driver software that complies with the IEEE 802.3-2002 or subsequent standards, incorporating support for IEEE 802.1q VLANs and IEEE 802.1d CoS. It may, without
15 providing tolerance to faults in the other hosts, be implemented using a switch that complies with the IEEE 802.3-2002 or subsequent standards, incorporating support for IEEE 802.1q VLANs and IEEE 802.1d CoS. Such switches being widely available both for commercial network installations and in rugged form for avionic applications, etc. The provision of tolerance to such system faults, as
20 excessive transmission of critical data, will require modification of the switches used in the system, but will not require modifications of or to any of the hosts, neither those that transmit and receive using the VCs described herein nor those that transmit and receive not using them, so long as these hosts and their software support VLANs and CoS. This method for supporting critical systems

- 29 -

may therefore be imposed, as an appliqué protocol, on an existing full duplex switched Gigabit Ethernet network with minimal impact to the existing host LRUs, and thus minimal costs.

Subsystems on some platforms may comprise multiple modules and
5 contain small switches to interconnect these and connect them to a core network switch or switches. Others can comprise a single module, and thus no switching. Such subsystem switches will still need to separate the best efforts and deterministic traffic (this is true for any appliqué protocol that adds deterministic transfer to best efforts traffic, such as TTE, as well). In the case of
10 the embodiment described herein, this requires support for both IEEE 802.1q and IEEE 802.1p, as it must recognize and use the PCP in the VLAN tag. For other appliqué protocols, the mechanisms for identifying deterministic data, and thus how the switch determines the priority of traffic, are generally specific to the protocols, and will thus generally require a bespoke switch in this context.

15 Embodiments of the data transfer system described herein can have the following advantages:

- It is cheap:
 - Cheap to implement or add as an upgrade – components can be COTS or MCOTS.
 - Cheap to own – it will remain compatible with future
20 subsystem components, i.e. Ethernet chipsets on processor cards, minimizing obsolescence issues associated with the network.

- 30 -

- Cheap to upgrade once implemented – minimizes impact of upgrading to higher physical layer bandwidth (e.g. 10 Gigabit Ethernet) and or a larger network (e.g. multiple switches).
- 5 • Impact on system and subsystem operation is minimized:
 - Neither prescribes nor proscribes methods for traffic scheduling – no need to schedule cyclically or synchronize transmissions by different subsystems – direct support for advanced scheduling methods, e.g. priority based
 - 10 preemptive scheduling.
 - Supports acyclic/sporadic transfers
 - Supports harmonically unrelated transfer rates (no need for divided rate methods).
- Ideally suited to information pipeline systems:
 - 15 ○ Data can be transferred as soon as processed (with preemptive schedule, may also be processed as soon as received).
- Low impact from changes - add traffic without changes to other subsystems:
 - 20 ○ Only need to recalculate buffer usages to confirm still no congestion and deadlines are still met.

- 31 -

- Compatible with standard higher layer protocols, e.g. Internet Protocols (UDP, TCP, etc.)

- 32 -

CLAIMS

1. A data transfer system including:

at least one switch device (102) configured, in use, to transfer data, directly or indirectly, between a plurality of computing devices (104),
- 5 wherein the at least one switch device is configured to provide a plurality of Virtual Local Area Networks (VLAN), each said VLAN being configured as a Virtual Channel (VC) by receiving data from a said computing device designated as a sole source of VC data over the VC, the VC data being transmitted from the sole source computing device as Ethernet frames (200)

10 tagged with a VLAN tag (202),

the VCs being routed through the at least one switch device according to a VLAN Id field (204) of the VLAN tag of a said frame, with fixed and pre-determined routing of the frame being determined by configuration of the at least one switch device.
- 15 2. A system according to claim 1, including a plurality of the switch devices (102), each said switch device using VLAN trunking to allow for transmission on, and reception from, a plurality of the VLANs, thereby allowing the plurality of switch devices to be connected together.
3. A system according to claim 2, wherein the switch devices (102) are
20 configured not to insert or strip the VLAN tag (202) of a said frame.
4. A system according to any one of the preceding claims, wherein a said VLAN tag (202) uses any available Tag Protocol ID (TPID) value and

- 33 -

designates at least one said computing device (104) other than the sole source computing device as at least one destination for the VC data.

5 5. A system according to any one of claims 1 to 3, wherein the VC is limited to one said computing device designated as a sole destination for the VC data to avoid destinations receiving data from other sources.

6. A system according to any one of the preceding claims, wherein at an output of a said switch device (102), the VCs can be assigned to and routed through a VC buffer and the routing through the VC buffer is by means of a Class of Service (CoS) protocol of IEEE 802.1d.

10 7. A system according to claim 6, wherein a Priority Code Point (PCP) field (206) in the VLAN tag (202) is used to identify the VC data for routing through the VC buffer.

8. A system according to claim 7, wherein the VC buffer is given priority over any other buffers associated with the output of the switch device (102),
15 such that no other traffic is transmitted from the output when there is any Ethernet data frame held in the VC buffer.

9. A system according to any one of the preceding claims, wherein each said VC is allocated a maximum bandwidth.

10. A system according to claim 9, wherein the at least one switch device
20 (102) employs traffic control measures that limit each said VC individually to its said allocated maximum bandwidth, e.g. by delaying and/or discarding any traffic that is in excess of the maximum bandwidth.

- 34 -

11. A system according to claim 10, wherein the traffic control measure comprises a token bucket algorithm or a leaky bucket algorithm.
12. A system according to any one of claims 9 to 11, wherein a maximum fill level for each of the said VC buffers is calculated and the maximum fill level is
5 compared with actual sizes of the buffers to avoid, or reduce a probability of, congestion within or between the buffers assigned to the VCs.
13. A system according to any one of the preceding claims, wherein the at least one computing device (104) comprises an Avionic Line Replaceable Unit (LRU).
- 10 14. A communications network including at least one switch device (102) and a plurality of said computing devices (104) according to any one of the preceding claims.
15. A method of providing data transfer between computing devices (104), the method including configuring at least one switch device (102), in use, to
15 transfer data, directly or indirectly between the computing devices according to any one of claims 1 to 13.




Fig. 2

Preamble	Start-of-frame delimiter	MAC destination	MAC source	VLAN Tag		Ether-type or Length	Payload (data and padding)	CRC32	Interframe gap
7 octets of 10101010	1 octet of 10101011	6 octets	6 octets	Tag Protocol Id	P C F P	2 octets	46-1500 octets	2 octets	12 octets of idles
				Bits		2 octets			
				3	1	12			

202 204 206 208

200

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2012/051321

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/46
ADD.

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

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 practicable, 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	US 2006/002370 A1 (RABIE SAMEH [CA] ET AL) 5 January 2006 (2006-01-05) paragraph [0023] - paragraph [0027] paragraph [0079] - paragraph [0084] paragraph [0090] paragraph [0095] paragraph [0106] - paragraph [0134] paragraph [0195] - paragraph [0196]; claims 1-13; figures 2,4,6,7,15,27,28 ----- -/-	1-15

☒ 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 application or patent 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

20 November 2012

Date of mailing of the international search report

03/12/2012

Name and mailing address of the ISA/

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

Authorized officer

Camba, Sonia

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2012/051321

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2006/070197 A2 (BRITISH TELECOMM [GB]; MCGUIRE ALAN [GB]; REID ANDREW BRYSON DICK [GB]) 6 July 2006 (2006-07-06) page 4 - page 9 page 13 page 17 - page 25 page 39 page 52 - page 56 claim 11; figures 2,6a-6c,7-9 -----	1-10,14, 15
X	US 2005/157737 A1 (RABIE SAMEH [CA] ET AL) 21 July 2005 (2005-07-21) paragraph [0006] - paragraph [0017] paragraph [0028] - paragraph [0030] paragraph [0033] - paragraph [0036] paragraph [0039] - paragraph [0059] paragraph [0062] - paragraph [0068]; claims 1-17; figures 1-2 -----	1,4-10, 12,14,15
X	US 7 369 495 B1 (LEMAIRE THOMAS A [US] ET AL) 6 May 2008 (2008-05-06) column 2 - column 3 column 11 - column 13; figures 4,6 -----	1,4-9, 14,15
A	"IEEE Standard for Local and Metropolitan Area Networks Virtual Bridged Local Area Networks;IEEE Std 802.1Q-2005 (Incorporates IEEE Std 802.1Q1998, IEEE Std 802.1u-2001, IEEE Std 802.1v-2001, and IEEE Std 802.1s-2002)", IEEE STANDARD, IEEE, PISCATAWAY, NJ, USA, 1 January 2006 (2006-01-01), pages _1-285, XP017603885, ISBN: 978-0-7381-4876-2 page 17 - page 29 page 34 - page 38 page 74 - page 78 page 96 -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2012/051321

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006002370	A1	05-01-2006	NONE
WO 2006070197	A2	06-07-2006	
		AU 2005321093 A1	06-07-2006
		BR PI0519612 A2	25-02-2009
		CA 2590669 A1	06-07-2006
		JP 2008527772 A	24-07-2008
		KR 20070095374 A	28-09-2007
		US 2008049621 A1	28-02-2008
		WO 2006070197 A2	06-07-2006
US 2005157737	A1	21-07-2005	
		CN 1910856 A	07-02-2007
		US 2005157645 A1	21-07-2005
		US 2005157721 A1	21-07-2005
		US 2005157737 A1	21-07-2005
		US 2011051723 A1	03-03-2011
US 7369495	B1	06-05-2008	
		US 7369495 B1	06-05-2008
		US 7835279 B1	16-11-2010
		US 2011019572 A1	27-01-2011