



US006717950B2

(12) **United States Patent**
Lui et al.

(10) **Patent No.: US 6,717,950 B2**
(45) **Date of Patent: Apr. 6, 2004**

(54) **METHOD AND APPARATUS FOR PRIORITY-BASED LOAD BALANCING FOR USE IN AN EXTENDED LOCAL AREA NETWORK**

6,147,971 A * 11/2000 Rochberger et al. 370/238
6,343,330 B1 * 1/2002 Khanna et al. 709/249
6,498,795 B1 * 12/2002 Zhang et al. 370/400

(75) Inventors: **King-Shan Lui**, Urbana, IL (US);
Whay Chiou Lee, Cambridge, MA (US)

(73) Assignee: **General Instrument Corporation**,
Wilmington, DE (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **10/206,553**

(22) Filed: **Jul. 26, 2002**

(65) **Prior Publication Data**

US 2003/0193959 A1 Oct. 16, 2003

Related U.S. Application Data

(60) Provisional application No. 60/350,572, filed on Jan. 20, 2002.

(51) **Int. Cl.**⁷ **H04L 12/28**

(52) **U.S. Cl.** **370/408; 709/252**

(58) **Field of Search** 370/351, 389,
370/392, 400, 401, 406, 408, 428; 709/252,
238, 240

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,018,137 A * 5/1991 Backes et al. 370/401
5,323,394 A * 6/1994 Perlman 370/238
5,734,824 A * 3/1998 Choi 709/224
5,850,397 A * 12/1998 Raab et al. 370/392
5,884,043 A * 3/1999 Teplitsky 709/238
6,081,511 A * 6/2000 Carr et al. 370/256

OTHER PUBLICATIONS

Lee, W.C.; Spanning tree method for link state aggregation in large communication Networks. INFOCOM '95. Fourteenth Conference of the IEEE Computer and Communications Societies. 1995, pp. 297-302 vol. 1.*

Lee, W.C.; Routing subject to quality of service constraints in integrated communication networks IEEE Network, vol.: 9 Issue: 4, Jul./Aug. 1995 pp.: 46-55.*

* cited by examiner

Primary Examiner—Hassan Kizou

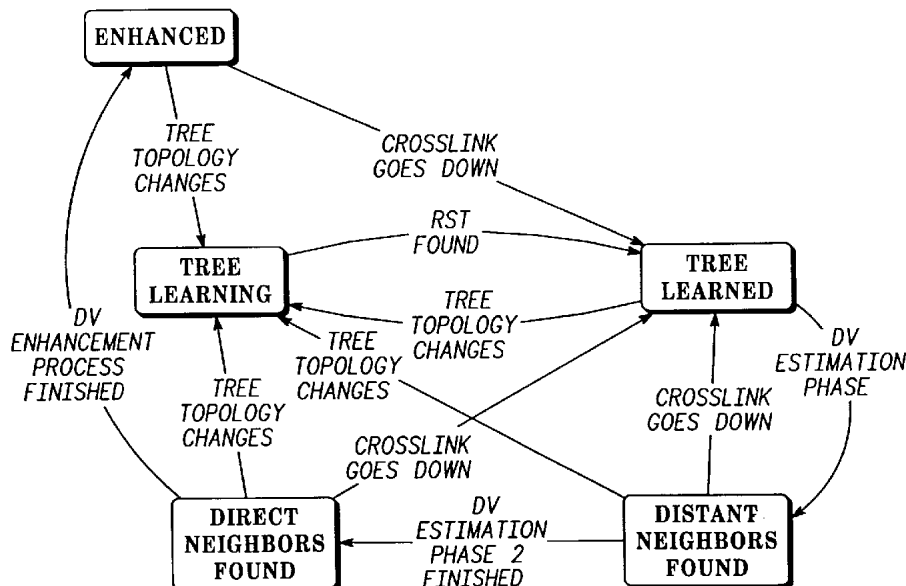
Assistant Examiner—Dmitry Levitan

(74) *Attorney, Agent, or Firm*—Volpe and Koenig, P.C.

(57) **ABSTRACT**

Bridges (10, 12, 14) are used to interconnect local area networks transparently. In the IEEE 802.1D standard for bridges, a spanning tree is built among the bridges for loop-free frame forwarding (FIG. 10). Although this approach is simple, it does not support all-pair shortest paths. A novel bridge protocol is employed that attempts to find and forward frames over alternate paths that are shorter than their corresponding tree paths on the standard spanning tree, and makes use of the standard spanning tree for default forwarding. The protocol, referred to as the Spanning Tree Alternate Routing (STAR) Bridge Protocol, is backward compatible with the IEEE 802.1D standard and has a complexity that is comparable to that of the standard and other existing protocols and further includes an ability to prioritize the forwarding of frames over at least two different paths responsive to a priority value carried by a frame to be forwarded.

33 Claims, 21 Drawing Sheets



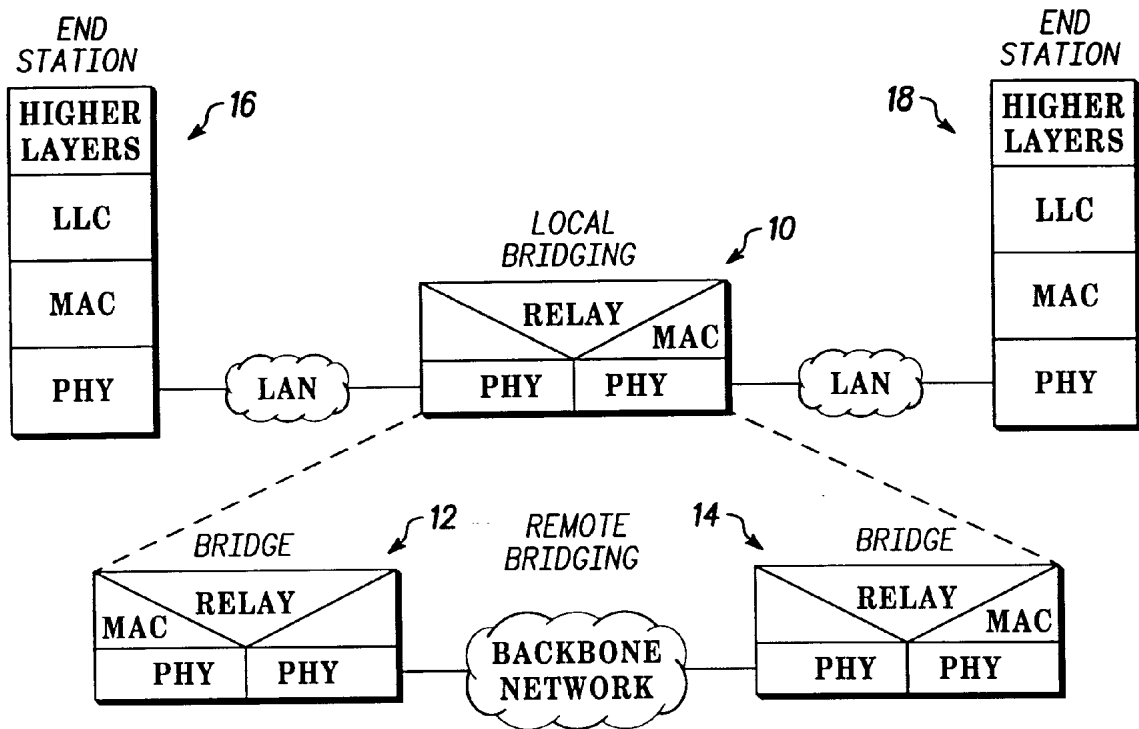


FIG. 1

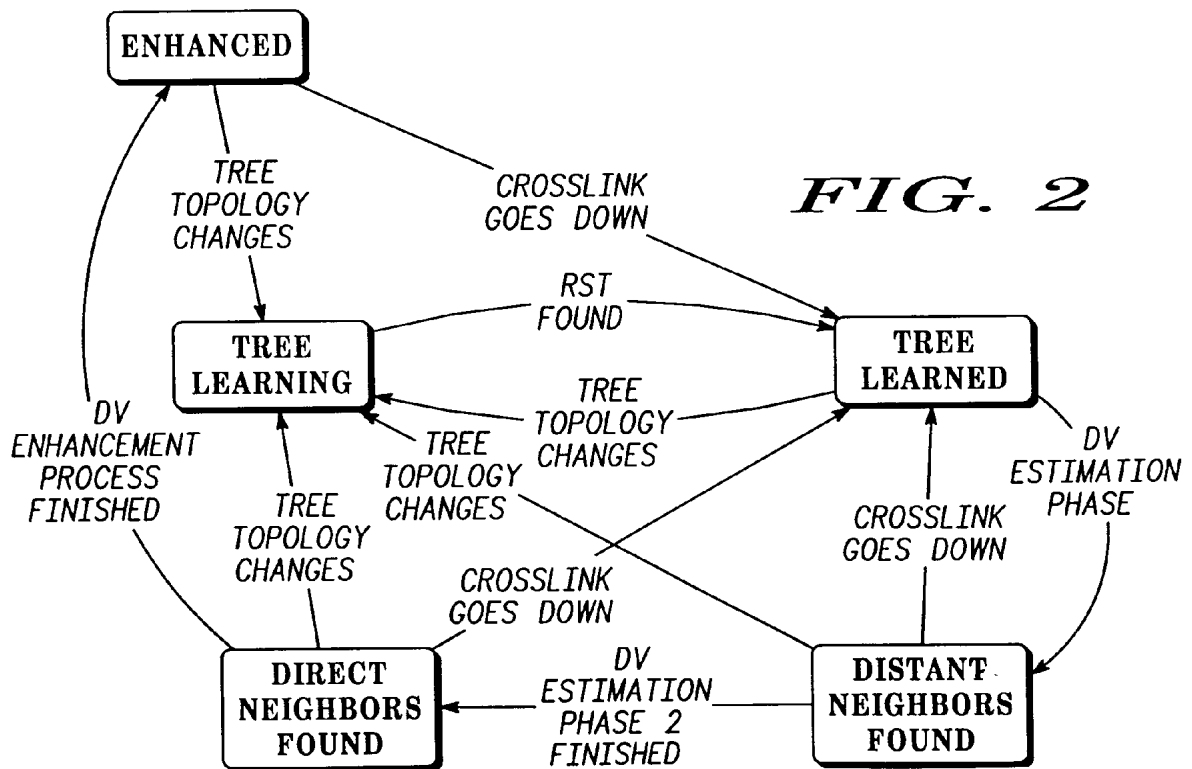


FIG. 2

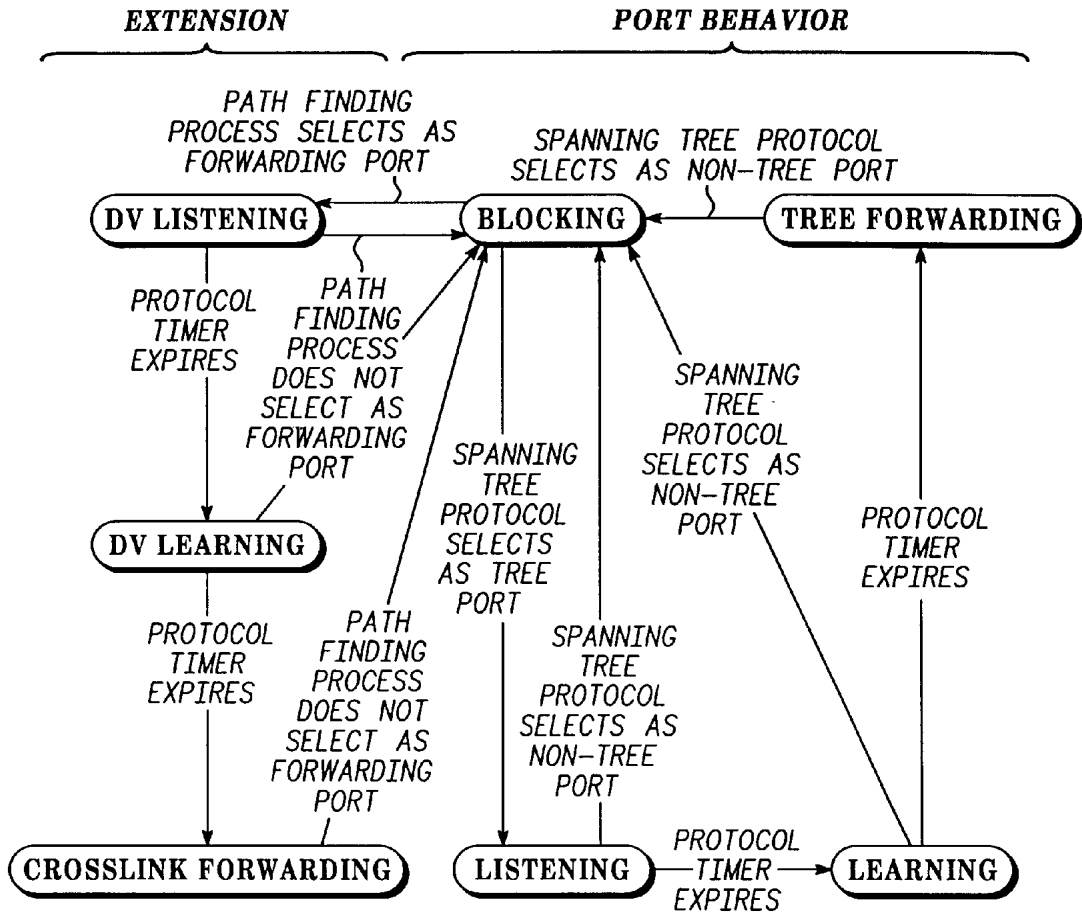


FIG. 3




FIG. 4

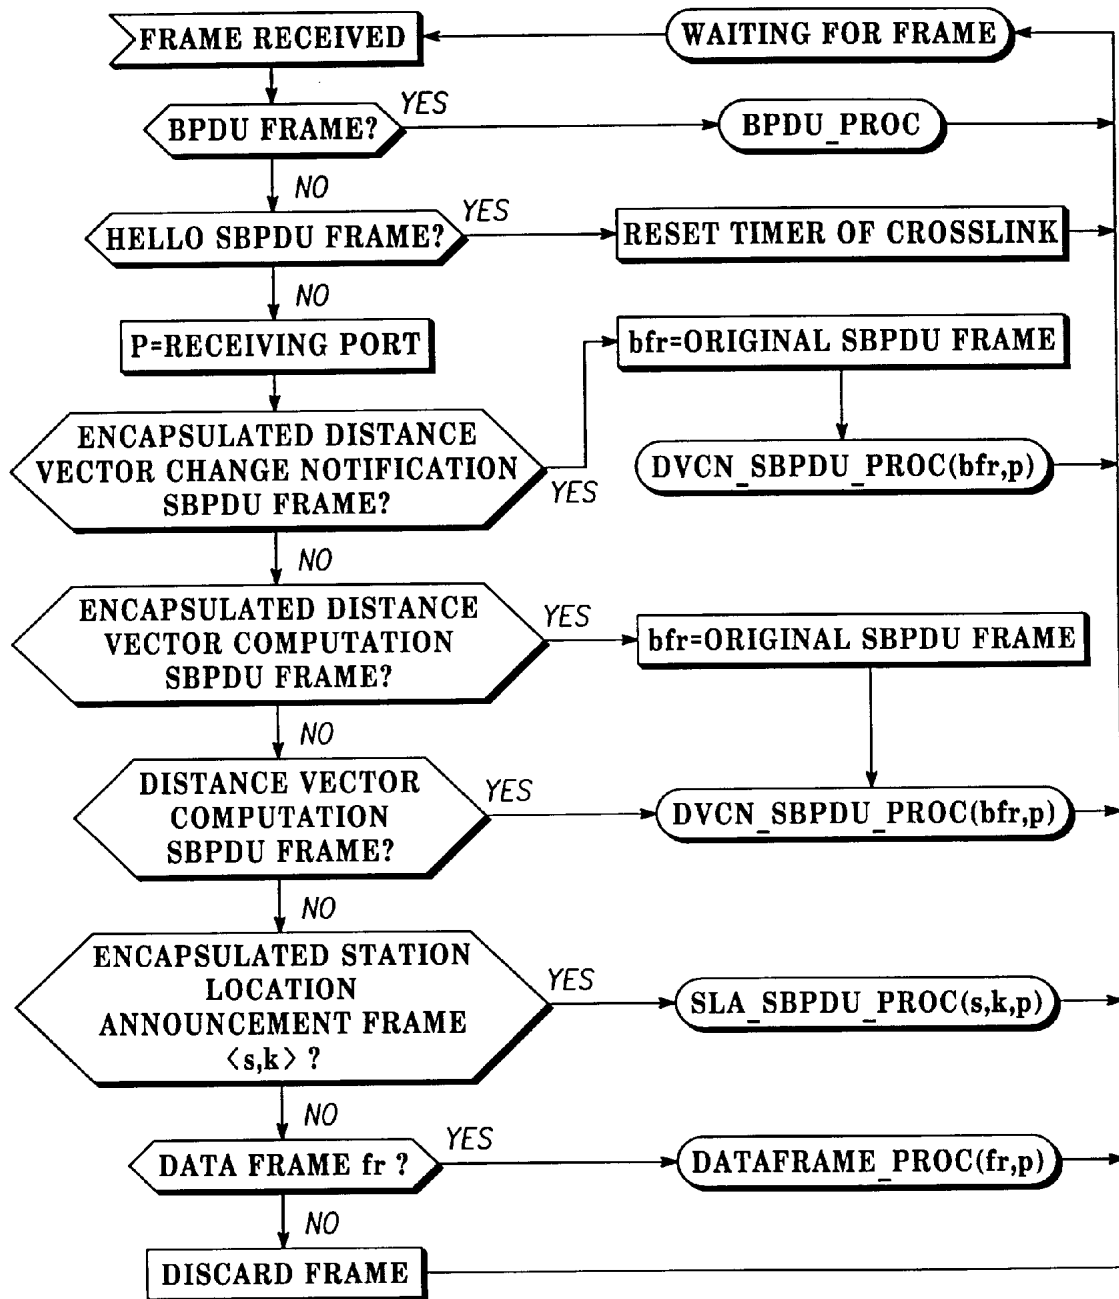


FIG. 5

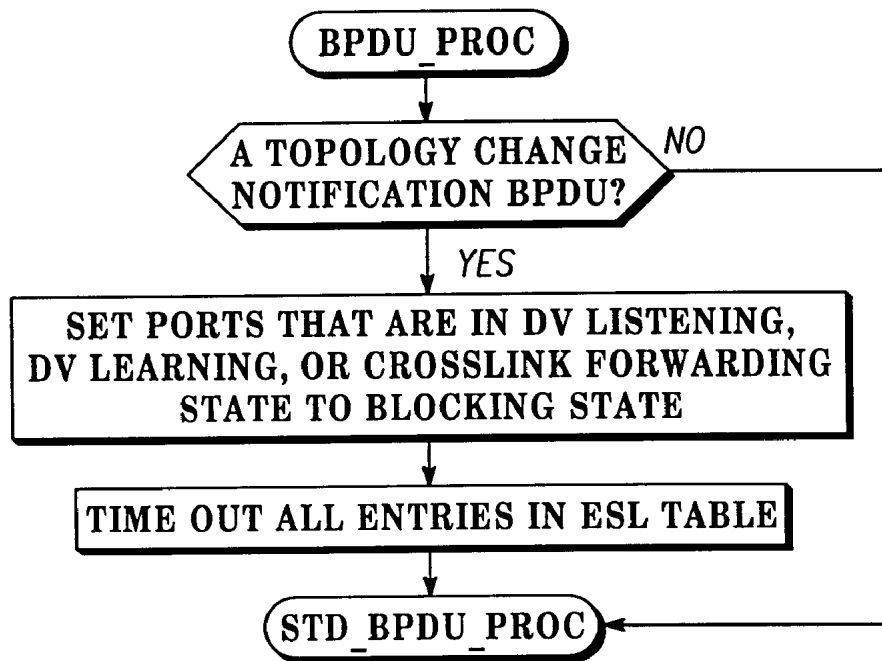


FIG. 6




FIG. 7

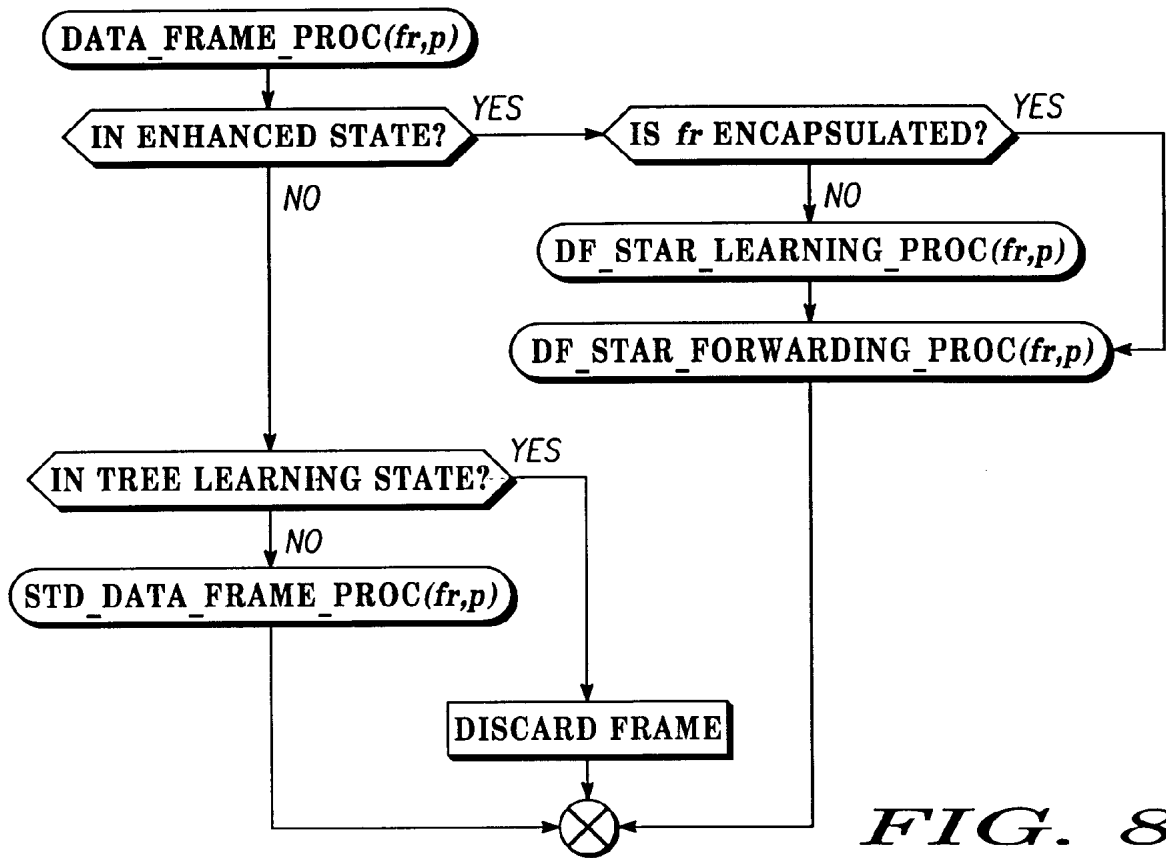


FIG. 8

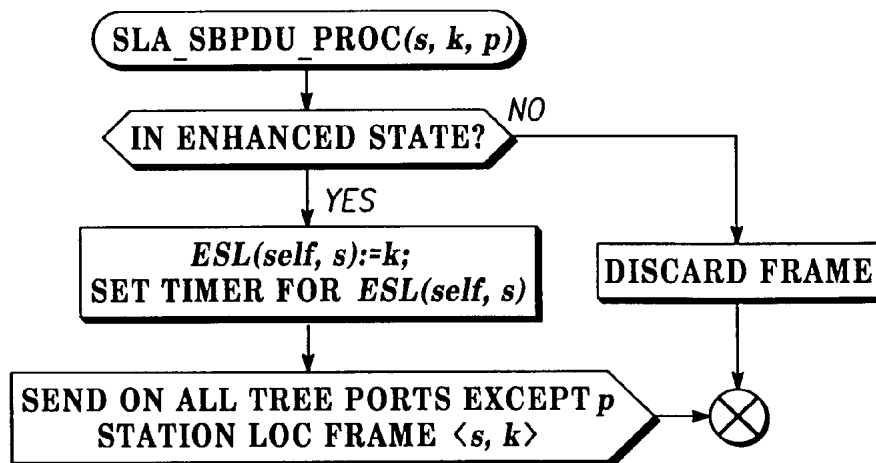

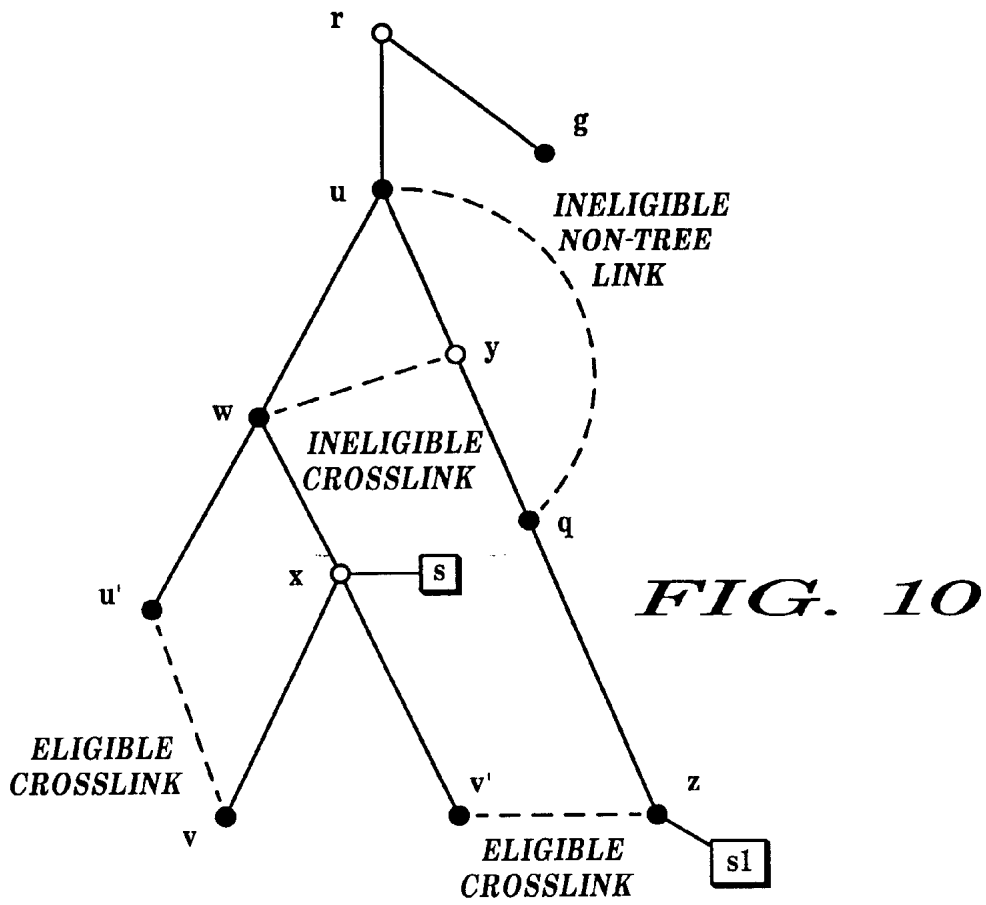


FIG. 9



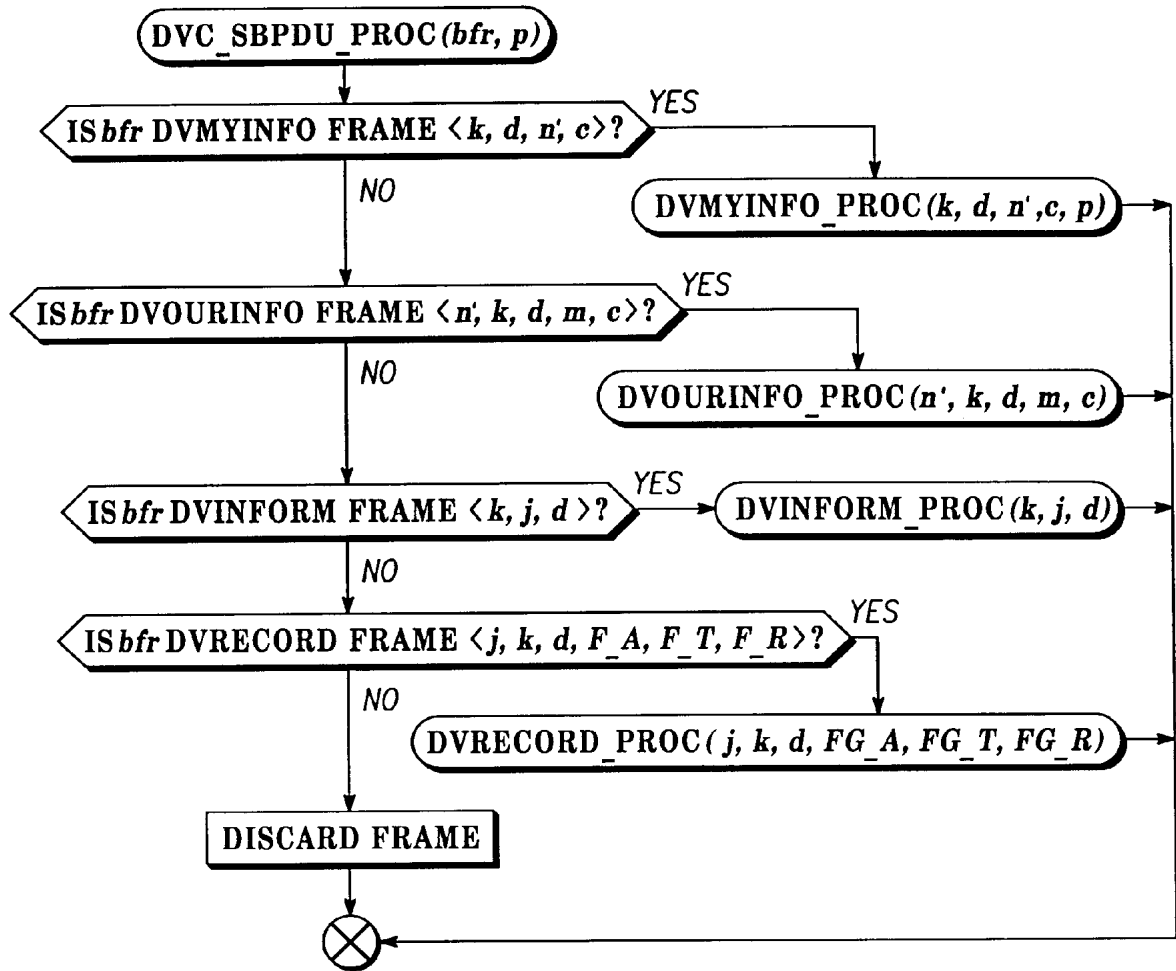


FIG. 12




FIG. 13

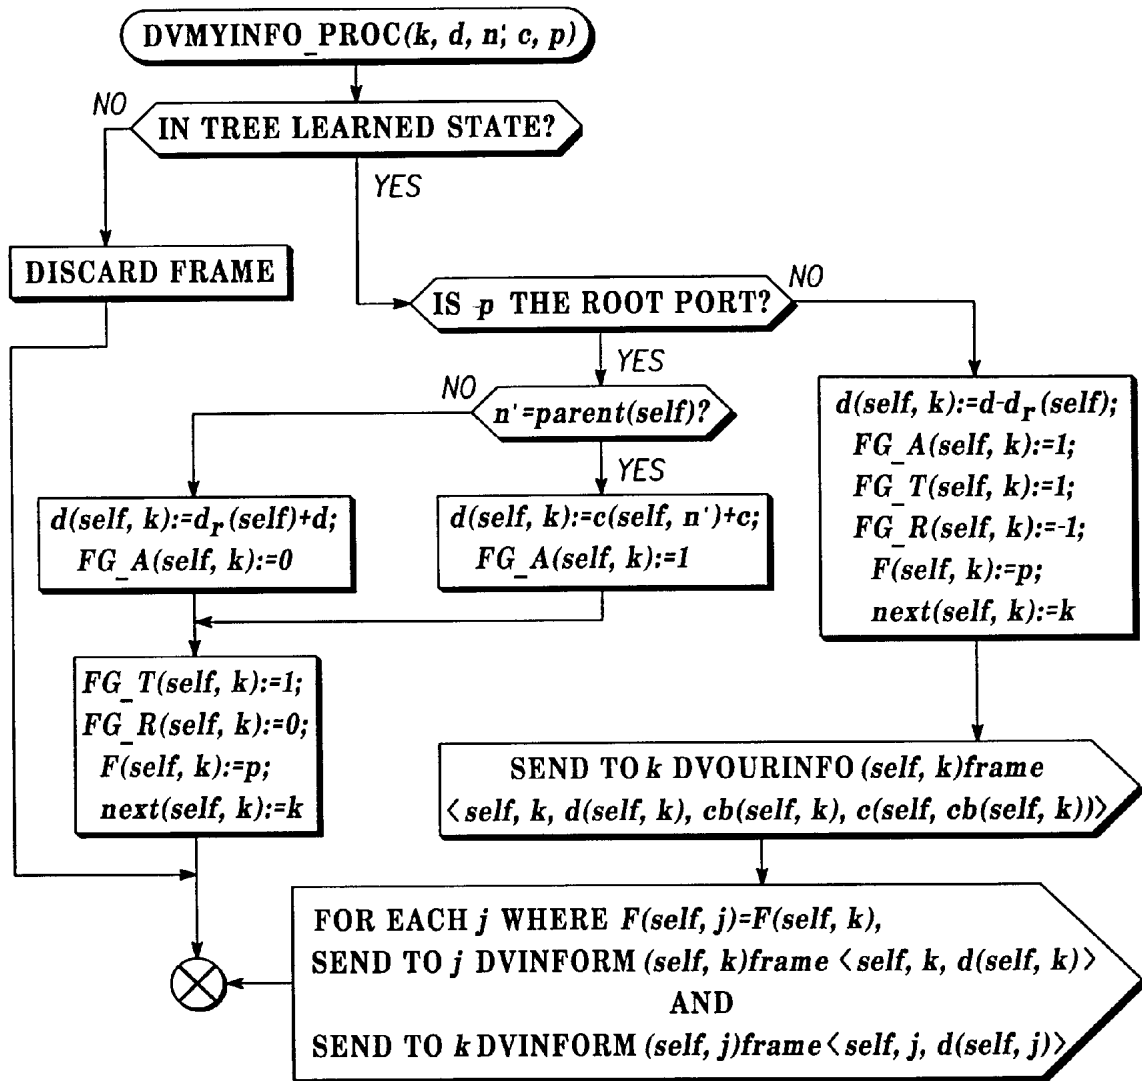


FIG. 14

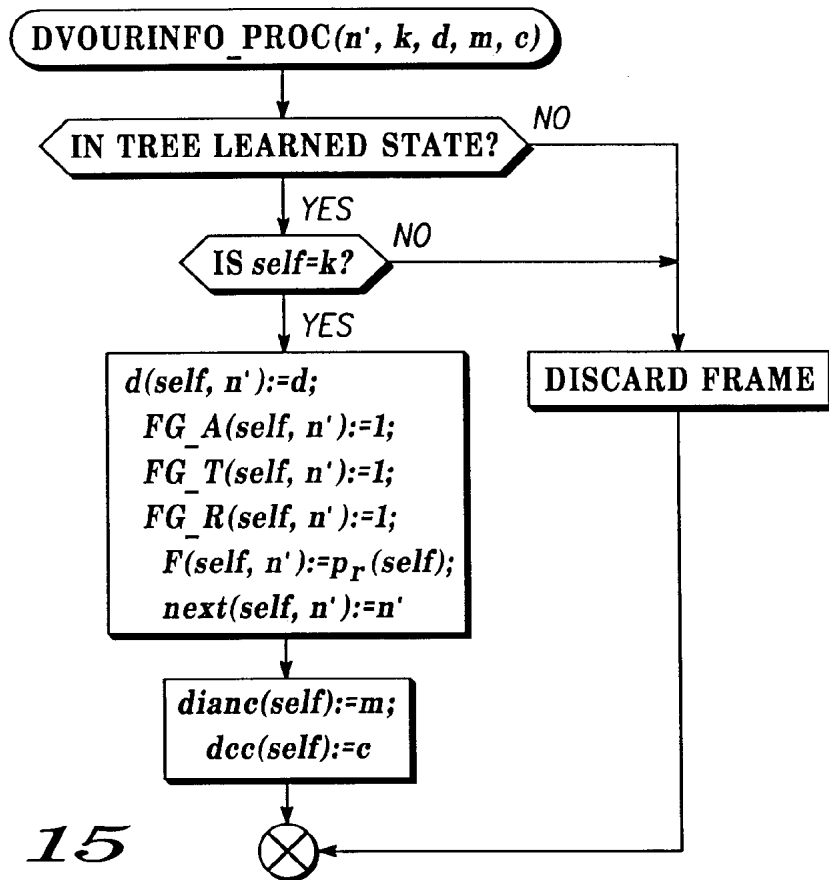


FIG. 15

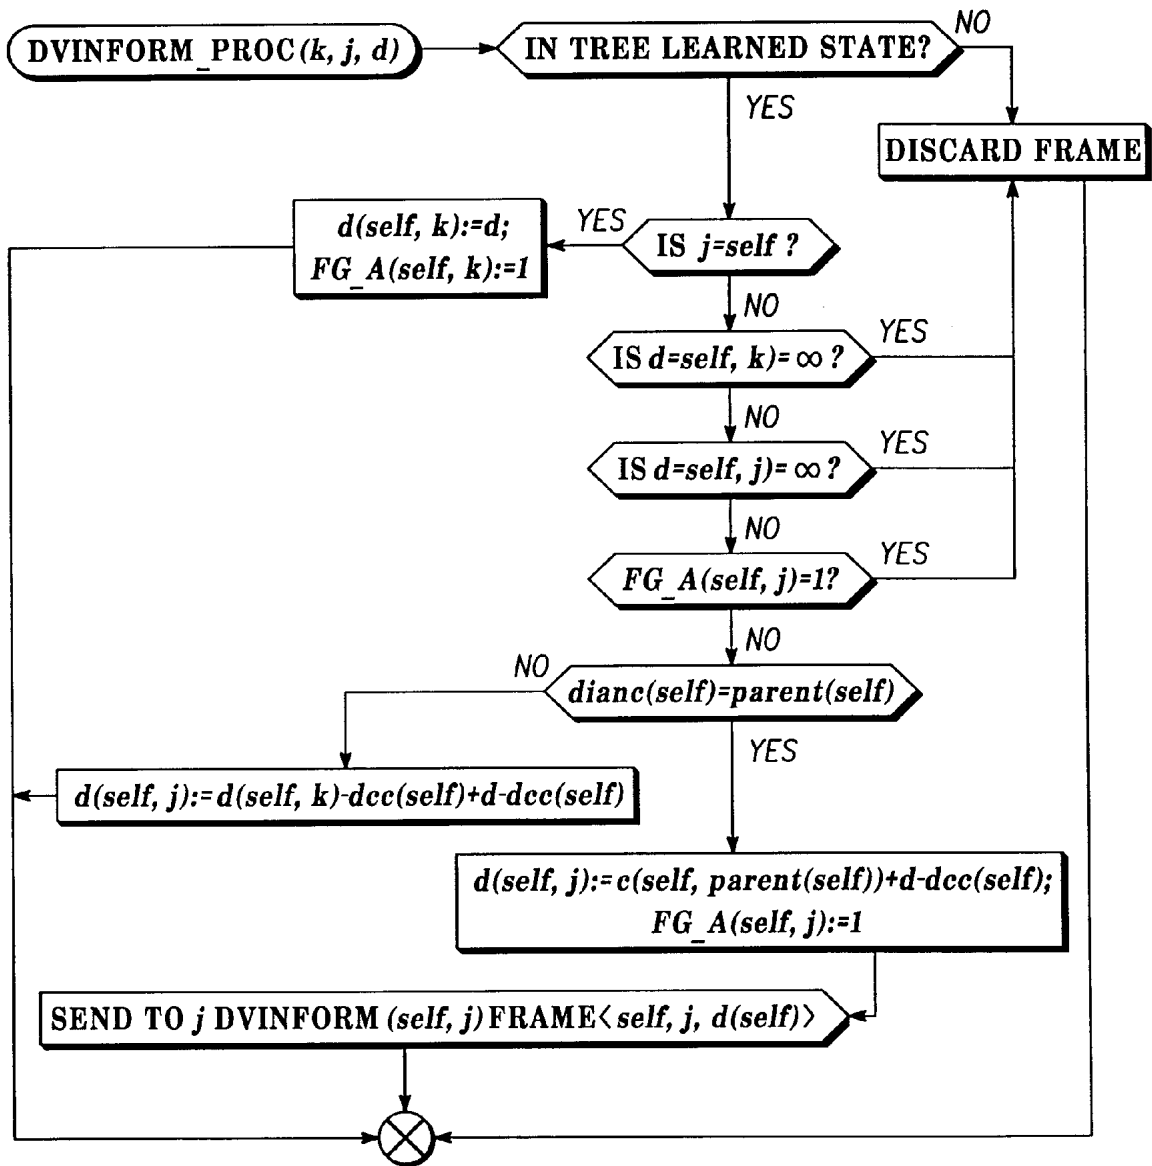


FIG. 16

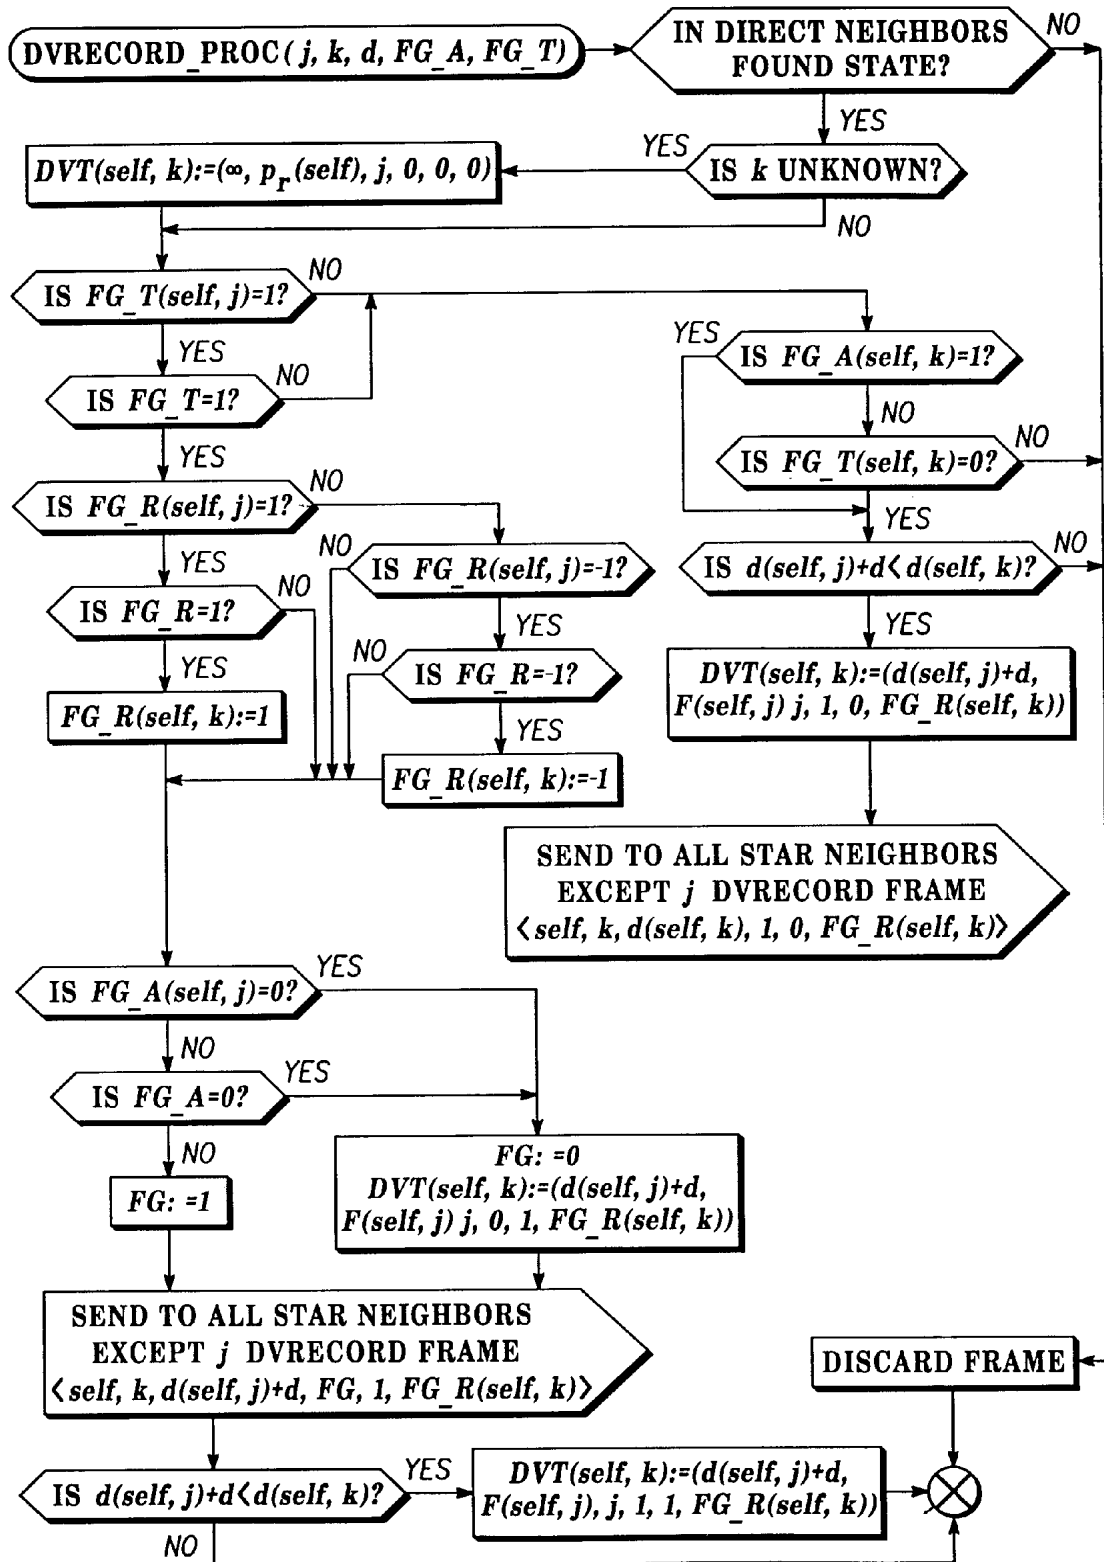


FIG. 17

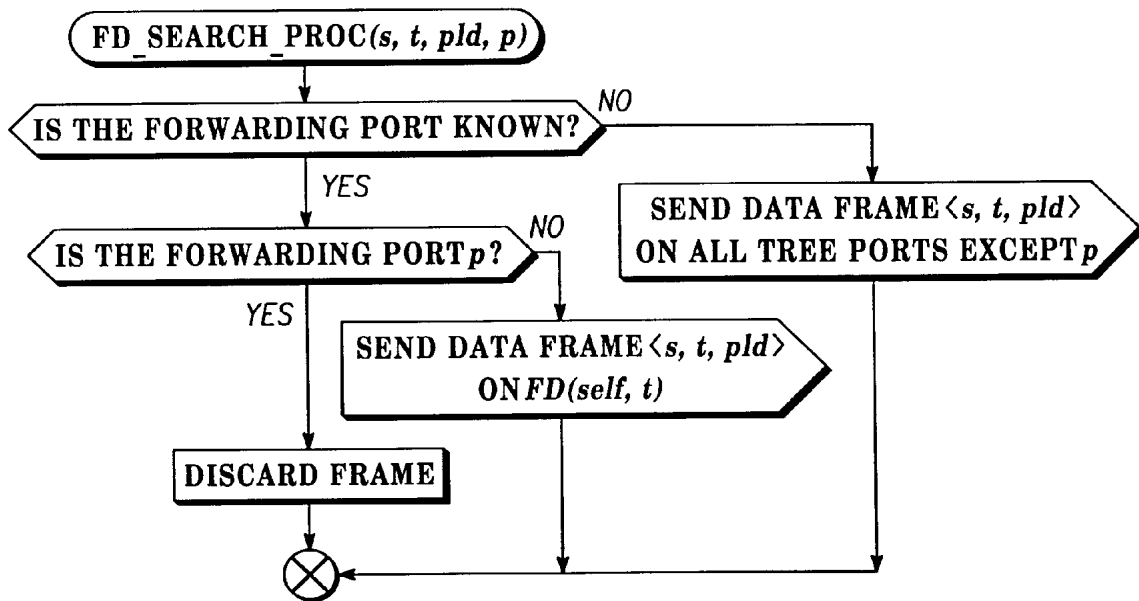
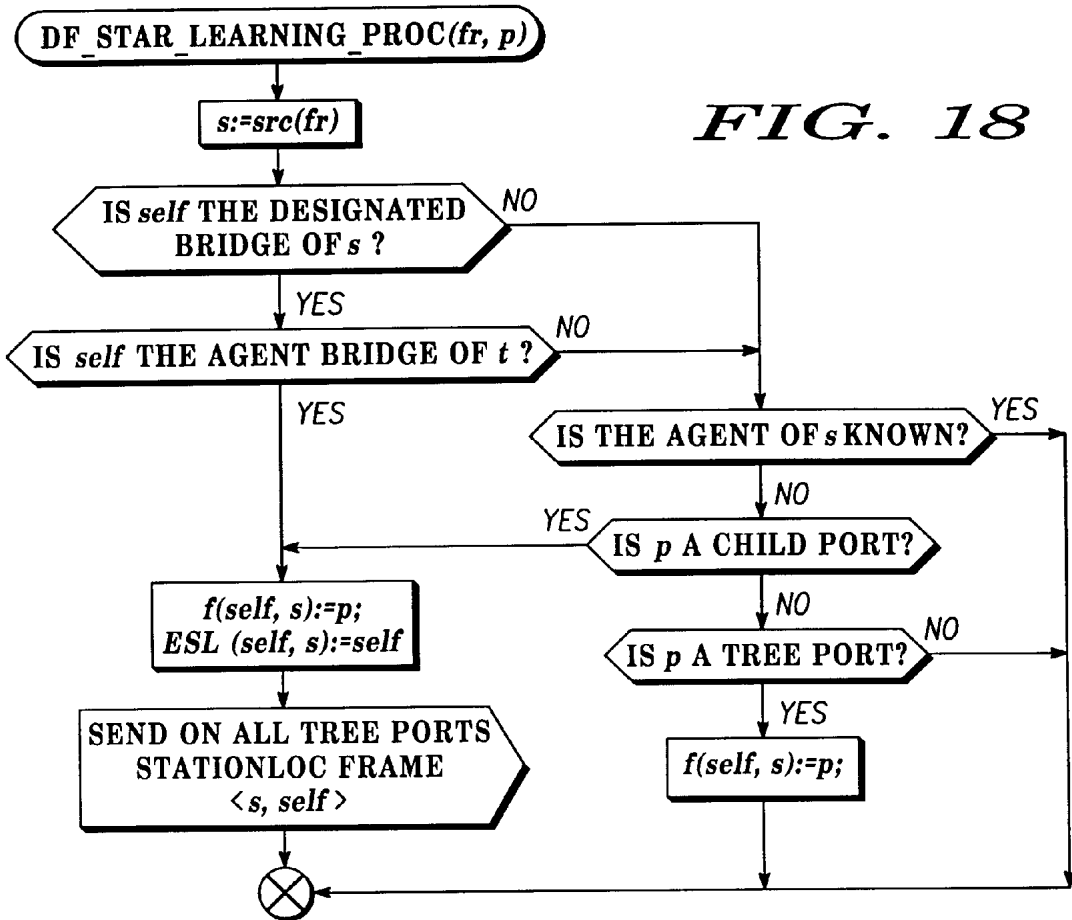


FIG. 19

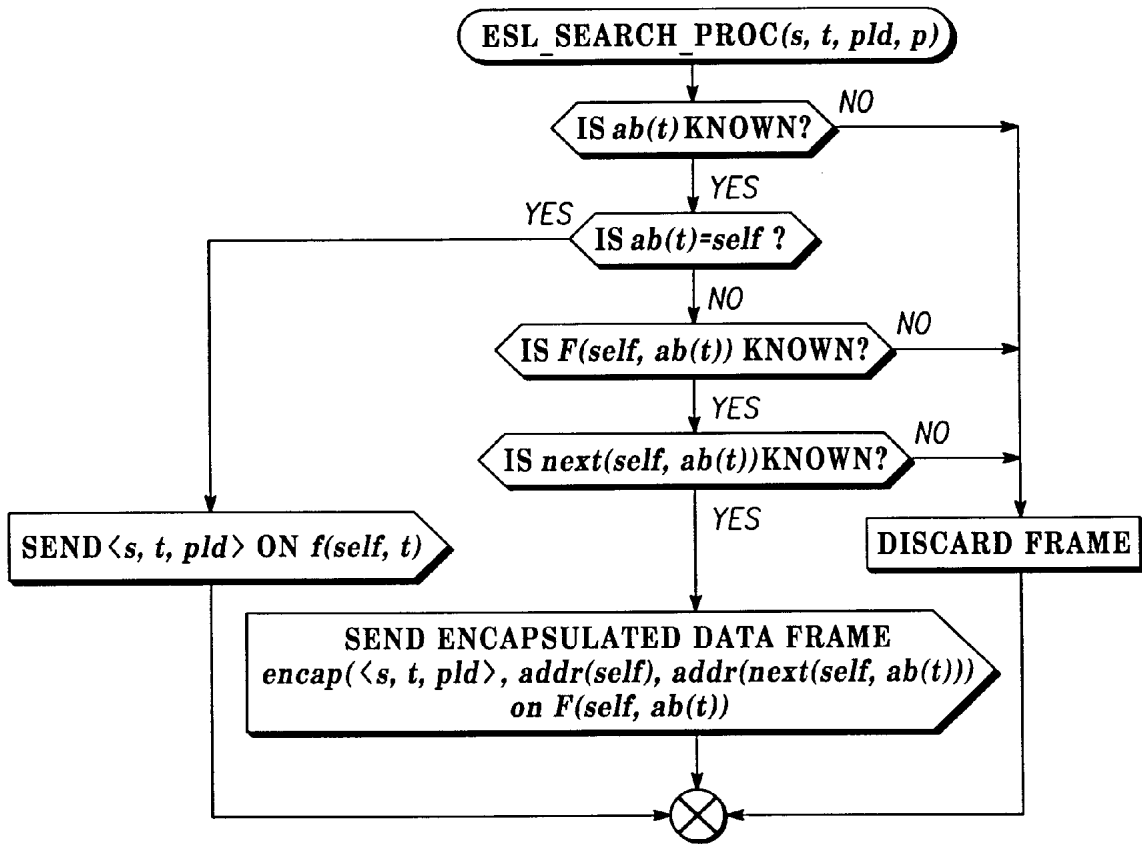


FIG. 20

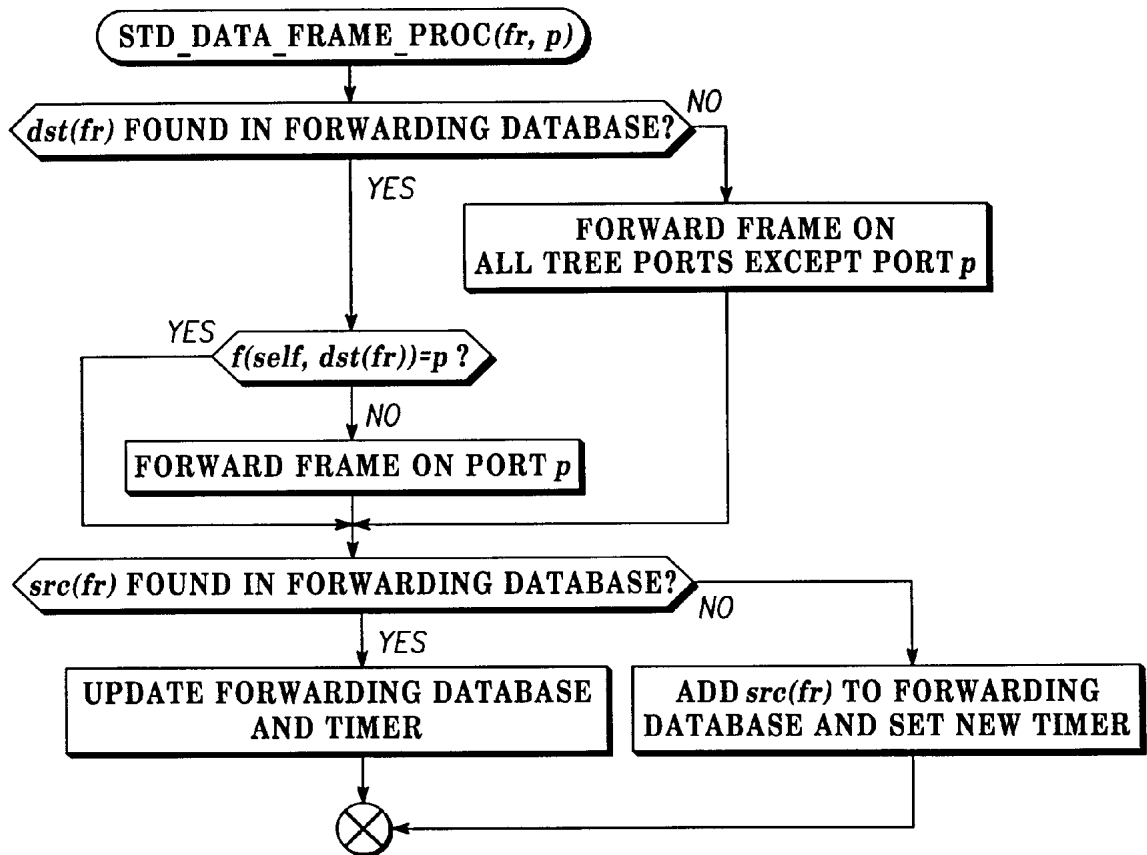


FIG. 21

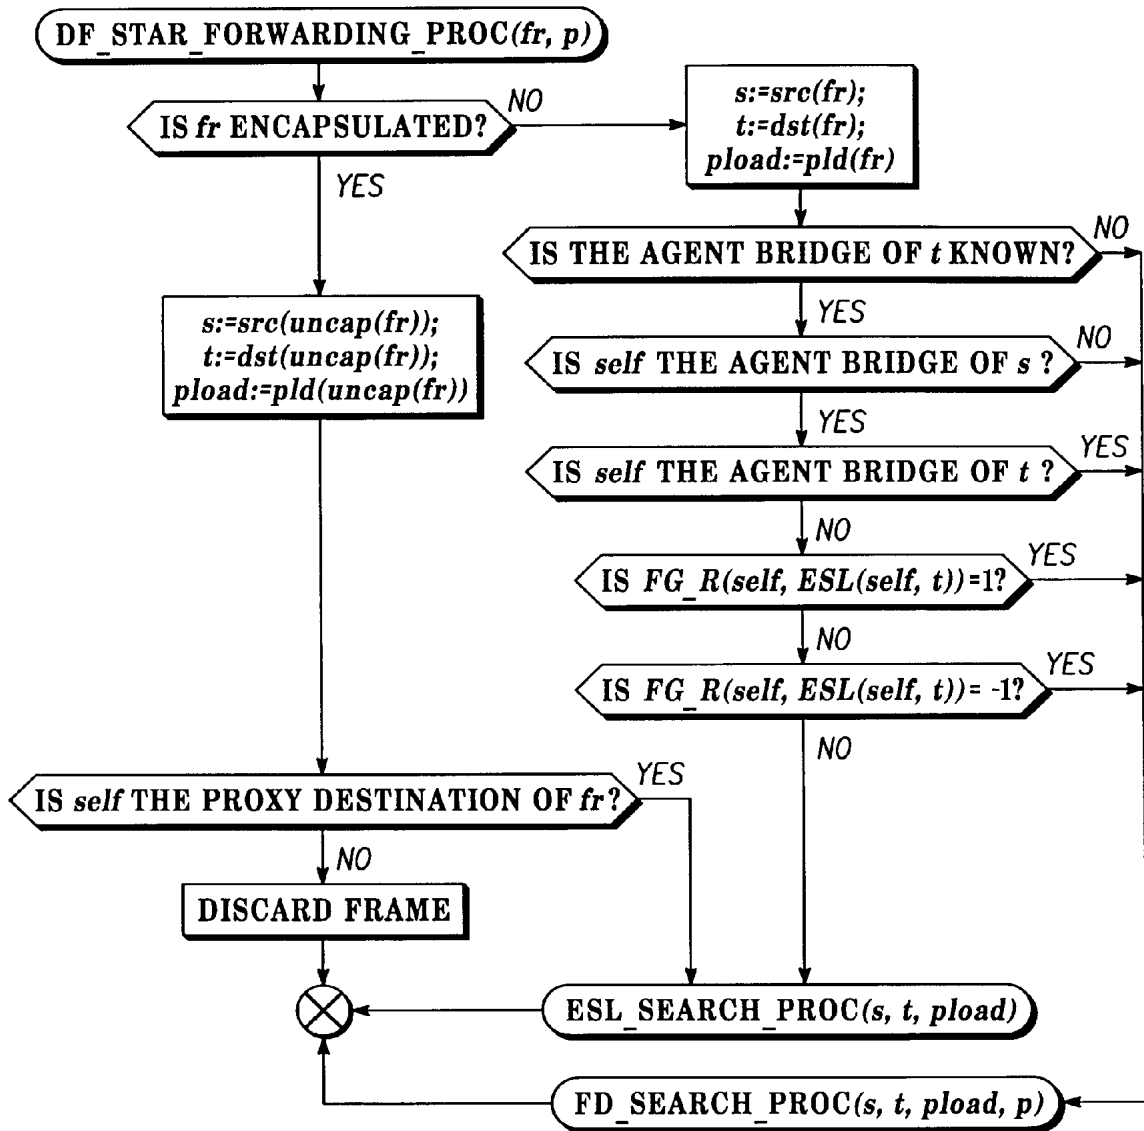


FIG. 22




FIG. 23

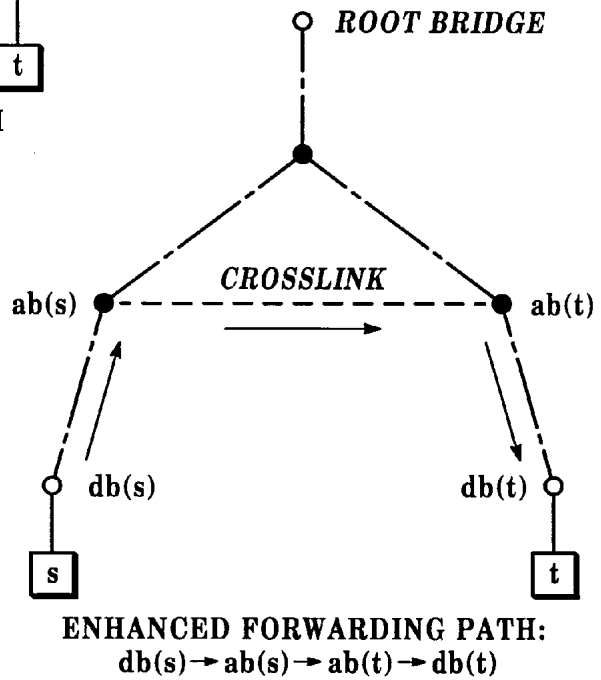


FIG. 24




FIG. 25




FIG. 26




FIG. 27




FIG. 28




FIG. 29




FIG. 30




FIG. 31




FIG. 32




FIG. 33




FIG. 34




FIG. 35




FIG. 36

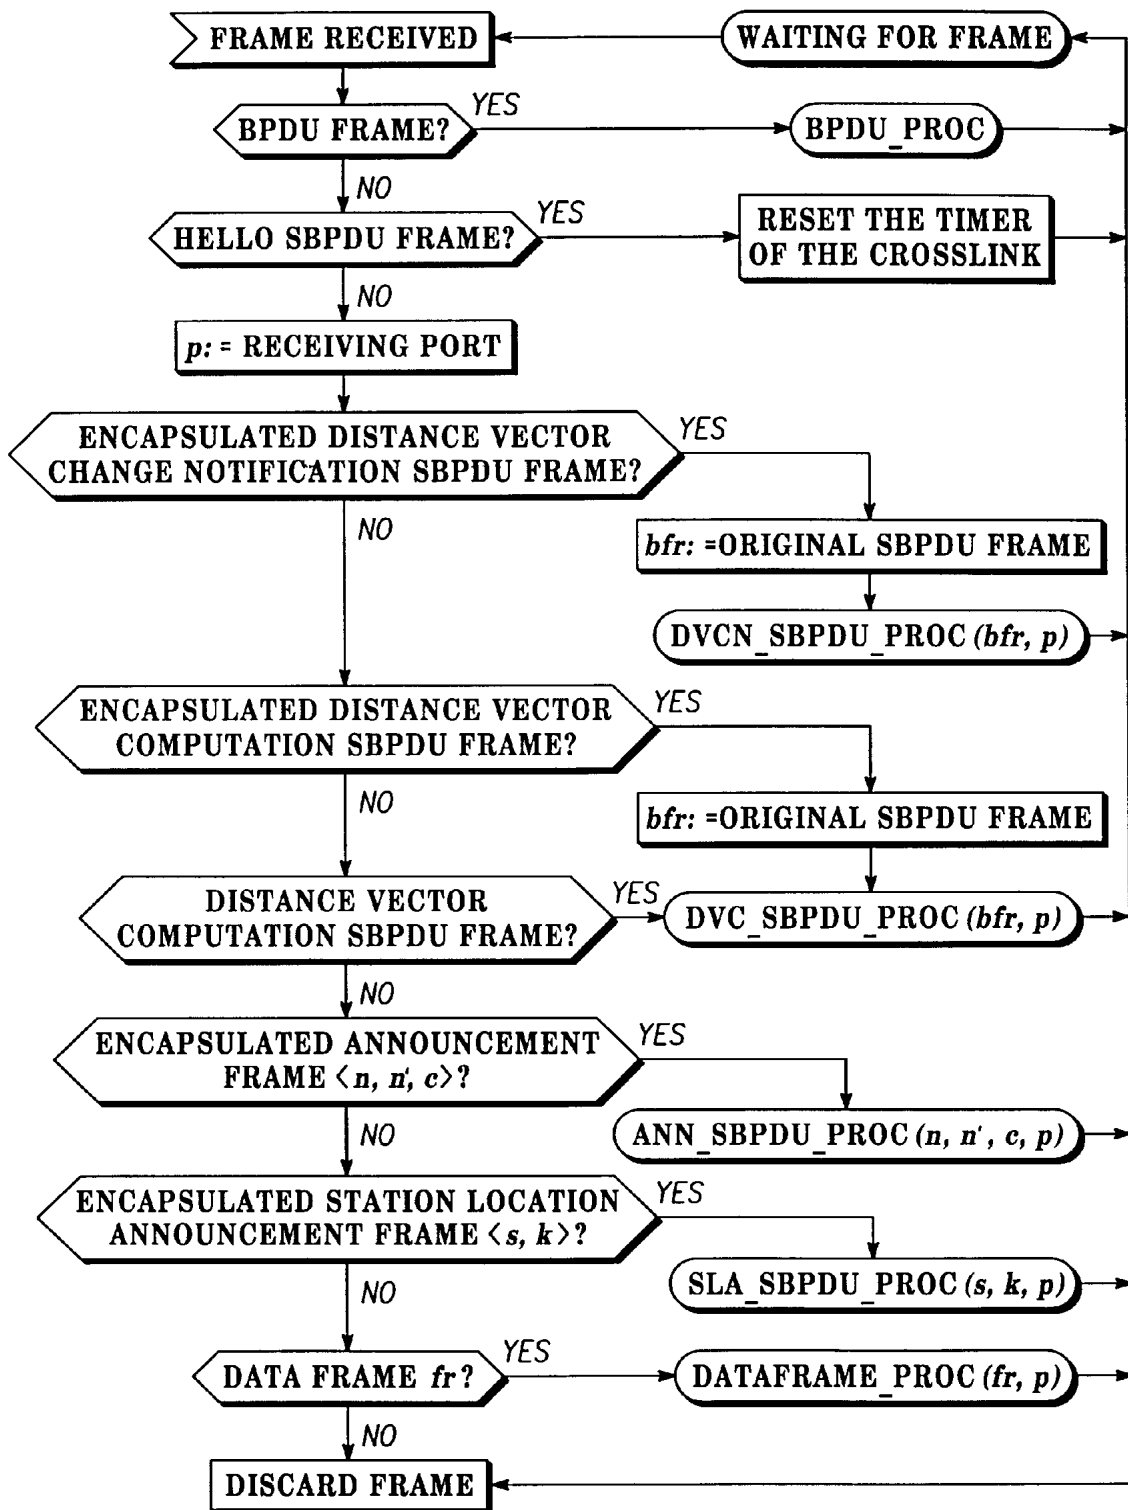


FIG. 37

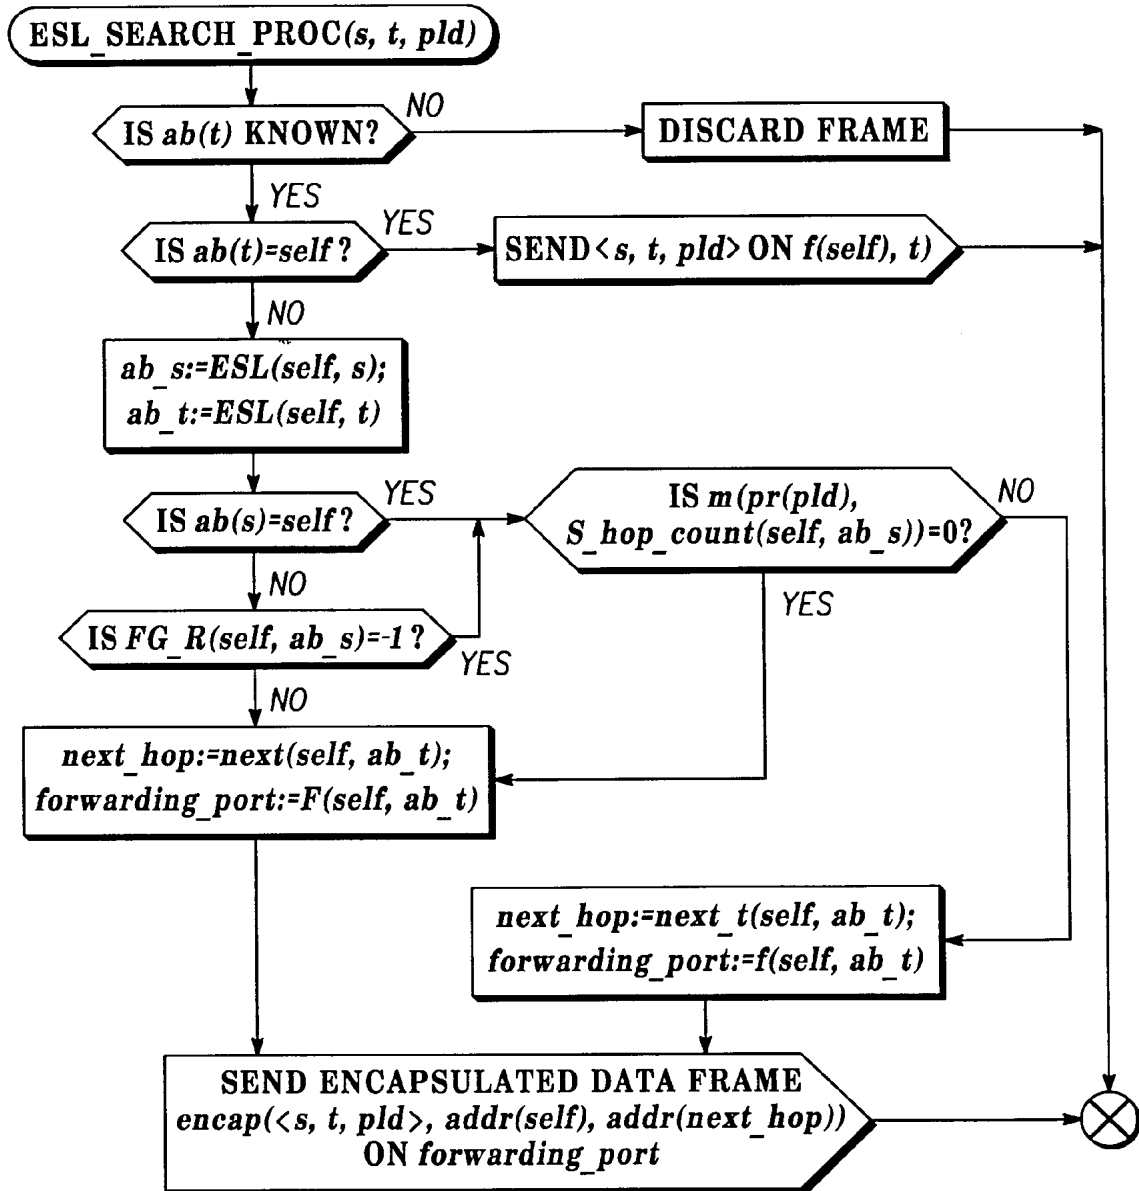


FIG. 38

METHOD AND APPARATUS FOR PRIORITY-BASED LOAD BALANCING FOR USE IN AN EXTENDED LOCAL AREA NETWORK

CROSS REFERENCE TO RELATED APPLICATION

This application claims priority from U.S. Provisional Patent Application No. 60/350,572, filed on Jan. 20, 2002, which is incorporated by reference as if fully set forth.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to employing a bridge protocol for interconnecting two or more local area networks (LANs). More particularly the invention relates to apparatus and method, which is backward compatible with existing 802.1D Spanning Tree Bridge Protocol, for improving routing capability of spanning tree forwarding without a significant increase in complexity and which enables traffic prioritized path differentiation and provides load balancing by supporting alternate paths selected according to one of two or more priority values carried by a frame to be forwarded.

2. Background Information

A Local Area Network (LAN) is used to connect end stations together to provide communications. A single LAN permits a limited number of end stations, a limited size, and a limited amount of offered load. In this respect, LANs cannot grow beyond a certain limit. LANs may be interconnected via internetworking devices such as bridges and routers. These devices have different advantages and disadvantages depending on the internetworking environment. In the early days of internetworking, bridges were popular because they were much cheaper and faster than routers. In addition, bridges were used to support heterogeneous network layer protocols. The primitive computing technology of those days favored off-loading of work to larger servers using protocols that were optimized for LANs.

IEEE 802 Standards Committee has specified two bridge protocols. IEEE 802.1 group has issued the IEEE 802.1D Spanning Tree Bridge Protocol and IEEE 802.5 group has issued the Source Routing Bridge Protocol. Among these two schemes, IEEE 802.1D offers a better solution and has been studied more intensively. This approach is transparent to end stations and requires no modifications to the MAC layer of end stations. All the routing related operations are done in the bridges. Today, the IEEE 802.1D Spanning Tree Bridge Protocol is widely used for interconnecting the family of IEEE 802 standard LANs.

A bridge has several ports connecting to different LANs. A frame sent from one LAN to the other will typically go through one or more ports and bridges. As bridges are capable of filtering frames, they are useful for dealing with unnecessary broadcast traffic. Such a broadcast containment capability renders bridging a simple solution to implementing a virtual LAN. This bridged LAN environment should be transparent and looks like a single LAN to end users. The basic function of bridges is to forward MAC (Medium Access Control) frames from one LAN to another, therein providing an extension to the LAN without requiring any modification to the communications software in the end stations attached to the LANs. Bridges do not modify the content or format of the MAC frames they receive. The operation of bridges should not disorder or duplicate frames. Upper-layer protocol transparency is an advantage

of bridging since bridges can rapidly forward traffic representing any network-layer protocol without having to examine upper-layer information.

The landscape for internetworking has evolved considerably with advances in high-speed network layer routing and data link layer switching technologies. Functionalities at the two layers are increasingly similar. While routers are generally more intelligent than bridges in terms of their dynamic routing capability, they are also more complicated to operate and costly to implement. Bridges have been designed to span a range of routing capabilities from dynamic source routing to static spanning tree forwarding, thereby allowing a trade-off between routing performance and protocol complexity. Although routers are becoming cheaper and faster than they used to be, they remain more complicated than bridges to operate because intermediate hops must still rise above the data link layer in the protocol stack. In spite of the common wisdom that IP has won the network layer, there are still going to be non-IP network layer protocols in the foreseeable future. On the other hand, while bridges are evolving to accommodate more and more network layer functionality, they will always support multiple network layer protocols.

An IP (Internet Protocol) address encodes both a network and a host on that network. Since it does not specify an individual end system, but a physical location in a network, the IP address of a host must change whenever it moves from one network to another. On the other hand, an IEEE 802 MAC address identifies an end system instead of a physical location in a network, and hence is always applicable to a host no matter where it is located in the network. Such portability of addresses is important particularly for mobility and the benefit of plug-and-play. Although new features, are emerging to minimize the need to configure and reconfigure IP addresses, these features can increase the cost and processing overhead of the system. DHCP (Dynamic Host Configuration Protocol), for example, provides a widely deployed framework for host registration and configuration. DHCP, however, was designed only for fixed hosts on physically secure LANs. DHCP is being extended to allow dynamic reconfiguration of a single host triggered by the DHCP server (e.g. a new IP address). Depending on the bandwidth of the network between server and client, the delay in the reconfiguration process can grow exponentially as failed retransmissions trigger exponential backoff.

The IEEE 802.1D specification defines a protocol architecture for MAC bridges and recommends formats for a globally administered set of MAC station addresses across multiple LANs.

FIG. 1 shows a bridge protocol architecture for a connection of two LANs via local **10** or remote **12, 14** bridging. Referring to the OSI (open systems interconnect) reference model, a bridge encompasses the first two layers, namely the Physical Layer (layer **1**) and the Data Link Layer (layer **2**). There are two sublayers in layer **2**: Medium Access Control (MAC) sublayer and Logical Link Control (LLC) sublayer. Bridges operate relay functions on the MAC sublayer and interface with the LLC sublayer above through LLC service access points. By using bridges, a growing LAN can be partitioned into self-contained units for administrative or maintenance reasons, as well as to improve performance via load balancing and fault isolation. Bridges are typically used to interconnect LANs of the same type, such as the family of IEEE 802 LANs. Translation among different link-layer protocols is needed, however, when the interconnected LANs are not homogeneous (e.g., IEEE 802.3 and IEEE 802.5 type LANs), and interoperability is achieved by appropriate frame encapsulation.

A bridge relays individual MAC user data frames between separate MAC protocols of the bridged LAN connected to the ports of the bridge. A MAC entity for each port handles all the media access method dependent functions, i.e., MAC protocol and procedures, as specified in the relevant IEEE 802 standard for that MAC technology. Each bridge port receives and transmits frames to and from the LAN to which it is attached using the services provided by the individual MAC entity associated with that port. Each bridge port also functions as an end station providing MAC service to the LLC layer. All MAC entities communicating across a bridged LAN are uniquely identified by their respective 48-bit MAC addresses. A bridge may use a 48-bit MAC address, or a 16-bit locally administered MAC address. This bridge address must be unique within the extended LAN, and a single unique bridge identifier (ID) is derived from it for the operation of a bridge protocol. Each frame transmitted from a source end station to a destination end station carries the MAC addresses of the end stations respectively in the source and destination address fields of the frame's MAC header. A frame that is to be relayed by every bridge to all its neighboring bridges in an extended LAN contains a bridge group MAC address in the destination address field of the frame's MAC header.

In the IEEE 802.1D standard, a shortest path spanning tree with respect to a predetermined bridge, known as a root bridge, is used to interconnect LANs to form an extended LAN. The spanning tree is determined by bridges in the extended LAN via a distributed spanning tree algorithm. A frame sent from one LAN to another could follow a longer path on the spanning tree than necessary when there exists an alternative shorter path connecting them. Note that non-tree links, which are links that have not been selected by the 802.1D spanning tree algorithm, are not used to share the load of the traffic. The load around the root bridge may be heavy, and throughput is severely limited.

The three basic functions set forth in the present standards of an IEEE 802.1D bridge are:

- 1) frame forwarding—forward a frame received from one port to another port
- 2) learning—"learn" and "remember" which port to forward a frame
- 3) spanning tree algorithm—make sure activated links form no loop, i.e., the bridges and links form a spanning tree

Functions (1) and (2) above are performed with the help of a Forwarding Database, or Filtering Database, (see FIGS. 7-4 of IEEE 802.1D 1998 Edition), within each bridge. Each bridge keeps a Forwarding Database, hereafter denoted FD, that specifies which port to forward a data frame with a particular destination. If there is no such entry in the FD, the bridge forwards the frame through all ports except the port from which the frame originates. Whenever a frame from source *s* arrives from port *p*, the bridge marks in its FD that the forwarding port of *s* is *p*. As the learning process is simple, if there is a loop in the bridged LAN, a frame may be forwarded indefinitely around the loop. To avoid this undesirable feature, function (3) mentioned above is used to make sure the active topology among the bridges is always a tree so that there is a unique path between each pair of bridges. We refer to such a path herein as a tree path.

The spanning tree algorithm builds a unique shortest path tree rooted at the root bridge in a distributed manner. This root bridge is typically selected using bridge identifiers. A path connecting the root bridge and another bridge over the spanning tree is referred to as a root path associated with the

bridge. By exchanging configuration messages, bridges identify the root bridge and select which ports to activate. For each LAN, a single bridge is elected among all bridges connected to the LAN to be the designated bridge, such that it is the bridge that is closest to the root bridge. In order to maintain an up-to-date spanning tree that reflects the underlying topology, the root bridge broadcasts configuration messages periodically over the spanning tree to all other bridges, for example, approximately every four (4) seconds.

The IEEE 802.1D standard defines two types of Bridge Protocol Data Unit (BPDU), namely Configuration BPDU and Topology Change Notification BPDU. Bridges send MAC frames containing Configuration BPDU to each other in order to communicate topology information and compute the spanning tree. Bridges send MAC frames containing Topology Change Notification BPDU up the spanning tree to inform the root bridge of topology changes. Each Configuration BPDU MAC frame includes a MAC header that contains a source MAC address and a destination MAC address. The source MAC address is the MAC address on the port of the bridge originating the Configuration BPDU MAC frame. The destination MAC address field carries the bridge group MAC address so that the Configuration BPDU MAC frame is received by all the bridges in the extended LAN. The information in the Configuration BPDU may be used by a bridge in preparing its own Configuration BPDU MAC frame. Each Configuration BPDU contains a BPDU Header and a set of BPDU Parameters. The BPDU Header consists of a Protocol Identifier field, a Protocol Version Identifier field and a BPDU Type field. The Protocol Identifier takes a specific value that identifies the Spanning Tree Bridge Protocol. The Topology Change Notification BPDU consists merely of a Protocol Identifier field, a Protocol Version Identifier field, and a BPDU Type field with a code reserved for this type. When a bridge detects a change in the active topology of the spanning tree, it sends a Topology Change Notification BPDU to the root bridge. The root bridge will then broadcast it to all bridges in the extended LAN. The encoding for the fields in the Configuration BPDU and the Topology Change Notification BPDU can be found in the IEEE standards.

In order to balance traffic load, extensions to the IEEE 802.1D Spanning Tree Bridge Protocol have been proposed to allow non-tree links to be used for frame forwarding under appropriate conditions. These extensions consider alternate paths that traverse at least one non-tree link.

To provide quality of service (QoS) in LANs, the IEEE 802.1p standard specifies priority queuing in bridges. This mechanism may be used to improve queuing delay experienced by some frames over that experienced by other frames. If the bulk of the delay on an end-to-end forwarding path is due to queuing delay in the bridges traversed by the path, then providing differentiated bridge queuing delay will be an effective means to support QoS in an extended LAN. On the other hand, if the bulk of the delay is due to cumulative transit delay across one or more LAN segments, then providing path differentiation for frames of different priorities may be a more effective way to enhance QoS since it depends on the number and performance characteristics of the LAN segments traversed.

A need remains for supporting multi-level traffic prioritization and improving load balancing in an extended LAN, while maintaining backward compatibility with the IEEE 802.1D standard.

A prior system, referred to as a Spanning Tree Alternate Route (STAR) Bridge Protocol was previously disclosed by the present inventors and is described in U.S. patent appli-

cation Ser. No. 09/977,115. The prior invention is a QoS-based bridge protocol that is backward compatible with the IEEE 802.1D standard, such that existing bridges need not be modified and new bridges operate seamlessly with the existing standard. This prior system attempts to find and forward frames over shorter alternate paths relative to tree paths on the standard spanning tree, and makes use of the standard spanning tree for default forwarding as well as multicast and broadcast operations. Shorter alternate paths are determined on a best effort basis, and the shortest alternate path identified by the STAR Bridge Protocol is referred to as an enhanced forwarding path. The enhanced forwarding path is not necessarily a shortest path, but is provably shorter than its corresponding tree path. In one embodiment of our prior system, all frames sent from a source bridge to a destination bridge are forwarded over either a standard tree path or an enhanced forwarding path, but not both. In another embodiment of the prior system, all frames sent from a source bridge to a destination bridge are forwarded over standard tree paths by default, while frames of a predetermined class are forwarded over enhanced forwarding paths if these alternate paths are available. In the event that an enhanced forwarding path leads to a destination bridge in a direction that is opposite to the normal tree path forwarding direction, a data frame is encapsulated to avoid unintended frame dropping. To enable the discovery of enhanced forwarding paths, the STAR bridge protocol defines three additional processes that are different from the three existing bridge processes set forth in the IEEE 802.1D standard. STAR bridges can execute both the standard and the new processes. One new bridge process allows a STAR bridge to find and estimate the distance of a path from itself to another STAR bridge. The other two new bridge processes enable extended learning and enhanced forwarding by the STAR bridges. The new processes operate based on control information carried in newly defined bridge protocol data units that are exchanged among STAR bridges. The STAR bridge protocol may be enabled to use the standard spanning tree for low priority traffic while forwarding high priority frames on enhanced forwarding paths whenever they are available. As IEEE 802.1p standard provides eight (8) different priorities, it is desirable to provide more than two forwarding paths to support different QoS levels when these paths are available and can be found. No prior system, except for the STAR Bridge Protocol, is known to support path differentiation in an extended LAN. However, the STAR Bridge Protocol supports at most two different forwarding paths between each pair of STAR bridges.

One prior art technique teaches a method of distributed load sharing (DLS) to allow non-tree links satisfying certain topological constraints to be selected for frame forwarding. The DLS method does not guarantee that a forwarding path is no worse than its corresponding tree path for any additive metric considered. Another prior art technique teaches an extended and simplified version of the original DLS method known as Generalized DLS (GDLS), wherein any non-tree link may be used. In the GDLS method, a non-tree link is not used if its "speed" is smaller than that of the corresponding tree path. Neither of the DLS method nor the GDLS method supports traffic prioritization.

In the Source Routing Bridge Protocol which is an IEEE standard, information on a forwarding path is included in the header of each frame and each bridge simply reads the header and forwards the frame accordingly. The responsibility of finding and selecting a suitable forwarding path is on the end stations, not the bridges. When an end station wants to send a frame, it first broadcasts a path discovery

frame. The path is recorded in the frame as it travels across the bridged LAN. When the frame reaches the destination end station, the desired path is found and the frame, which has a complete path set forth therein, is sent back to the source end station. As this mechanism of finding forwarding paths is basically a flooding mechanism, it may generate a large number of frames for each forwarding path discovered. With this path forwarding mechanism of the Source Routing Bridge Protocol, a frame can always travel through a shortest path. However, transparency of end stations is sacrificed in order to achieve all pair shortest paths and load sharing.

Another prior art technique teaches a bridge architecture that has IP routing features. The proposed bridges, referred to as SmartBridges, exchange topology information to obtain a complete topology of bridges and LAN segments in the extended LAN. Once the complete topology is synchronized, the shortest path to every LAN can be found. This method employs a mechanism to determine which LAN each end station is connected to, and detects the movement of each end station from one LAN to another. In accordance with this technique, frames that are forwarded between end stations of known locations follow a shortest possible path in the extended LAN. This method does not support traffic prioritization, and is not backward compatible with the IEEE 802.1D standard.

Still another method dynamically creates a shortest path tree rooted at a given bridge using the IEEE 802.1D standard spanning tree as a default. In accordance with this method, which is referred to as Source Dependent Spanning Tree (SDST), some non-tree links are activated and some tree links are disabled on demand according to observed packet transit delay over those links. Information kept in the SDST bridges grows quadratically with the number of ports in the bridges. This method does not support traffic prioritization, and is not backward compatible with the IEEE 802.1 standard.

In still another method, referred to as Optimal-Suboptimal Routing (OSR), alternate paths that are shorter than their corresponding tree paths can be identified in a bridged LAN, provided all the bridges in the bridged LAN implement and execute the method. In each of these bridges, the cost to each known end station is kept for each port of the bridge. OSR is similar to a distance vector routing method. Topology information is kept by every bridge for every port. In order to be backward compatible with the IEEE 802.1D standard, OSR bridges are required to encapsulate original data frames in a new special frame format that is recognized only by OSR bridges. Without such a tunneling mechanism, the frames can interfere with the normal operation of the bridges in accordance with the IEEE 802.1D standard. Due to data frame encapsulation, a path found by the method may be longer than its corresponding tree path when there are bridges that do not execute the method. Also, this method does not support traffic prioritization.

Another method maintains distance vectors in bridges showing the length and forwarding direction of each shortest path to a particular LAN, not an end station. Mapping tables are used to map end stations to LANs. When a frame is received by a bridge, the bridge maps the destination end station to the destination LAN to which the destination end station is attached, and finds the forwarding port from the distance vector. Mapping tables are exchanged by means of flooding, which results in inefficient utilization of network resources. This method does not support traffic prioritization, and is not backward compatible with the IEEE 802.1D standard.

A need remains for a method of multipriority path differentiation and load balancing in an extended LAN, wherein

the method can support more than two priorities and spreads traffic of different priorities over multiple alternate forwarding paths when alternate paths that are shorter than their corresponding tree paths can be found.

SUMMARY OF THE INVENTION

The present invention is a novel protocol which is able to forward frames on two or more alternate forwarding paths (including a tree path) according to frame priorities, provided at least one alternative forwarding path that is shorter than the tree path can be found. If no such alternative forwarding path can be found, the tree path is used as a default forwarding path. It is possible to have alternate forwarding paths from a source end station to a destination end station when agent bridges of the end stations exist and are located on different branches of the spanning tree wherein an agent bridge of an end station is the STAR bridge that is closest to the end station along the standard tree forwarding path to the root bridge on the spanning tree. In this case, a frame forwarded from the designated bridge of the source end station must go upstream on the tree to the nearest common ancestor of the designated bridges of the source and destination end stations. A key feature of the present invention is that the higher the priority of a frame, the smaller the number of hops over the upstream tree path, starting from the designated bridge of the source end station, the frame has to make before it is allowed to be forwarded over an available shorter alternate forwarding path. In addition, for a frame of a given priority, it should be more likely to be forwarded over an available shorter alternate forwarding path the more hops it has made. The invention ensures that the forwarding path of a higher priority frame between a pair of end stations is shorter than or the same as the forwarding path of a lower priority frame between the same pair of end stations.

The present invention is an extension of the STAR Bridge Protocol, which as described herein in detail, attempts to find and forward frames over enhanced forwarding paths, and makes use of the standard spanning tree for default forwarding. In the STAR Bridge Protocol, a frame that is sent from the agent bridge of a source end station to the agent bridge of a destination end station is forwarded over a tree path if an enhanced forwarding path is unavailable. Otherwise, in one embodiment, the frame is forwarded over an enhanced forwarding path leading from the agent bridge of the source end station, and in another embodiment, the frame is forwarded over the said enhanced forwarding path provided the frame belongs to a selected class. Should the frame be forwarded over a tree path, intermediate STAR bridges along the tree path are not able to forward the frame over any of the enhanced forwarding paths leading from themselves toward the destination bridge. In accordance with the present invention, each frame that is sent from the agent bridge of a source end station to the agent bridge of a destination end station may go over a tree path, an enhanced forwarding path leading from the agent bridge of the source end station, and a hybrid forwarding path, which is a concatenation of a partial tree path and an enhanced forwarding path leading from an intermediate STAR bridge. The hybrid forwarding path is shorter than its corresponding tree path. The bridge processes and storage requirements specified in the STAR Bridge Protocol are generally applicable to the present invention. However, the invention extends the path finding process and the STAR forwarding process, but not the STAR learning process in the STAR Bridge Protocol.

The present invention extends the path finding process in the STAR Bridge Protocol. A new SBPDU frame known as

an Announcement frame, is sent once by every STAR bridge in the beginning of the path finding process. This frame is broadcast over the spanning tree. When a first STAR bridge receives an Announcement frame originated from a second STAR bridge, that first STAR bridge knows the tree forwarding port, next hop STAR bridge, and the number of intermediate STAR bridges on the tree path leading to that second STAR bridge and keeps this information in the Bridge Forwarding (BF) Table. In accordance with the present invention, this information is used in the STAR forwarding process.

The present invention modifies the STAR forwarding process in the STAR Bridge Protocol so that more than two alternate paths may be used for forwarding frames. In the original STAR Bridge Protocol, only agent bridges may decide which of a tree path and an enhanced forwarding path is used to forward a data frame. In this invention, some intermediate STAR bridges are required to make similar forwarding decisions. When a STAR bridge receives a frame from a child port and discovers that the frame has been traversing a tree path, it decides whether to forward the frame upstream to a tree neighbor or forward the frame over an enhanced forwarding path leading from said STAR bridge. It determines the forwarding port by the priority of the frame and the length of the tree path the frame has traversed. If the STAR bridge sends the frame to an ancestor STAR bridge, that ancestor STAR bridge must decide on its own whether to forward the frame to an enhanced forwarding path leading from itself to the agent bridge of the destination end station. Subsequently, each STAR bridge that receives the frame must also decide on its own whether to forward the frame to an enhanced forwarding path leading from itself to the agent bridge of the destination end station, provided that the frame is not already being forwarded on an enhanced forwarding path leading from a different STAR bridge. If a frame reaches the agent bridge of the destination end station without a detour over an enhanced forwarding path from any intermediate STAR bridge, the frame will end up being forwarded over a tree path.

The present invention supports multi-level traffic prioritization and load balancing in an extended LAN by means of path differentiation for traffic of different priorities. The present invention is also backward compatible with the IEEE 802.1D standard. The present invention supports multi-level traffic prioritization by extending the STAR Bridge Protocol to make use of hybrid forwarding paths as well as enhanced forwarding paths to forward frames in accordance with frame priorities. Since frames are forwarded over a plurality of alternate forwarding paths, the present invention also offers improved load balancing, thereby relieving congestion on certain tree links.

BRIEF DESCRIPTION OF THE FIGURES

The present invention will hereinafter be described in conjunction with the appended drawing figures, wherein like numerals denote like elements, and:

FIG. 1 is a diagram showing bridge protocol architecture for connecting two LANs via remote or local bridging;

FIG. 2 is a bridge state transition diagram for the STAR Bridge Protocol;

FIG. 3 is a port state transition diagram showing both the standard states and the additional states according to the STAR Bridge Protocol;

FIG. 4 shows an encapsulated data frame;

FIG. 5 is a STAR bridge operation flow chart;

FIG. 6 is a flow chart for the process BPDU_PROC;

FIG. 7 is a flow chart showing the procedure for processing DVCN_SBPDU frames;

FIG. 8 is a flow chart for the process Data_Frame_Proc;

FIG. 9 is a flow chart for process SLA_SBPDU_Proc.;

FIG. 10 is a bridged LAN graph useful in explaining the non-tree links between STAR bridges of the present invention;

FIG. 11 is an example of a STAR bridge graph;

FIG. 12 is a flow chart of the process DVC_SBPDU_Proc.;

FIG. 13 is a graph useful in explaining the computation of a tree path;

FIG. 14 is a flow chart of the process DVMy Info_Proc.;

FIG. 15 is a flow chart of the process DVOur Info_Proc.;

FIG. 16 is a flow chart of the process DV Inform_Proc.;

FIG. 17 is a flow chart of the process DV Record_Proc.;

FIG. 18 is a flow chart of the process DF_STAR_Learning_Proc.;

FIG. 19 is a flow chart of the process FD_Search_Proc.;

FIG. 20 is a flow chart of the process ESL_Search_Proc.;

FIG. 21 is a flow chart of the process Std_Data_Frame_Proc.;

FIG. 22 is a flow chart of the process DF_STAR_Forwarding_Proc.;

FIGS. 23–25 are diagrams useful in explaining the possible paths between end stations, showing tree paths and enhanced forwarding paths;

FIGS. 26–29 are diagrams useful in explaining examples of forwarding paths when at least one of the agent bridges of the end stations is not defined;

FIGS. 30–32 are diagrams useful in explaining examples of selection of tree paths when the agent bridges of both end stations are defined;

FIG. 33 is a diagram useful in explaining the selection of a tree path when the agent bridges of the end stations are defined and are different; and

FIG. 34 is a diagram useful in explaining the selection of a tree path when the agent bridges of the end stations are defined, are not the same and one is the ancestor of the other.

FIG. 35 is a bridge LAN graph useful in explaining the operation of the present invention.

FIG. 36 is a flow chart of the process Ann_SBPDU_Proc (n, n', c, p).

FIG. 37 is a STAR bridge operation flow-chart process incorporating the process Ann_SBPDU_Proc (n, n', c, p).

FIG. 38 is a flow chart of the process ESL_Search_Proc (s, t, pld).

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

I. Spanning Tree Alternate Routing (Star) Bridge Protocol I.A. Bridge States

The ensuing detailed description provides preferred exemplary embodiments only, and is not intended to limit the scope, applicability, or configuration of the invention. Rather, the ensuing detailed description of the preferred exemplary embodiments will provide those skilled in the art an enabling description for implementing a preferred exemplary embodiment of the invention. It being understood that various changes may be made in the function and arrangement of elements without departing from the spirit and scope of the invention as set forth in the appended claims.

FIG. 2 shows a bridge state transition diagram for the STAR Bridge Protocol. There are five states: Tree Learning, Tree Learned, Distant Neighbors Found, Direct Neighbors Found, and Enhanced states. The first two states capture the behavior of an old bridge in accordance with the IEEE 802.1D standard, while the other three states reflect the additional functionality of a STAR bridge. When the spanning tree is being built, each bridge (both old and STAR bridges) is in the Tree Learning state. In this state, all data frames are dropped as in the IEEE 802.1D standard. After the spanning tree is built, the bridge goes to the Tree Learned state. A path finding process starts and data frames are forwarded using the standard forwarding and learning processes. There are several phases in the path finding process. The bridge first goes to the Distant Neighbors Found state, then the Direct Neighbors Found state, and finally the Enhanced state in different phases. The transitions will be explained further hereinafter in the Path Finding Process section. When the bridge is in the Enhanced state, the path finding process is completed. STAR forwarding and learning processes are executed when a data frame is received in this state.

I.B. Port States

In the IEEE 802.1D standard, there are four port states: Blocking, Listening, Learning, and Forwarding. In the STAR Bridge Protocol, there are three additional port states for the path finding process, which are similar to the Listening, Learning, and Forwarding states. These port states, and the transactions among the port states are shown in FIG. 3. These new port states are distance vector listening (DV Listening), distance vector learning (DV Learning), and STAR Forwarding. The IEEE 802.1D Spanning Tree Bridge Protocol activates tree ports while the path finding process activates other useful non-tree ports. The transition among the four states in the standard remains the same for the STAR bridges. A port changes from the Blocking state to the Listening state when it is selected as a tree port. After an appropriate protocol timer expires, the port moves to the Learning state. In this state, learning is enabled but data frames will not be forwarded. The port is in the Forwarding state if the timer expires again. A port changes from the Blocking state to the DV Listening state if the path finding process selects it. A protocol timer is started when the port enters the DV Listening state. When this timer expires, the port moves to the DV Learning state. Unlike the Learning state in the present IEEE standard, in which a port learns the locations of end stations, a port in DV Learning state does not do this. It is because all data frames are still forwarded on tree paths and a port in DV Learning state must be a non-tree port. Hence, there should be no data frame arriving on that port. A similar protocol timer is used to time out the DV Learning state, and the next transition is into the Forwarding state.

I.C. Storage

Each bridge keeps an FD for the forwarding process. In STAR bridges, three tables are additionally used provided: bridge forwarding table (BF Table), end-station location table (ESL Table), and bridge address table (BA Table). These tables are preferably stored in memory devices already resident in non-STAR bridges, such as memory for storing FD. A BF Table indicates for each other STAR bridge in the bridged LAN the port that leads to the next hop along the “best” path found. BF Tables are obtained in the path finding process by a modified distance vector method as will be described in further detail hereinafter. An ESL Table is used to map an end station to a STAR bridge near it. The STAR learning process is responsible for filling this table.

Therefore, if the ESL Table of a STAR bridge has a record for an end station, it is unnecessary for the FD of the STAR bridge to have a record for the same end station in most situations. This implies that the FD in a STAR bridge is no larger than that in a standard bridge. The BA Table is a mapping between STAR bridge identifiers and STAR bridge MAC addresses. Every bridge has its own unique MAC address and this MAC address is used in the STAR forwarding process. It is not necessary for a STAR bridge to know the MAC address of all other STAR bridges. It should be noted that $\text{addr}(n)$ will be the MAC address of a bridge with a bridge ID n .

I.D. Protocol Data Units

The STAR Bridge Protocol recognizes two types of protocol data units, namely BPDU (Bridge Protocol Data Units), which is specified in the IEEE 802.1D standard, and SBPDU (STAR Bridge Protocol Data Units), which are specified below for the STAR Bridge Protocol.

The SBPDU contains an SBPDU Header, which has the same format as the BPDU Header, and a set of SBPDU Parameters. The Protocol Identifier in the SBPDU has its own unique value to identify the STAR Bridge Protocol. SBPDU MAC frames assume a format similar to that of BPDU MAC frames. There are four kinds of SBPDU frames: Hello SBPDU, Distance Vector Change Notification SBPDU (DVCN_SBPDU), Distance Vector Computation SBPDU (DVC_SBPDU), and Station Location Announcement SBPDU (SLA_SBPDU). Hello SBPDU is used for monitoring crosslink failures. DVCN_SBPDU is used to notify STAR bridges about topology changes. DVC_SBPDU is used in the path finding process to compute the distance vectors and BF Tables. SLA_SBPDU is generated in the STAR learning process to fill the ESL Tables.

Since any BPDU frame with an unknown Protocol Identifier will not be forwarded by old bridges, an SBPDU frame sent by a STAR bridge must be encapsulated as a data frame if it is expected to traverse at least one old bridge. The source address of an encapsulated SBPDU frame is the MAC address of the STAR bridge that is responsible for the encapsulation. The destination MAC address is either the STAR bridge group MAC address or the unique MAC address of the intended STAR bridge recipient. Which destination MAC address is applicable depends on the specific type of SBPDU MAC frame being sent. In the proposed protocol, most SBPDU MAC frames are sent over the spanning tree. It will be explained herein that some SBPDU MAC frames may be sent to direct STAR neighbors over selected non-tree links. Except in the only case described in with reference to crosslink maintenance and Path Finding Process, an SBPDU MAC frame received by a STAR bridge will not be forwarded by the STAR bridge.

I.E. Data Frames

Each data frame at the LLC Sublayer contains LLC information and is encapsulated within a MAC frame at the MAC layer of the end station. We refer a MAC data frame generated by an end station as a normal MAC data frame, or simply a normal data frame when the context is clear. As will be explained in greater detail hereinafter with reference to the frame dropping problem, a normal data frame may have to be encapsulated by a bridge for forwarding purpose. When a normal MAC data frame is encapsulated, we refer to it as an encapsulated MAC data frame, or simply an encapsulated data frame when the context is clear. We refer to the intended recipient of an encapsulated data frame as a proxy destination, and its address as a proxy destination address. In the rest of this document, the phrase "data frame"

refers to either a normal data frame or an encapsulated data frame. We assume that a MAC data frame fr has the form $\langle \text{src}(fr), \text{dst}(fr), \text{pld}(fr) \rangle$ where $\text{src}(fr)$, $\text{dst}(fr)$, and $\text{pld}(fr)$ represent the source MAC address, the destination MAC address, and the payload data respectively.

A normal MAC data frame consists of a MAC Header, an LLC Header, a layer 3 packet, and a MAC Trailer. When a normal MAC data frame is encapsulated, an additional MAC Header and LLC header are put in the front of the normal data frame and an additional MAC Trailer is put at the end as shown in FIG. 4. Each STAR bridge has to distinguish among the following frames: normal data frames, encapsulated data frames, and encapsulated SBPDU frames. All these frames appear as normal data frames to old bridges. The present invention specifies frame type information in the Protocol Type field of the encapsulating LLC header so that a STAR bridge can correctly identify the frame type and process the frame.

In most bridged LANs, frames are subject to a maximum transfer unit (MTU) constraint. An encapsulated MAC frame of a given size may have to be fragmented before encapsulation, or the resulting frame will violate the MTU constraint. In this respect, STAR bridges must implement a fragmentation and reassembly mechanism to accommodate the encapsulation that may be needed for forwarding certain data frames. It follows that each encapsulated frame must include an additional field to carry an appropriate sequence number. It is sufficient to fragment an oversized frame into two fragments since the encapsulation overhead is much less than an MTU. Hence, only one bit is needed to identify a pair of fragments. Fragmentation and reassembly mechanisms are beyond the scope of the present invention since, frame encapsulation is already an existing function in present day bridges.

I.F. Crosslink Maintenance

Since a crosslink is not part of the active topology of a spanning tree, topology changes that involve crosslinks normally will not trigger a bridge to send out a Topology Change Notification BPDU. Therefore, a mechanism is needed for monitoring and updating the status of each crosslink selected by the proposed protocol for supporting an enhanced forwarding path. Hello SBPDUs and DVCN_SBPDU are used for this purpose. A Hello SBPDU consists merely of a Protocol Identifier field, a Protocol Version Identifier field, and an SBPDU Type field with a code reserved for this type. A DVCN_SBPDU contains an SBPDU Header and IDs of both bridges on the ends of the crosslink that is being identified to have failed.

In order to detect crosslink failures, STAR bridges on both ends of each crosslink exchange Hello SBPDUs periodically. These Hello SBPDUs are not forwarded by the STAR bridges to their neighbors. Each STAR bridge uses a pre-determined timer for each of its crosslink neighbors to time out pending Hello SBPDUs. If a STAR bridge receives a Hello SBPDU from a STAR neighbor over a crosslink before an appropriate timer expires, the STAR bridge resets the timer. Otherwise, the STAR bridge assumes that the crosslink has failed and then transits to the Tree Learned state. In addition, the STAR bridge also multicasts a Distance vector change Notification SBPDU over the IEEE 802.1D spanning tree to all STAR bridges.

When a crosslink is recovered from a recent failure or a new crosslink link is enabled, the STAR bridges on both ends of the crosslink will first check to determine if the crosslink is eligible for supporting an enhanced forwarding path based on information in their respective BF Tables. If

the crosslink is eligible, the STAR bridges will each independently multicast a DVCN_SBPDU over the IEEE 802.1D spanning tree to all STAR bridges.

Upon receiving a DVCN_SBPDU, a STAR bridge forwards it to all of its tree neighbors, and transits to the Tree Learned state. The STAR bridge makes use of a timer to remember the identity of the affected crosslink for a predetermined time-out period. When a DVCN_SBPDU identifying a given crosslink is received by a STAR bridge, and if no other DVCN_SBPDU identifying the same crosslink has been received within a current time-out period, the STAR bridge resets the timer and forwards the SBPDU to all its tree neighbors. Otherwise, the SBPDU is dropped.

I.G. Bridge Operation

As mentioned earlier, there are three kinds of MAC frames a STAR bridge would receive in this protocol: BPDU frames, SBPDU frames, and data frames. FIG. 5 is the STAR bridge operation flow-chart. When a MAC frame is received, it invokes different procedures for different kinds of frame.

BPDU frames received are processed by a procedure BPDU Proc. FIG. 6 depicts the flow-chart for BPDU_Proc. When a Topology Change Notification BPDU is received, the bridge has to deactivate the non-tree ports selected by the path finding process and invalidate the entries of its ESL table. After that, standard BPDU processing is executed. In FIG. 6, Std_BPDU_Proc refers to the standard BPDU processing procedure. As the standard BPDU processing procedure can be found in the IEEE 802.1D standard which is incorporated herein by reference, the details of that procedure have been omitted for purposes of brevity.

There are four kinds of SBPDU frames as described in the Protocol Data Units Section. There is a procedure for processing each kind of SBPDU frames. DVCN_SBPDU Proc procedure is for processing DVCN_SBPDU frames, and the flow-chart for this procedure is shown in FIG. 7. The details of DVC_SBPDU_Proc procedure for processing DVC_SBPDU frames and SLA_SBPDU Proc procedure for processing SLA_SBPDU frames will be discussed hereinafter with reference to Section V and Section IV respectively. The flow-charts for these procedures are shown in FIG. 12 and FIG. 9 respectively. A STAR bridge invokes a Data_Frame_Proc procedure when a data frame is received. The flow-chart for Data_Frame_Proc is shown in FIG. 8. Different procedures are executed depending on the current state of the bridge. The processing of data frames will be discussed further in Section IV and Section V.

II. Model

In this section, we describe the mathematical model we use in this invention. We will also define the notations for the proposed protocol. A summary can be found in Table A in the Appendix.

In the present invention the bridged LAN is represented as an undirected graph $G=(V, E)$ where V is the set of all bridges and E is the set of links connecting the bridges. Each link $(x, y) \in E$ is assumed to have a non-negative cost $c(x, y)$. For convenience, we let $c(x, y)=\infty$ if $(x, y) \notin E$. If there are several links between bridge x and bridge y , $c(x, y)$ should be the minimum among the costs of the links. A path in G is a loop-free tandem concatenation of links in E . The length of a path is the sum of the costs of all the links along the path. The distance between a pair of nodes, x and y , is the length of a shortest path connecting the nodes.

Bridge x is a direct neighbor of bridge y , and vice versa, if $(x, y) \in E$. $T=(V, E_T)$ is a tree subgraph of G representing an RST, wherein $(x, y) \in E_T$ if and only if (x, y) is an activated

link in the RST. The links in E_T are referred to as tree links and the links in $E \setminus E_T$ as non-tree links. If $(x, y) \in E_T$, x and y are tree neighbors. A path in T is a tree path. A tree path originating at bridge s and terminating at bridge t is denoted $treepath(s, t)$. The distance of this tree path is denoted $d_T(s, t)$. Note that $d_T(x, y)=c(x, y)$ if x and y are tree neighbors. We refer to $treepath(s, t)$ are referred to as an old bridge tree path if it has at least one intermediate bridge (i.e., one other than the source and destination bridges) and every intermediate bridge on the path is an old bridge. B represents the set of STAR bridges. If s and t are STAR bridges, that is, $s, t \in B$, and there is an old bridge tree path between them, s is a distant STAR neighbor of t , and vice versa. If, in addition, s is an ancestor of t , then s is a unique distant STAR ancestor neighbor of t . Henceforth, the distant STAR ancestor neighbor of a bridge t will be referred to by $dsan(t)$. The set of distant STAR neighbors of n is represented by $N'_B(n)$.

The nearest common ancestor of x and y is the highest-level bridge on a tree path between x and y . If x is an ancestor of y , then x is necessarily the nearest common ancestor of x and y . Let the nearest common ancestor of x and y be denoted $nca(x, y)$. We say x and y are on different branches if $nca(x, y) \neq x$ and $nca(x, y) \neq y$. We call $(x, y) \in E \setminus E_T$ a crosslink if x and y are on different branches. x and y are crosslink neighbors then. FIG. 10 is an example of an undirected graph of a bridged LAN. Node r is the root. The solid lines are links in E_T and the dotted lines are non-tree links. Link (u, q) is a non-tree link but not a crosslink while (w, y) , (u', v) and (v', z) are all crosslinks. Therefore, u and q are direct neighbors but neither tree neighbors nor crosslink neighbors. They are distant STAR neighbors though since the path $u \rightarrow y \rightarrow q$ is an old bridge tree path. Table 1 summarizes the definitions of different kinds of neighbors.

TABLE 1

Neighbor Types	
Neighbor Type	Definition
Direct neighbors	$x, y \in V, (x, y) \in E$
Tree neighbors	$x, y \in V, (x, y) \in E_T$
Crosslink neighbors	$x, y \in V, (x, y) \in E \setminus E_T$ and $nca(x, y) \neq x$ and $nca(x, y) \neq y$
Direct STAR neighbors	$x, y \in B, (x, y) \in E$
Distant STAR neighbors	$x, y \in B, treepath(x, y)$ is an old bridge tree path

Since an old bridge sets the port associated with each of its non-tree links to blocking state, there is no way to use any such a link to forward data frames, even though the other side of the link is a STAR bridge. Therefore, a non-tree link may be used only if it connects two STAR bridges. Even so, such a link may not support any shortest path. There is yet another reason that a non-tree link is useless. Consider FIG. 10 where both bridges u and q are STAR bridges. This non-tree link between them is useless because the distance of this link must be larger than the tree path from u to q ; otherwise, the spanning tree algorithm would have set q to be a child of u . In FIG. 10, (u, q) and (w, y) are ineligible links. A non-tree link that is obviously useless, as just described, for supporting any shortest path is termed an ineligible link. A non-tree link is termed eligible otherwise. In FIG. 10, (u', v) and (v', z) are eligible crosslinks.

A STAR bridge graph is defined as $G_B=(B, E_B)$ where B is the set of STAR bridges and $(x, y) \in E_B$ if and only if x and y are STAR neighbors, either direct or distant. $c(x, y)$, the cost of link (x, y) , is defined by the following formula.

$$c(x,y) = \begin{cases} c(c,y) & \text{if } x \text{ and } y \text{ are direct but not distant STAR neighbors} \\ \min(d_T(x,y), c(x,y)) & \text{if } x \text{ and } y \text{ are both direct and distant STAR neighbors} \\ d_T(x,y) & \text{if } x \text{ and } y \text{ are distant but not direct STAR neighbors} \end{cases}$$

The STAR bridge graph of FIG. 10 is shown in FIG. 11.

III. Path Finding Process

The goal of this process is to compute the BF Table. In the best case, the BF Table has next hop and forwarding port information associated with a shortest path to a STAR bridge. The STAR bridge graph contains all tree paths among STAR bridges and all eligible non-tree links in the original bridged LAN. Therefore, the shortest path in G_B between a pair of STAR bridges x and y would be the best path which can be achieved in the bridged LAN after pruning ineligible links. Ideally, if every $c'(x, y)$ can be computed correctly, each STAR bridge can compute its own BF Table based on distance vectors. Nonetheless, we will still use a distance vector approach for updating the BF Table, except that some modification as described below is needed.

In a conventional distance vector update protocol, each node initializes its distance vector with distances to all its neighbors. It then sends the distance vector to all its neighbors. When a neighbor receives the distance vector, the neighbor updates its own distance vector if any shorter path is found. This neighbor then sends its update to all its own neighbors. The procedure keeps on going until the algorithm converges.

As there are old and STAR bridges in the bridged LAN, the conventional distance vector update protocol cannot be applied directly. Each bridge knows only the cost to its direct neighbors. If x and y are distant STAR neighbors, both x and y do not know $d_T(x, y)$ since there are one or more old bridges between them. Therefore, a bridge $x \in B$ has to determine $d_T(x, y)$ if y is a distant STAR neighbor. Unfortunately, due to the limitation of old bridges, STAR bridges may not be able to determine every distance correctly. Nonetheless, the distances may be estimated such that each estimated distance is at least its corresponding real distance.

The process for path finding in the proposed protocol consists of two procedures:

Distance Vector Estimation

Distance Vector Enhancement

In these procedures, STAR bridges communicate using Distance Vector Computation SBPDUs. In the Distance Vector Estimation procedure, each STAR bridge initializes the distance to its neighbors in the STAR bridge graph of the bridged LAN. It involves discovery of distant STAR neighbors and computation of the tree path distances as described above. In the Distance Vector Enhancement procedure, which follows the Distance Vector Estimation procedure, STAR bridges exchange their distance vectors, discover other non-neighbor STAR bridges, and find the shortest path to them.

In the process, a STAR bridge n maintains a distance vector only for other STAR bridges known to n . As unknown STAR bridges are discovered by n in the process of the algorithm, a new entry is created in the distance vector maintained by n for each newly discovered STAR bridge. When the process ends, n should have discovered all other STAR bridges and its distance vector will consist of one entry for each STAR bridge $n' \in B \setminus \{n\}$. Each entry in the

distance vector of n consists of a tuple of seven fields. The entry associated with n' in the distance vector of n is denoted as $DVT(n, n')$. As summarized in Table 2, the information contained in $DVT(n, n')$ provides an estimated distance between n and n' , indicates whether the estimated distance is actually accurate, and enables STAR bridge n to know its forwarding port for n' , its next hop STAR bridge neighbor on the forwarding path to n' , indicates whether the forwarding path is a tree path, as well as whether n' is an ancestor or a descendant. Incidentally, $d(n, n')$, the current estimated distance from n to n' , will be appropriately initialized as described in Sections III.A and III.B.

TABLE 2

Fields in $DVT(n, n')$ for a Path from n to n'	
Field	Definition
N'	ID of destination STAR bridge
$d(n, n')$	Estimated distance between n and n'
$F(n, n')$	Forwarding port for n'
$next(n, n')$	ID of the next hop STAR bridge neighbor on the path from n to n'
$FG_A(n, n')$	Distance accuracy flag with a value 1 if $d(n, n')$ is accurate, and 0 otherwise
$FG_T(n, n')$	Tree path flag with a value 1 if the path from n to n' is a tree path, and 0 otherwise
$FG_R(n, n')$	Relation flag with a value 1 if n' is an ancestor of n , -1 if n' is a descendant of n , and 0 otherwise

III.A Distance Vector Estimation

In the DV Estimation procedure, a STAR bridge n discovers all its STAR neighbors, both direct and distant, that is, the fields in the distance vector of STAR bridge n are filled out for each STAR bridge $n' \in B \setminus \{n\}$ where $(n, n') \in E_B$. There are two phases in the Distance Vector Estimation procedure. In the first phase, STAR bridge n estimates $d_T(n, n')$ if n' is a distant STAR neighbor and n fills out the entry for n' in $DVT(n, n')$. Note that $d(n, n')$ is equal to the estimated $d_T(n, n')$ in this phase. In the second phase, n determines and fills out the entry for k if $k \in N_B(n)$. If k is a distant STAR neighbor of n , then n replaces $d(n, k)$ by $c(n, k)$ and other fields accordingly only if it is appropriate. This phase will be discussed in more detail later.

Before the spanning tree algorithm starts, each bridge, either old or STAR, should know its own ID and the cost of the link to each of its direct neighbors in the bridged LAN. After the tree has been built, every bridge k will also know its tree links, as well as the root bridge and the root path distance, $d_r(k)$, where bridge r is the root bridge. Incidentally, $d_r(r)=0$. Table 3 is the topology information of bridge v in FIG. 10 after the spanning tree computation. The column Old/STAR is the information obtained by STAR bridges only while all bridges, either old or STAR, obtain all other columns.

TABLE 3

Topology Database of Bridge v				
Bridge	Type	Distance	P(v)	Old/STAR
r	Root	$d_r(v)$	$p_r(v)$	N/A
x	Parent	$c(v, x)$	$p(v, x)$	Old
u'	Non-tree neighbor	$c(v, u')$	$p(v, u')$	STAR

There are three kinds of DVC_SBPDU frames in the first phase—DVMyInfo, DVOurInfo, and DVInform frames. Table 4 shows the formats of these frames. DVMyInfo frames are used for a STAR bridge to inform other STAR bridges of its own topology information. DVOurInfo frames carry information related to both the source and the destination STAR bridges. DVInform frames allow STAR bridges to pass on topology information of other STAR bridges. DVRecord frames are used in Distance Vector Enhancement and will be described in Section III.B DVC_SBPDU_Proc procedure first identifies the frame and then invokes corresponding procedures. The flow-chart for DVC_SBPDU_Proc is shown in FIG. 12.

Each STAR bridge k sends a DVMyInfo(k) frame on its root link only if its parent is an old bridge. If there is any STAR bridge along the root path of k, the one that is nearest to k, say n, will receive the DVMyInfo(k) frame from a child link. Note that n and k are distant STAR neighbors. n is the distant STAR ancestor neighbor of k and $n=dsan(k)$. Bridge n can determine the tree path distance between k and n, wherein $d_T(n, k)=d_r(k, n)$ is the difference of their root path distances. Then, n informs k of the distance between them by a DVOurInfo(n, k) frame and stops forwarding the DVMyInfo(k) frame. In this case, n and k are on the same branch and n is an ancestor of k. If the STAR bridges are on different branches, like v and v' in FIG. 10, and there is no STAR bridge on the tree path between them, then, since DVMyInfo frames are multicast frames addressed to all STAR bridges, v' will receive the DVMyInfo(v) frame of v and vice versa. However, there is no way for them to calculate the real tree path distance between them using root path distance alone. In the case of v and v', if they know that they have the same parent, they can determine the accurate tree path distance. Therefore, the DVMyInfo(v) frame also contains the information of the parent of v. When v and v' receive each other's DVMyInfo frame through a root link and find out they are siblings, they may calculate the distance between them correctly by adding $c(v, parent(v))$ and $c(v', parent(v'))$.

TABLE 4

Format of DVC SBPDU Frames			
SBPDU Frame	Format	Source	Destination
DVMyInfo(n)	$\langle n, d_r(n), parent(n), c(n, parent(n)) \rangle$ where $parent(n) \in V \setminus B$	$n \in B$	Multicast address of B
DVOurInfo(n, n')	$\langle n, n', d(n, n'), cb(n, n'), c(n, cb(n, n')) \rangle$ where $n' \in B$, and $cb(n, n') \in V \setminus B$	$n \in B$	$n' \in B$
DVInform(n, n')	$\langle n, n', d(n, n') \rangle$ where $n' \in B \setminus \{n\}$	$n \in B$	$k \in B \setminus \{n\}$ and $k \in N_B'(n)$
DVRecord(n, n')	$\langle n, n', d(n, n'), FG_A(n, n'), FG_T(n, n'), FG_R(n, n') \rangle$ where $n' \in B \setminus \{n\}$	$n \in B$	$k \in B \setminus \{n, n'\}$ and $k \in N_B'(n)$

FIG. 13 shows an example in which it is not sufficient to calculate the correct distance. Nevertheless, n and j may obtain an overestimate of the true distance between them by simply adding $d_T(k, n)$ and $d_T(k, j)$. It can be done by requiring k, the distant STAR ancestor neighbor, to send the information $d_T(k, j)$ to n using a DVInform(k, j) frame. If n also knows $c(k, m)$, where m is the child of k on the path to n, it can get an even better estimate. Therefore, m and $c(k, m)$ are sent by k in a DVOurInfo(k, n) frame. Note that $m \in V \setminus B$. We denote by $cb(k, n)$ the child bridge of k on a tree path leading from k to n. Let $dsanc(n)$ be used by bridge n to keep track of $cb(k, n)$, the child of the $dsan(n)$, and $dcc(n)$ to denote the doubly counted cost of the link between $cb(k, n)$ and k. In this case, $dsanc(n)$ and $dcc(n)$ are respectively set to $cb(k, n)$ and $c(k, dsanc(n))$.

Upon receiving a DVMyInfo(k) frame $\langle k, d, n', c \rangle$ from tree port p, a STAR bridge self, with its root path distance $dr(n)$, processes the frame as in Pseudocode 1. FIG. 14 shows the flow-chart corresponding to Pseudocode 1 (DVMyInfo_Proc). When p is a child port, k is a descendant of elf and self can calculate the tree path distance between them and inform k. There are two cases of interest in the processing of a DVMyInfo frame when p is a root port. In Case 1.1, bridges self and k are siblings with a common old bridge parent. In Case 1.2, self and k are on different branches but they are not siblings and the distance $d(self, k)$ evaluated in this case may be incorrect. It is correct only if the nearest common ancestor of self and k is the root bridge.

If n and k do have a STAR common ancestor n, n can estimate $d(n, k)$ better when it receives the DVOurInfo(n', n) and DVInform(n', k) frames from n'. STAR bridge self processes DVOurInfo frame $\langle n', k, d, m, c \rangle$ as in Pseudocode 2. FIG. 15 shows the flow-chart corresponding to Pseudocode 2 (DVOurInfo_Proc).

Each STAR bridge n receives at most one DVOurInfo(n', n) frame, and this frame must be sent from $dsan(n)$, which is n'. In Pseudocode 2, the frame DVOurInfo(n', k) is dropped when $self \neq k$ because the recipient of the frame is self and there is no other STAR bridge along the path from self to n'. In Pseudocode 1, a DVOurInfo(n', n) frame is sent before sending any DVInform frame. As a DVInform frame will not propagate beyond its distant STAR neighbors and bridges do not reorder frames, we can conclude that n receives the DVOurInfo(n', n) frame before receiving any DVInform frame sent by n'. Therefore, when bridge n receives a DVInform(n', j) frame $\langle n', j, d \rangle$, it should have $dsanc(n)$ and $dcc(n)$ correctly assigned. The pseudocode for bridge self to execute when the DVInform(k, j) frame $\langle k, j, d \rangle$ is received is shown in Pseudocode 3. FIG. 16 shows the flow-chart corresponding to Pseudocode 3.

ASTAR bridge self processes a DVInform(k, j) frame <k, j, d> according to the following different situations:

- Case 3.1: self=j, that is, k is informing self the distance between them
- Case 3.2: self≠j, and, either d(self, k) or d(self, j) is unknown
- Case 3.3: self≠j, d(self, j) is estimated
 - Case 3.3a: parent(self)=dsanc(self)
 - Case 3.3b: parent(self)=dsanc(self)
- Case 3.4: self≠j, d(self, k) is known, and d(self, j) is accurate

If d(self, j) is accurate, the bridge self does not have to do anything. In Case 3.3 when FG_A(n, j)=0, it means d(self, j) is an estimate and so a better estimate can be obtained. Case 3.3a is the situation where dsanc(self) is the nearest common ancestor of self and j. As a result, self can calculate the distance correctly. Since j may not be able to do so (see

k) yet and so initializes it to be (c(n, k), p(n, k), k, 1, 0, 0). On the other hand, if k is also a distant STAR neighbor, d(n, k) should, if possible be assigned to be min(d_T(n, k), c(n, k)), and other fields accordingly. If the estimated d_T(n, k) is surely correct, that is, if FG_A(n, k)=1, it is trivial. Unfortunately, d_T(n, k) may be incorrect. Since we are not sure whether the direct link (n, k) is shorter than the tree path from n to k, d_T(n, k) won't be replaced in order to avoid selecting a link with a larger distance than its corresponding tree path. The same applies whenever the distance vector is enhanced. When the tree path distance between a pair of STAR bridges is only an estimate, the tree path won't be replaced between them by a non-tree path. Therefore, when the distance vector becomes stable, d(n, n') ≤ d_T(n, n') for all n, n' ∈ B such that n ≠ n'.

PROCEDURE: DVMyInfo Proc(k, d, n', c, p), also see FIG. 14.

```

Begin
  If p = pr(self) /* p is the root port, self and k are on different
branches */
    If parent(self) = n' /* Case 1.1: self and k are siblings
*/
      d(self, k) := c(self, n') + c;
      FG_A(self, k) := 1
      /* Case 1.2: self and k are not
siblings */
    Else
      d(self, k) := dr(self) + d; /* overestimated d(self, k)
*/
      FG_A(self, k) := 0
    Endif
    F(self, k) := p;
    next(self, k) := k;
    FG_T(self, k) := 1;
    FG_R(self, k) := 0
  Else /* p is a child link port, i.e., self is an ancestor of k */
    DVT(self, k) := (d - dr(self), p, k, 1, 1, -1);
    Send DVOurInfo(self, k) frame to k
    <self, k, d(self, k), cb(self, k), c(self, cb(self, k))>
    For each bridge j where F(self, j) = F(self, k)
      /* j and k are from the same child
port */
      Send to j DVInform(self, k) frame
      <self, k, d(self, k)>
      Send to k DVInform(self, j) frame
      <self, j, d(self, j)>
    endif
  end
end

```

FIG. 13), self has to inform j by a DVInform(self, j) frame. When j receives that, which is Case 3.1, it can enhance the distance. Case 3.3b is the situation that accurate distance cannot be found but the estimated distance can be improved by taking out the doubly counted cost dcc(n). Since DVInform frames are sent after DVOurInfo and DVMyInfo frames, Case 3.2 occurs only if there is an error.

Ultimately, the first phase must terminate. When it terminates, n should have a correct or an overestimate d_T(n, n') for every distant STAR neighbor n'. We then proceed to the second phase, that is, each STAR bridge n will fill out DVT(n, k) if k is a direct STAR neighbor of n. If k is a tree neighbor, n should know c(n, k) from the topology database, it can initialize the entry DVT(n, k) to be (c(n, k), p(n, k), k, 1, 1, 1) if k is a parent and (c(n, k), p(n, k), k, 1, 1, -1) if k is a child. If k is a crosslink neighbor, there are two cases. If k is not a distant STAR neighbor, n doesn't have DVT(n,

Pseudocode 1: DVMyInfo_Proc

PROCEDURE: DVOurInfo_Proc(n', k, d, m, c), also see FIG. 15

```

55 Begin
  If self ≠ k /* error */
    Drop the frame
  Else
    d(self, n') := d;
    F(self, n') := pr(self);
    next(self, n') := n';
60 FG_A(self, n') := 1;
    FG_T(self, n') := 1;
    FG_R(self, n') := 1;
    dsanc(self) := m;
    dcc(self) := c
  endif
65 end

```

Pseudocode 2: DVOurInfo_Proc

PROCEDURE: DVInform_Proc(k, j, d), also see FIG. 16.

```

Begin
If j = self                               /* Case 3.1 */
  d(self, k) := d ;
  FG_A(self, k) := 1
Else If d(self, k) is unknown or d(self, j) is unknown /* Case 3.2 */
  Drop the frame
Else if FG_A(self, j) = 0                 /* Case 3.3
*/
  If parent(self) = dsanc(self)          /* Case 3.3a */
    d(self, j) := c(self, parent(self)) + (d - dcc(self)) ;
    FG_A(self, j) := 1 ;
    Send DVInform frame < self, j, d(self, j)> to j
  Else                                    /* Case 3.3b */
    d(self, j) := (d(self, k) - dcc(self)) + (d - dcc(self))
  endif
endif
end

```

Pseudocode 3: DVInform_Proc

III.B Distance Vector Enhancement

The DV Enhancement procedure is similar to the distance vector exchange procedure in the traditional approach except that we can replace a tree path only if its exact distance is known. That is, $d(n, n')$ in the distance vector can be replaced only if $FG_A(n, n')=1$ or $FG_T(n, n')=0$. On the other hand, if a tree path from n to n' is found and $d_T(n, n')$ is only an estimate, it replaces $d(n, n')$ in the distance vector. After the DV Estimation procedure, bridge n only knows the tree distance, either correct or estimate, to its tree neighbors and distant STAR neighbors. In order to let a bridge identify whether a path to a formerly unknown bridge is a tree path and whether the tree path distance is correct, the accuracy flag and the tree path flag must be put in the DVRecord frames. The format of a DVRecord(n, n') frame sent by bridge n is $\langle n, n', d(n, n'), FG_A(n, n'), FG_T(n, n'), FG_R(n, n') \rangle$. Pseudocode 4 is the DV Enhancement procedure of STAR bridge self upon receiving DVRecord(j, k) frame $\langle j, k, d, FG_A, FG_T, FG_R \rangle$. FIG. 17 shows the flow-chart corresponding to Pseudocode 4 (DVRecord_Proc).

In Pseudocode 4, when the path from self to j and the path from j to k are both tree paths, the path from self to k through

j must be a tree path too. This is Case 4.1. $FG_R(self, k)$ flag is updated accordingly and the tree path information is sent to other STAR neighbors so that the other STAR tree neighbors discover the tree path leading to k . If the tree path distance is an estimate, it always replaces the existing $DVT(self, k)$; otherwise, it replaces only if it is better in terms of distance. Case 4.2 is the case where a non-tree path is found. If the existing $DVT(self, k)$ is not an estimated tree path, it can be replaced if the newly found path is shorter. After the algorithm converges, the distance vector can be reduced to the BF Table. Bridge n has to keep $F(n, n')$ and $next(n, n')$ for every bridge $n' \in B \setminus \{n\}$ for forwarding purpose. A path in the BF Table between a pair of STAR bridges is referred to as a STAR forwarding path. Note that a STAR forwarding path may be a standard tree path or an enhanced forwarding path when it can be identified.

PROCEDURE: DVRecord_Proc(i, k, d, FG_A, FG_T, FG_R), also see FIG. 17.

```

Begin
  If k is unknown                               /* k is newly discovered
*/
    DVT(self, k) := ( $\infty$ , pt(self), j, 0, 0, 0) /* initialize DVT (self k) */
  Endif
  If FG_T(self, j) = 1 and FG_T = 1             /* Case 4.1: a tree path to k is
found */
    If FG_R(self, j) = 1 and FG = 1             /* k is an ancestor of self
*/
      FG_R(self, k) := 1
    Else if FG_R (self, j) = -1 and FG_R = -1   /* k is a descendant of self
*/
      FG_R(self, k) := -1
    endif
  If FG_A (self, j) = 0 or FG_A = 0 /* tree path distance is a estimate
*/
    FG := 0
    DVT(self, k) := (d(self, j) + d, F(self, j), j, 0, 1, FG_R (self, k))
  Else
    FG := 1
  Endif
  Send DVRecord <self, k, d(self, j) + d, FG, 1, FG_R (self, k)>
  to all STAR neighbors except j
  If d(self, j) + d < d(self, k)                /* tree path distance is correct

```

-continued

```

and better */
      DVT(self, k) := (d(self, j) + d, F(self, j), j, 1, 1, FG_R(self, k))
    endif
  Else
    /* Case 4.2: a non-tree path is
found */
    If FG_A(self, k) = 1 or FG_T(self, k) = 0 /* original path can be
replaced */
      If d(self, j) + d < d(self, k)
        DVT(self, k) := (d(self, j) + d, F(self, j), j, 1, 0,
FG_R(self, k))
      Send DVRecord < self, d(self, k), 1, 0, FG_R(self,
k)> to
        all STAR neighbors except j
      endif
    endif
  endif
end

```

Pseudocode 4: DVRecord_Proc

IV. STAR learning process

When the path finding process is done, a STAR bridge should have filled its BF Table. It then starts the STAR learning process and the STAR forwarding process to forward data frames. The BF Table alone is not enough to forward a data frame since the data frame contains the address of an end station instead of a bridge. Therefore, the STAR learning process has to learn the location of end stations and store the information. An ESL Table is used to map end stations to STAR bridges.

IV.A Designated Bridge and Agent Bridge

Each LAN has a designated bridge and this bridge is also the designated bridge of all end stations attached to that LAN. The designated bridge of an end station s is designated to be $db(s)$. A designated bridge may be an old bridge or a STAR bridge.

If an end station can be mapped to its "closest" STAR bridge, as the distance vector is correct, the bridges will forward the frame according to a STAR forwarding path. This "closest" bridge is referred to as the agent bridge of the end station. This agent bridge must be a STAR bridge for the BF Table contains STAR bridges only. Once a STAR bridge identifies itself to be an agent bridge of a formerly unknown end station, it is responsible to tell other bridges so that they know where to forward a frame destined for that station. This can be done by using Station Location Announcement SBPDU frames. For convenience, we refer a Station Announcement SBPDU frame as a StationLoc frame. The destination address of all StationLoc frames is the multicast address of bridge group B. Let the format of a StationLoc frame be <end station address, agent bridge ID>. The expression $ESL(n, s)$ represents $ab(s)$ in the ESL Table of n .

The designated bridge is a suitable candidate to be an agent bridge if it is a STAR bridge. If the designated bridge is a STAR bridge, it can announce itself to be the agent bridge. However, an old bridge does not do that. As a result, we need to find a STAR bridge that is reasonably close to the old, designated bridge of the end station to be the agent. When an old bridge forwards a data frame, it may send it to more than one child link but at most one root link. Therefore, it would be undesirable to have a STAR bridge on the downstream to be an agent for we may end up having more than one agent. Among all the STAR bridges along the upstream, the one closest to the old bridge is preferred. However, there are situations that no agent bridge is identified. For example, if the designated bridge is the root that is an old bridge, all the STAR bridges will be on downstream and so no bridge will declare as the agent bridge. In that case, the tree path is used for forwarding and we record the forwarding port in the FD as what the old bridges do. We say

the agent bridge is undefined for an end station if there is no bridge declared as the agent. In other words, the agent bridge is unknown to all STAR bridges.

IV.B End Station Location Table

Each entry in the ESL Table of STAR bridge n is a tuple $(s, ab(s))$ where s is an end station and $ab(s)$ is the agent bridge of s . Each such entry in the ESL Table of n is created when an unknown end station is newly discovered by n . Each entry in the FD of STAR bridge n indicates a forwarding port of n for an end station s , that is $f(n, s)$. Each such entry in the FD of n is created when an unknown end station is newly discovered by n .

The STAR learning process of STAR bridges is responsible for filling out these entries. STAR bridge n fills out the ESL Table using the information in the StationLoc frames received. The FD can be filled as in the IEEE 802.1D standard, in which case n records the port from where a normal data frame arrives. The set of end stations in the ESL Table of n is $H(n)$ and the set of end stations in the FD of n is $S(n)$. An end station s is referred to as a known end station with respect to n if $s \in H(n) \cup S(n)$; s is unknown otherwise. Note that, after timing out old entries in the FD, $H(n) \cap S(n)$ is a set of end stations where the agent bridge of each of the end stations is n . In addition to filling out the ESL Table and FD, n should be able to identify whether it is the agent bridge of an end station s when a normal data frame from s is received.

IV.C Procedures for STAR Learning Process

Whenever a STAR bridge is in the Enhanced state, it executes the STAR Learning process. The bridge invokes a SLA_SBPDU_Proc procedure upon receiving a StationLoc frame, and a $DF_STAR_Learning_Proc$ procedure upon receiving a normal data frame. The pseudocodes for these procedures are shown in Pseudocode 5 and Pseudocode 6 respectively. In both pseudocodes, whenever an entry in the ESL Table or the FD is updated or created, its corresponding timer is reset. FIG. 18 and FIG. 19 show the flow-charts corresponding to Pseudocode 5 and Pseudocode 6 respectively.

The SLA_SBPDU_Proc procedure is used by a STAR bridge to update its ESL Table and to propagate agent bridge information to its STAR neighbors. Specifically, upon receiving a StationLoc(s) frame < s, k > from tree port p , a STAR bridge self assigns k , the agent bridge of s , to $ESL(self, s)$, and then forwards the StationLoc(s) out of all its tree ports except p .

The $DF_STAR_Learning_Proc$ procedure is used by a STAR bridge to update its FD and ESL Table upon receiving a normal data frame fr . The procedure is also used by the STAR bridge to discover if it is an agent bridge for $src(fr)$, and if so, the STAR bridge forwards a StationLoc($src(fr)$) out of appropriate tree ports.

As specified in Pseudocode 6, when self is the designated bridge but not the agent bridge of the source end station s (Case 6.1), it sends a StationLoc(s) frame on all tree ports and updates the FD entry and ESL Table entry for s. If ab(s) is unknown (Case 6.2), self has to check whether it is ab(s). If $ab(s) \neq db(s)$, then ab(s) is the closest upstream STAR bridge and so ab(s) must receive the normal data frame of s from a child link port (Case 6.2a). The agent bridge sends the StationLoc(s) frame on all tree ports. Therefore, those STAR bridges in higher levels of the spanning tree will receive the StationLoc(s) frame before the normal data frame. They should have an entry for s in their ESL Tables by the time they receive the normal data frame and won't send out another StationLoc(s) frame.

PROCEDURE: SLA_SBPDU_Proc(s, k, p), see also FIG. 9

```

Begin
  ESL(self, s) := k;          /* fill the ESL Table */
  Send StationLoc(s) frame on all tree ports except p
end

```

Pseudocode 5: SLA_SBPDU_Proc

PROCEDURE: DF_STAR_Learning_Proc(fr, p), see also FIG. 18

```

/* fr is a normal data frame */
Begin
  s := src(fr)
  If db(s) = self and ab(s) ≠ self      /* Case 6.1: self is the designated
  bridge                                and current agent bridge is not
                                        self*/
    Send StationLoc(s) <s, self> frame on all tree ports
    f(self, s) := p;
    ESL(self, s) := self
  Else if ab(s) is not found            /* Case 6.2: agent bridge
  not known */                          /* Case 6.2a: p is a child link port
  */
    If p is a child port
      Send StationLoc(s) <s, self> frame on all tree ports
      f(self, s) := p;
      ESL(self, s) := self
    Else if p is a root port
      f(self, s) := p
end

```

Pseudocode 6: DF_STAR_Learning_Proc

V. STAR Forwarding Process

STAR bridges execute the STAR forwarding process after the STAR learning process when a data frame is received. Having received a data frame destined for an end station t, a STAR bridge n first checks its ESL Table to determine if it knows ab(t), the agent bridge of t. If ab(t) is found, n will then find out from its BF Table the forwarding port of ab(t). If no entry for ab(t) is found in the BF Table, the data frame is dropped because it indicates an error in the BF Table. Since the agent bridge is a STAR bridge, the BF Table should have a record showing how to get there. If ab(t) is unknown, n will proceed to check its FD. If end station t is unknown, STAR bridge n will forward the data frame on all tree ports except the incoming one, just as the IEEE 802.1D standard.

V.A Frame Duplication Problem

In the IEEE 802.1D standard, although a bridge may forward the same frame on more than one port, only one port leads to the destination and the designated bridge of the destination since there is a unique path from any source to

any destination on a spanning tree. Therefore, an old bridge can never receive the same data frame more than once. As the STAR bridge graph may not be a tree, two STAR bridges may receive the same data frame and try to forward it to the destination using different paths. For example, suppose that the destination end station is attached to u' and the source end station is attached to x in FIG. 10. If x forwards one copy of the data frame to w and one to v, U' may receive two copies of the same data frame, one from w and one from v. The STAR Bridge Protocol avoids this by allowing only w, the agent bridge of s, to forward the frame using, in this case, an enhanced forwarding path. Since the agent bridge is unique, at most one copy of the frame may be sent to the destination.

V.B Frame Dropping Problem

Another problem is due to the existence of old bridges. Although the STAR bridges know how to forward a frame on an enhanced forwarding path after knowing the location of the destination, old bridges don't. In some cases, old bridges will drop a frame trying to pass through. In FIG. 10, bridges w, x, and v' are on the same branch and there is a crosslink between v' and z. Let z be the designated bridge of end station s1. As x and z are on different branches, in the FD of x, it marks the root port as the forwarding port of s1. However, if w wants to send a frame to s1 and finds out the

shortest path to z is going through v', it sends the frame with destination address s1 to x. The frame will be dropped by x since it is coming from the forwarding direction. To fix this problem, we will encapsulate the normal data frame with an appropriate proxy destination address so that x will forward the frame towards v', but not other directions as in Section I.E. An old bridge may drop a data frame only if the data frame is being forwarded on a tree path. It also implies that the frame is trying to go from one STAR bridge to another through an old bridge tree path between them. If the destination address of a frame is the MAC address of the next hop bridge, all the old bridges along the tree path will forward the frame to the next hop as desired. In the present example, w encapsulates the frame with the MAC address of v' as a proxy destination address, such that x, upon receiving the encapsulated frame, will forward the frame to v' without dropping it. In general, when the next hop STAR bridge is not a tree neighbor, the sender STAR bridge will encapsulate the data frame. Since the encapsulated data frames are encapsulated using the sender STAR bridge MAC address as

the source address, old bridges will learn the forwarding directions to the sender STAR bridge. In this respect, no additional control message is needed to enable old bridges to learn the forwarding direction to any STAR bridge.

A BA Table is used to keep the MAC addresses of all distant STAR neighbors. Since STAR bridge n puts its MAC address as the source address of the DVMyInfo and DVOurInfo frames it sends, we don't need an extra SBPDU frame to fill this table.

V.D Redundant Traversal Problem

The encapsulation approach described in Section V.B may prevent a designated old bridge from identifying a normal data frame that is destined to it. Referring to the configuration in FIG. 10, let end station s be attached to the old bridge x. According to the protocol, STAR bridge w will declare itself to be the agent bridge of s, so that all frames destined for s will be forwarded as though they were destined for w. Suppose that s1 wants to send a frame to s. STAR bridge z, which is db(s1), will forward the frame via a crosslink to v' then. Since the next hop to a STAR bridge is w and it is on a tree path, v' will encapsulate the frame with w as a proxy destination. When old bridge x receives the frame, it will think that it is not a frame addressed to itself and forward it to the proxy destination w. When w receives it, it will strip off the encapsulation header and send the normal data frame back to x. Then, x can identify the frame and send it to the destination end station. Therefore, the data frame traverses a redundant path from the designated bridge to the agent bridge and back to the designated bridge. In the case where the agent bridge of source end station, ab(s), and the agent bridge of the destination end station, ab(t), are on different branches, we will show later in Section VII.C. that the total distance traversed is still no worse than the corresponding tree path. However, when ab(s) and ab(t) are on the same branch, the total distance traversed may be longer than the tree path. We avoid this by not encapsulating the normal data frame in this situation. To let the next hop STAR bridge know whether a frame is intended to be forwarded on an enhanced forwarding path or a tree path, the agent bridge always encapsulates a frame that is going to be forwarded on an enhanced forwarding path.

V.E Procedure for STAR Forwarding Process

To avoid unnecessary frame dropping, STAR bridge n encapsulates a normal data frame as discussed in Section I.E. In the encapsulated data frame, the MAC address of n is used as the source address and the MAC address of the next hop STAR bridge is used as the destination address. Given a data frame fr, if it is a normal data frame, src(fr) and dst(fr) is the source end station address and the destination end station address respectively. If fr is an encapsulated data frame, it must have been encapsulated by a STAR bridge

normal data frame fr where srcbridge and dstbridge are respectively the source and destination addresses associated with the encapsulated frame. We use uncap(fr) to represent the normal data frame that an encapsulated data frame fr is carrying in its payload. Pseudocode 7 presents the FD Search procedure for finding information from the FD. Pseudocode 8 presents the ESL_Search procedure for finding information from the ESL Table. In all the pseudocodes of this section, self is the STAR bridge executing the process, p is the receiving port of the data frame, s is the source end station, t is the destination end station, and pld is the payload portion of a normal data frame. FIG. 19 and FIG. 20 show the flow-charts corresponding to Pseudocode 7 and Pseudocode 8 respectively.

In Pseudocode 8 (ESL_Search_Proc(s, t, pld, p)), when the agent bridge of end station t is not known (Case 8.1). It is an error case and the frame should be dropped. This is an error because DF_STAR_Forwarding_Proc calls ESL_Search_Proc only when self knows that ab(t) is defined (Case 9.2) or ab(t) is known by some STAR bridge that encapsulates the frame (Case 9.1). When self itself is the agent bridge of t (Case 8.2), it sends the frame to the forwarding port leading to t. In Case 8.3, the agent bridge is another STAR bridge. In this case, the BF Table should give the forwarding port and the next hop information. If not, there is error and the frame is dropped.

Pseudocode 9 is a complete new forwarding procedure of bridge self FIG. 22 shows the flow-chart corresponding to Pseudocode 9. When the data frame is encapsulated, it must be sent from another STAR bridge n', which has the information of ab(t) in its ESL Table and BF Table. Therefore, self should search from its ESL Table to forward the frame. On the other hand, when the data frame is not encapsulated, there are several situations. If ab(t) is unknown (Case 9.2a), of course, self should try to look at the FD. Both Case 9.2b and Case 9.2c are the cases where ab(t) is known. In Case 9.2b, self is the agent bridge of the source end station. In this case, when ab(t) and ab(s) are on the same branch, the normal data frame should not be encapsulated as explained in Section V.C. We don't search the ESL Table in Case 9.2c because of the frame duplication issue discussed in Section V.A. Only a tree path can be used to forward the data frame in this case. FIG. 8 and FIG. 21 show the procedures for processing a data frame in accordance with the STAR Bridge Protocol and the IEEE 802.1D Spanning Tree Bridge Protocol respectively.

PROCEDURE: FD_Search_Proc(s, t, pld, p), also see FIG. 19

```

Begin
  if f(self, t) is not found      /* Case 7.1: End station t is unknown */
    send data frame <s, t, pld> on all tree ports except p
  else if f(self, t) = p        /* Case 7.2: Forwarding port is the coming port */
    drop the frame
  else
    send data frame <s, t, pld> on f(self, t)
  endif
end

```

whose MAC address is src(fr), and dst(fr) is the address of the intended STAR bridge recipient. We let encaps(fr, srcbridge, dstbridge) as the encapsulated data frame of a

Pseudocode 7: FD_Search_Proc

PROCEDURE: ESL_Search_Proc(s, t, pld, p), also see FIG. 20

```

Begin
  if ESL(self, t) is not found      /* Case 8.1: error */
    drop the frame
  Else if ESL(self, t) = self      /* Case 8.2: ab(t) = self */
    Send data frame <s,t,pld> on f(self,t)
  Else                             /* Case 8.3: ab(t) is known and ab(t) ≠
self */
    a := ab(t)
    if F(self, a) is not found or next(self, a) is not found /* Case 8.3a: error
*/
      drop the frame
    else                             /* Case 8.3b: encapsulation
necessary */
      send data frame encaps(<s, t, pld>, addr(self), addr(next(self,
a)))
      on F(self, a)
    endif
  endif
end

```

Topology Change Notification BPDUs inform other bridges to recompute the spanning tree. End station information is updated by a timeout mechanism. Each entry in the

Pseudocode 8: ESL_Search_Proc

PROCEDURE: DF_STAR_Forwarding_Proc(fr, p),²⁵ also see FIG.

FD is assigned a timer and the information is forgotten when the timer expires.

```

Begin
  If fr is encapsulated           /* Case 9.1: fr is an encapsulated frame
*/
    s := src(uncap(fr))
    t := dst(uncap(fr))
    pload := pld(uncap(fr))
    if dst(fr) ≠ addr(self)       /* self is not the proxy destination
*/
      drop the frame
    else
      ESL_Search_Proc(s, t, pload)
    endif
  else                             /*Case 9.2: fr is not encapsulated
*/
    s := src(fr)
    t := dst(fr)
    pload := pld(fr)
    if ESL(self, t) is not found   /* Case 9.2a: ab(t) is unknown */
      FD_Search_Proc(s, t, pload, p)
    Else if ESL(self, s) = self    /* Case 9.2b: ab(s) = self */
      If ESL(self, t) = self      /* ab(s) = ab(t) */
        FD_Search_Proc(s, t, pload, p)
      If FG_R(self, ESL(self, t)) = 1 /* ab(t) is an
ancestor of ab(s) */
        FD_Search_Proc(s, t, pload, p)
      Else if FG_R(self, ESL(self, t)) = -1 /* ab(t) is an descendant of
ab(s) */
        FD_Search_Proc(s, t, pload, p)
      else                         /* ab(t) and ab(s) are on different
branches */
        ESL_Search_Proc(s, t, pload)
      Else                         /* Case 9.2c: ab(s) is unknown or
ab(s) ≠ self */
        FD_Search_Proc(s, t, pload, p)
      endif
    endif
  endif
end

```

Pseudocode 9: DF_STAR_Forwarding_Proc

VI. Update

In the IEEE 802.1D standard, the root bridge sends a BPDUs message periodically to update the spanning tree. When a bridge detects a topological change, it sends a

The STAR Bridge Protocol keeps topological information in the BF Table, which is built upon the spanning tree and eligible crosslinks. Therefore, if either the tree changes or any crosslink changes, the BF Table must be recomputed. The procedure for detecting any change of the spanning tree is available in the above-mentioned standard. The BF Table

is recomputed after the spanning tree becomes stable again. The mechanism of detecting crosslink failures is described in Section I.F. A STAR bridge transits back to the Tree Learned state when a crosslink fails. In the mean time, the STAR bridge executes the standard forwarding process and the standard learning process instead of the new ones to forward data frames.

In the STAR Bridge Protocol, the information needed for reaching end stations is kept in the ESL Table and the FD. They both time out in the same way as the FD in the old bridges do. This is necessary because no bridge can detect the relocation of an end station. Since bridge addresses do not change frequently, the BA Table does not need to be timed out.

VII. Performance

In this section, we analyze the storage, message complexity, and path length of the STAR Bridge Protocol.

VII.A Storage

Each old bridge keeps only one table for forwarding, which is the FD. One entry is necessary for each known end station. Therefore, the space required is $O(|M|)$, where M is the set of all end stations in the extended LAN. In addition to an FD, there are three new tables in each STAR bridge: BF Table, ESL Table, and BA Table. Table 5 is a summary of these tables in STAR bridge n .

TABLE 5

Storage Requirements in STAR Bridges		
Name	Content	Space Required
BF Table	$\langle n', F(n, n'), \text{next}(n, n'), \text{FG_R}(n, n') \rangle$, $n' \in B \setminus \{n\}$, $F(n, n') \in P(x)$, $\text{next}(n, n') \in B \setminus \{n\}$	$O(B)$
ESL Table	$\langle s, \text{ab}(s) \rangle$, $s \in M$, $\text{ab}(s) \in B$	$O(M)$
BA Table	$\langle n', \text{addr}(n') \rangle$, $n' \in B \setminus \{n\} \setminus N_B(n)$	$O(B)$
FD	$\langle s, f(n, s) \rangle$, $s \in M$, $f(n, s) \in P_T(x)$	$O(M)$

After the STAR learning process has been executed for some time and old entries in the FD have been timed out, an end station s appears in both the ESL Table and the FD of STAR bridge n only if $\text{ab}(s)=n$. Therefore, the total memory needed for the ESL Table and the FD in STAR bridges together would be about the same as in the old bridges. We do need extra space for the BF Table and the BA Table. However, as the number of entries of both tables is at most $|B|$ which is far less than $|M|$, we can conclude that the storage requirement in a STAR bridge is comparable to that in an old bridge.

VII.B Message Complexity

In the IEEE 802.1D standard, BPDU frames are sent periodically to build and maintain a spanning tree. In the STAR Bridge Protocol, SBPDU frames are introduced and they are described in Section I.D. Hello SBPDUs are sent on

eligible crosslinks only and a Hello SBPDU frame will not propagate beyond the crosslink that it is sent on. Therefore, Hello SBPDUs do not put extra message overhead on tree links. Distance Vector Change Notification SBPDUs are sent only when distance vectors have to be recomputed and they are sent over the spanning tree. As a result, under stable configuration, there will normally be no Distance Vector Change Notification SBPDUs generated. Table 6 summarizes the format of Distance Vector Computation SBPDUs and Station Location Announcement SBPDUs. The path finding process generates Distance Vector Computation SBPDUs and the STAR learning process generates Station Location Announcement SBPDUs.

For each Distance Vector Computation SBPDU frame generated by the path finding process, there is at most one recipient on each port. Obviously, there should be more DVRecord frames than other Distance Vector Computation SBPDU frames in this process. The number of DVRecord frames needed for each pair of STAR bridges depends on the length of the enhanced forwarding path between them. The path length is bounded by the diameter of the tree. The number of messages generated by the spanning tree is related to the diameter of the tree too. Therefore, we can conclude that the number of messages generated by the path finding process is about $|B|$ times the number of the messages needed in building the spanning tree. The path finding process will not generate any DVRecord frame after building the BF Table. Nevertheless, the root bridge will keep on generating BPDU messages periodically after the spanning tree has been built. Therefore, for a stable bridged LAN, the extra number of messages generated by the path finding process is negligible.

TABLE 6

Summary of SBPDU Frames used in the STAR processes		
SBPDU Frame	Recipients responsible for processing the frame	Number of copies for each recipient
DVMyInfo	$n' \in N'_B(n)$	1
DVOurInfo	$n' \in N'_B(n)$ s.t. $\text{nca}(n, n') = n$	1
DVInform	$n' \in N'_B(n)$ s.t. $\text{nca}(n, n') = n$	$ \{k \mid k \in \{N'_B(n') \text{ s.t. } k \text{ and } n' \text{ are on different branches}\} $
	$n' \in N'_B(n)$ s.t. $\text{nca}(n, n') \neq n$	1 when the accurate distance is calculated by n but not n' ; 0 otherwise
DVRecord	$n' \in N'_B(n) \cup N_B(n)$	At most $ B $ for each $k \in B \setminus \{n, n'\}$
StationLoc	$n' \in B \setminus \{n\}$	$ \{m \mid m \in M \text{ s.t. } \text{ab}(m) = n\} $

Location information is necessary in all algorithms in the Section entitled Background Information that are applicable for any additive metric. In those algorithms, every bridge has to know the location of all end stations. In the STAR Bridge Protocol, a STAR bridge keeps only the location of end station s provided $\text{ab}(s)$ is defined. Therefore, the location messages generated by the STAR Bridge Protocol are less than those generated by the algorithms in the prior art.

VII.C Path Length

In this section, we would like to show that the length of a STAR forwarding path is always less than or equal to the corresponding tree path. In the following discussion, we denote the length of the STAR forwarding path between two bridges x and y as $\text{len}(x, y)$. In all figures referring to this section, a black node represents a STAR bridge, a white

node represents an old bridge, and a dot-dash line represents a tree path. We first establish the following lemmas.

Lemma 1:

When a STAR bridge n receives a normal data frame originated by an end station s , the STAR bridge may encapsulate the frame only if $n=ab(s)$.

Proof of Lemma 1:

We observe that encapsulation is executed only in the `ESL_Search_Proc` procedure (Pseudocode 8) and the `DF_STAR_Forwarding_Proc` procedure (Pseudocode 9) in which the `ESL_Search_Proc` procedure is invoked. In `DF_STAR_Forwarding_Proc` procedure, when n receives a normal data frame, the `ESL_Search_Proc` procedure is invoked by n only when $n=ab(s)$ (Case 9.2b). Note that if $ab(s)$ is not defined, n can never be $ab(s)$.

Lemma 2:

A normal data frame originated by an end station s will always be forwarded over a tree path if the frame is not encapsulated by $ab(s)$.

Proof of Lemma 2:

According to Case 8.3b in the `ESL_Search_Proc` procedure, the normal data frame is never forwarded over a crosslink, unless it is encapsulated. By Lemma 1, if the normal data frame is not encapsulated by $ab(s)$, it will remain as a normal data frame for the rest of its forwarding journey. According to Case 9.2a and Case 9.2c of the `DF_STAR_Forwarding_Proc` procedure, having received the normal data frame, a STAR bridge n that is not $ab(s)$ may forward the frame only to tree neighbors. Having received the normal data frame, an old bridge may forward the frame only to tree neighbors. Therefore, the normal data frame will be forwarded over a tree path if the frame is not encapsulated by $ab(s)$.

Lemma 3:

If a frame is forwarded from an end station s to another end station t over an enhanced forwarding path, the path must traverse at least one crosslink, and the frame must be encapsulated by $ab(s)$, which is necessarily defined.

Proof of Lemma 3:

If s and t are on the same branch, then the tree path from s to t is necessarily a shortest path. Given that s and t are on different branches, the forwarding may traverse two or more different branches. In the first case, the forwarding path must traverse a crosslink, or it would not be an enhanced forwarding path. In the second case, it must traverse at least one crosslink because the forwarding path cannot be a normal tree path. According to Case 8.3b in the `ESL_Search_Proc` procedure, the frame will be encapsulated when it is forwarded over any crosslink. By Lemma 2, the frame must be encapsulated by $ab(s)$.

Lemma 4:

If a STAR forwarding path, along which a frame is sent from an end station s to another end station t , is an enhanced forwarding path, then $ab(s)$ must be defined and is the first STAR bridge on the enhanced forwarding path.

Proof of Lemma 4:

By Lemma 3, the frame must be encapsulated by $ab(s)$, which is necessarily defined. Since $ab(s)$ is by definition the first STAR bridge on the root path of $db(s)$. If $ab(s)=db(s)$, then the proof is complete. Otherwise, $db(s)$ is an old bridge, and will send a copy of the frame along the root path of $db(s)$ regardless of its knowledge of the end station t . This copy of the frame will be received $ab(s)$, which will encapsulate the frame so that it will be forwarded over the enhanced forwarding path. Duplicate copies of the frame will be dropped in accordance with the protocol.

Lemma 5:

If a STAR forwarding path, along which a frame is sent from an end station s to another end station t , is an enhanced forwarding path, then $ab(t)$ must be defined and is the last STAR bridge on the enhanced forwarding path.

Proof of Lemma 5:

By Lemma 3, the forwarding path must traverse at least one crosslink. According to Case 8.3b in the `ESL_Search_Proc` procedure, the frame will be encapsulated when it is forwarded over each crosslink. $ab(t)$ must be defined because, according to Case 9.2a of `DF_STAR_Forwarding_Proc` procedure, a frame will not be encapsulated otherwise. If $ab(t)=db(t)$, the proof is complete. Otherwise, $ab(t)$ must be an ancestor of $db(t)$ by definition, and $db(t)$ must receive the frame without encapsulation. It suffices to show that $ab(t)$ is on the enhanced forwarding path and all intermediate STAR bridges on the enhanced forwarding path, except $ab(t)$, will forward the frame with encapsulation. $ab(t)$ is on the enhanced forwarding path because, according to Case 8.3b of the `ESL_Search_Proc` procedure, each intermediate STAR bridge forwards the encapsulated frame through its forwarding port leading to $ab(t)$. According to Case 8.2 and Case 8.3 in the `ESL_Search_Proc` procedure, a STAR bridge will forward the frame without encapsulation only if it is $ab(t)$.

Lemma 6:

When $ab(s)$ and $ab(t)$ are both defined, and they are on different branches of the spanning tree, $len(db(s), db(t)) \leq d_T(db(s), db(t))$.

Proof of Lemma 6:

FIGS. 23–26 respectively show various exemplary scenarios for this lemma. Since $ab(s)$ and $ab(t)$ are on different branches, s and t must be on different branches. If the forwarding path from $db(s)$ to $db(t)$ is a tree path (FIG. 23), then the proof is complete. Otherwise, the forwarding path is an enhanced forwarding path. By Lemma 4, $ab(s)$ is the first STAR bridge on the enhanced forwarding path. By Lemma 5, $ab(t)$ is the last STAR bridge on the enhanced forwarding path. Therefore, the enhanced forwarding path consists of three disjoint segments. The first segment, which is a tree path from $db(s)$ to $ab(s)$, has a path length $d_T(db(s), ab(s))$. The second segment, which is an enhanced forwarding path from $ab(s)$ to $ab(t)$, has a path length $d(ab(s), ab(t))$. The third segment, which is a tree path from $ab(t)$ to $db(t)$, has a path length $d_T(ab(t), db(t))$. Therefore, $len(db(s), db(t))$, the length of the enhanced forwarding path from $db(s)$ to $db(t)$, satisfies the following inequality.

$$\begin{aligned} len(db(s), db(t)) &= d_T(db(s), ab(s)) + d \\ &(ab(s), ab(t)) + d_T(ab(t), db(t)) \leq d_T(db(s), ab(s)) + d_T(ab(s), ab \\ &(t)) + d_T(ab(t), db(t)) \\ &= d_T(db(s), db(t)) \end{aligned}$$

Lemma 7:

A STAR forwarding path for a frame sent by an end station s to another end station t is a tree path if at least one of $ab(s)$ and $ab(t)$ is not defined.

Proof of Lemma 7:

FIGS. 26–29 respectively show various exemplary scenarios for this lemma. By Lemma 2, if $ab(s)$ is not defined, the STAR forwarding path must be a tree path because the frame will never be encapsulated. According to Case 9.2a of `DF_STAR_Forwarding_Proc` procedure, a frame will not be encapsulated unless $ab(t)$ is defined.

Lemma 8:

When $ab(s)$ and $ab(t)$ are both defined, and $ab(s)=ab(t)$, then $len(db(s), db(t)) < d_T(db(s), db(t))$.

Proof of Lemma 8:

FIGS. 30–32 show various exemplary scenarios for this lemma. When $ab(s)=ab(t)$, there are two cases: $ab(s)=db(s)$ and $ab(s) \neq db(s)$. In the first case, according to Case 8.2 of the `ESL_Search_Proc` procedure, $ab(s)$ will send the frame according to the tree path, and the proof is complete. In the

second case, if $db(t)$ is an ancestor of $db(s)$, when the normal data frame travels upstream from $db(s)$, it will reach $db(t)$ before $ab(s)$ and so $ab(s)$ won't receive that. In the second case, if $db(t)$ is a descendant of $db(s)$, by Lemma 8, the STAR forwarding path must be a tree path, and the proof is complete.

Lemma 9:

When $ab(s)$ and $ab(t)$ are both defined, $ab(s) \neq ab(t)$, and $ab(s)$ is an ancestor of $ab(t)$, then $len(db(s), db(t)) \leq d_r(db(s), db(t))$.

Proof of Lemma 9:

FIG. 33 shows an exemplary scenario for this lemma. When $ab(s)$ is an ancestor of $ab(t)$, there are two cases: $ab(s) = db(s)$ and $ab(s) \neq db(s)$. In the first case, $ab(s)$ may encapsulate the normal data frame and send it downstream along the tree path. When $ab(t)$ receives an encapsulated data frame, it will strip off the header and trailer, and send the normal data frame according to the information in the FD. Therefore, the STAR forwarding path between $db(s)$ and $db(t)$ is a tree path. The proof is complete. When $ab(s) \neq db(s)$, $ab(s)$ must be an ancestor of $db(s)$ and $db(s)$ must be an ancestor of $db(t)$.

Lemma 10:

When $ab(s)$ and $ab(t)$ are both defined, $ab(s) \neq ab(t)$, and $ab(s)$ is a descendant of $ab(t)$, then $len(db(s), db(t)) \leq d_r(db(s), db(t))$.

Proof of Lemma 10:

FIG. 34 shows an exemplary scenario for this lemma. We have discussed the issue of the scenario in which $ab(t)$ is an ancestor of $ab(s)$ in Section V.D. When $ab(s)$ finds out that $ab(t)$ is an ancestor, it sends out the normal data frame on its root port. Since only $ab(s)$ is allowed to encapsulate the normal data frame, all STAR bridges on the tree path between $ab(s)$ and $ab(t)$ will send the normal data frame on tree ports only. Then, the normal data frame will reach $db(t)$ along the tree path without any redundant traversal and the proof is complete. All duplicate frames sent over the spanning tree will eventually be dropped.

We now present a theorem on the path length of the STAR bridge protocol.

Theorem 1:

A STAR forwarding path for a frame sent by an end station s to another end station t leads from $db(s)$ to $db(t)$, and has a path length $len(db(s), db(t))$ that satisfies the following inequality.

$$len(db(s), db(t)) \leq d_r(db(s), db(t))$$

Proof of Theorem 1:

We first observe that any STAR forwarding path is either a tree path or an enhanced forwarding path. If it is a tree path, the proof is complete. Otherwise, we will show that the inequality still holds.

We divide all situations into the following scenarios:

1. $ab(s)$ is not defined
2. $ab(s)$ is defined
 - 2.1. $ab(t)$ is not defined
 - 2.2. $ab(t)$ is defined
 - 2.2.1. $ab(s)$ and $ab(t)$ are on different branches
 - 2.2.2. $ab(s)$ and $ab(t)$ are on the same branch

By Lemma 7, the STAR forwarding path is always a tree path except for the cases in scenario 2.2. By Lemma 6, the inequality holds for scenario 2.2.1. By Lemma 9 through Lemma 11, the inequality holds for scenario 2.2.2.

Star Bridges With Prioritized Path Differentiation Capabilities

The STAR bridge protocol as described above in detail, makes use of processes specified in the IEEE 802.1D standard for supporting the three basic functions of the spanning tree bridge protocol. In addition, three new processes are further specified: path finding process, STAR learning process, and STAR forwarding process, also described above in detail. The path finding process allows a STAR bridge to find and estimate the distance of a path from itself to another STAR bridge. The STAR forwarding process and the STAR learning process are modified versions of the forwarding process and the learning process respectively specified in the standard. All STAR bridges can execute both the standard and the new processes.

In the IEEE 802.1D standard, a rooted spanning tree is built before the forwarding and learning processes start. Similarly, a rooted spanning tree is found by standard and STAR bridges together before the execution of the new processes. After the rooted spanning tree is found, the path finding process is started. Before the path finding process ends, STAR bridges and standard bridges execute the standard forwarding process and the standard learning process to forward data frames on tree paths. When the path finding process is completed, each STAR bridge begins to execute the STAR learning process and the STAR forwarding process instead of the standard ones to forward data frames on identified enhanced forwarding paths.

In the IEEE 802.1D standard, each bridge keeps a forwarding database (FD) for the forwarding process. In STAR bridges, three tables are additionally used: bridge forwarding table (BF Table), end-station location table (ESL Table), and bridge address table (BA Table). A BF Table indicates for each other STAR bridge in a bridged LAN the port that leads to the next hop along the "best" path found. BF Tables are obtained in the path finding process. An ESL Table is used to map an end station to a STAR bridge near it. The STAR learning process is responsible for filling this table. Therefore, if the ESL Table of a STAR bridge has a record for an end station, it is unnecessary for the FD of the STAR bridge to have a record for the same end station in most situations. The BA Table provides mapping between STAR bridge identifiers and STAR bridge MAC addresses. Every bridge has its own unique MAC address and this MAC address is used in the STAR forwarding process.

In the IEEE 802.1D standard, each LAN has a designated bridge and this bridge is also the designated bridge of all end stations attached to that LAN. In accordance with the STAR bridge protocol, each end station can be mapped to its "closest" STAR bridge along the standard tree forwarding path to the root bridge on the spanning tree, provided such a STAR bridge exists. Said "closest" STAR bridge to an end station is referred to as an agent bridge of the end station. Once a STAR bridge identified itself to be an agent bridge of a formerly unknown end station, it informs other STAR bridges so that they know where to forward a frame destined for that end station. Specifically, enhanced forwarding paths between agent bridges may be identified. In the event that no STAR bridge is located upstream of the designated bridge of an end station, the agent bridge for said end station is undefined. In this case, the tree path is used for forwarding frames originated from said end station.

The inventive elements that extend the STAR Bridge Protocol in order to support multi-level traffic prioritization and load balancing in an extended LAN by means of path differentiation for traffic of a plurality of priorities are described below.

Three new fields are added in the Bridge Forwarding (BF) Table. The new fields in the BF Table of a first STAR bridge include, for each other STAR bridge, (1) the tree forwarding port leading from the first STAR bridge to the other STAR bridge, and (2) the next hop STAR bridge neighbor on the tree path from the first STAR bridge to the other STAR bridge, and (3) the number of STAR bridges on the tree path from the first STAR bridge to the other STAR bridge. The first two new fields are needed because the present invention requires each STAR bridge to maintain both a tree path and an enhanced forwarding path, if available, to each other STAR bridge, whereas the original STAR Bridge Protocol requires each STAR bridge to maintain only one of a tree path and an enhanced forwarding path, if available, to each other STAR bridge. The third new field is needed because the present invention makes use of the information contained in this field to implement prioritized frame forwarding.

Table 7 is a summary of an entry in the BF Table of a STAR bridge n in accordance with the present invention. The entry as shown contains forwarding information associated with another STAR bridge n'. Rows (c), (d) and (e) are the new fields which are utilized for prioritizing.

TABLE 7

Fields in an entry associated with STAR bridge n in the Bridge Forwarding Table of STAR bridge n		
Field	Definition	
(a) F(n, n')	Forwarding port for n' on the enhanced forwarding path	
(b) next(n, n')	ID of the next hop STAR bridge neighbor on the enhanced forwarding path from n to n'	
(c) f(n, n')	Forwarding port for n' on the tree path	
(d) next_t(n, n')	ID of the next hop STAR bridge neighbor on the tree path from n to n'	
(e) S_hop_count(n, n')	Number of STAR bridges on the tree path from n and n'	
(f) FG_R(n, n')	Relation flag with a value 1 if n' is an ancestor of n, -1 if n' is a descendant of n, and 0 otherwise	

To obtain the information needed to fill in the above new fields in the BF Table, every STAR bridge broadcasts an Announcement frame over the tree to other STAR bridges in

the beginning of the path finding process, as shown in FIG. 36. This Announcement frame carries the identity of the STAR bridge that originates the frame as well as a value of a counter, whose initial value is set to zero. This counter keeps track of the STAR bridge hop count. When a first STAR bridge receives an Announcement frame from a second STAR bridge from a child port, the first STAR bridge learns that the second STAR bridge is a STAR bridge on the downstream and the number of STAR bridge hops from the first STAR bridge to the second STAR bridge is given by the value of the counter in the Announcement frame. The first STAR bridge then increments the counter value by 1 and forwards the modified Announcement frame over all tree ports except the incoming one. The Announcement frame also carries the identity of the last STAR bridge that receives the frame. Thus, a STAR bridge knows the next STAR tree hop leading to the STAR bridge that originates the Announcement frame.

Consider an Announcement frame, ANNOUNCEMENT <n, n', c>, which is originated by a first STAR bridge (denoted n) and forwarded by a second STAR bridge (denoted n'), wherein the count value of the counter is c. The following is the pseudocode to be executed by a third STAR

bridge (denoted self) after it receives the Announcement frame over a forwarding port p of the third STAR bridge: PROCEDURE: Ann_SBPDU_Proc(n, n', c, p), see FIG. 36

```

Begin
  f(self, n) := p;      /* Assigning p to forwarding port from self to n
on the tree path */
  next_t(self, n) := n'; /* Assigning n' to be the ID of next hop
STAR bridge neighbor on
the tree path from n to self */
  S_hop_count(self, n) := c; /* Assigning c to be the number of STAR
bridges on the tree path
from n to self */
  Send ANNOUNCEMENT <n, self, c+1> on all tree ports except p
End
    
```

The present invention does not modify the distance vector computation in the path finding process of the original STAR bridge protocol, and the process identifies enhanced forwarding paths to other STAR bridges as in the original STAR protocol described above in detail. In accordance with the present invention, the STAR learning process works in

the same manner as in the STAR bridge protocol to identify agent bridges of end stations.

In the original STAR Forwarding Process, for a frame sent from a source end station to a destination end station, the agent bridge of the source end station may decide whether that data frame should be sent over the enhanced forwarding path or the tree path. If the agent bridge decides to send the frame on the tree path, an intermediate STAR bridge on the tree path cannot forward the frame over the enhanced forwarding path leading from itself to the agent bridge of the destination end station. In other words, there are at most two alternate paths that may be used from the source end station to the destination end station, i.e., one is the tree path and the other is the enhanced forwarding path identified. The agent bridge of the source end station decides which one of them to be used for forwarding the frame. Intermediate STAR bridges have to follow that decision and forward the frame accordingly. To avoid out-of-order delivery of frames, the agent bridge must forward frames belonging to a flow consistently either over a tree path or an enhanced forwarding path, but not both. This consistency is ensured by forwarding only frames of a selected class over the enhanced forwarding path, provided it is available.

In the present invention, the agent bridge of a source end station, and other intermediate STAR bridges along a tree path, decide only the next forwarding hop, instead of an entire forwarding path from said agent bridge to the agent bridge of the destination end station. The next hop of a tree path is referred to as the next tree hop and the next hop of an enhanced forwarding path as the next enhanced path hop. Both next tree hop and next enhanced path hop are STAR bridge. Note that if the tree path is the shortest path identified, the next tree hop is the same as the next enhanced path hop. If a STAR bridge decides to send a frame over the next enhanced path hop, the next enhanced path hop must forward the frame over its own next enhanced path hop leading to the agent bridge of the destination end station. On the other hand, if a STAR bridge decides to send a frame over the next tree hop, when the frame arrives at the next tree hop, that STAR bridge can decide on its own whether to forward it on its next tree hop or its next enhanced path hop. That is, the forwarding path of a frame from the agent bridge of a source end station to the agent bridge of a destination end station will be either a tree path, an enhanced forwarding path leading from the agent bridge of the source end station, or a path that consists of a tree path from the agent bridge of the source end station to an ancestor STAR bridge of said agent bridge, and an enhanced forwarding path leading from the ancestor STAR bridge to the agent bridge of the destination end station. The higher the priority of a frame, the smaller the number of hops on the tree path it traverses. For a frame of the highest priority, the agent bridge will send it over an enhanced forwarding path right away so that the frame will not traverse any more tree hops. For a frame of the next lower priority, the agent bridge might send it over its next tree hop so that an intermediate STAR bridge on the tree path may forward the frame over an enhanced forwarding path leading from the intermediate STAR bridge. A frame of successively lower priority may traverse more tree hops before it is forwarded on an enhanced forwarding path. Finally, a frame of the lowest priority is forwarded only over the tree path, as no bridge will forward it over any enhanced path hop.

If the designated bridges of the source end station and the destination end station of a frame are on the same branch, the tree path is the shortest path and there is no path differentiation. The number of different paths that can be used for forwarding a frame from the designated bridge of the source end station to the designated bridge of the destination end station, where the designated bridges are on different branches, depends on the number of STAR bridges on the

tree path from the designated bridge of the source end station to the nearest common ancestor of the designated bridges of the source and destination end stations. It is because only the STAR bridges on the tree path from the designated bridge of the source end station to the said nearest common ancestor (said STAR bridges include the agent bridge of the source end station but exclude the said nearest common ancestor) can be the first STAR bridge that decides to forward the frame on a next enhanced path hop. When a frame, which has been forwarded over tree hops, arrives at the said nearest common ancestor, as the said nearest common ancestor and the agent bridge of the destination end station are on the same branch and there is no enhanced forwarding path, the said nearest common ancestor will forward the frame downstream to the agent bridge of the destination end station and the frame will follow a tree path to the agent bridge of the destination end station afterwards.

In the original STAR forwarding process, procedure `ESL_Search_Proc` is responsible for forwarding frames on enhanced forwarding paths. In accordance with the present invention, this procedure is extended so that intermediate STAR bridges can determine whether to forward a frame on a next tree hop or a next enhanced path hop. In the original `ESL_Search_Proc`, if a frame has been forwarded over an enhanced forwarding path by an agent bridge, each intermediate STAR bridge would forward the frame on an enhanced forwarding path whenever it is available. In accordance with the present invention, a STAR bridge uses the priority of the frame and the number of STAR tree hops that the frame has traversed to determine which port to forward the frame. Specifically, the present invention introduces a new function in the `ESL_Search_Proc` so that it takes in two (2) integer parameters, `priority_value` and `S_hop_count`, and returns an integer for decision making. `priority_value` represents the priority value of a frame. Priority value is a non-negative integer and a smaller priority value implies a higher priority. Priority value 0 is used to indicate the highest priority. `S_hop_count` is the number of STAR tree hops a frame has traversed. The integer returned by the function is then used for deciding which port to forward a frame. We refer to this function as the forwarding decision function and the returned value of the function as the forwarding decision value. If the forwarding decision value is greater than zero, we forward the frame on the next tree hop, otherwise, we forward the frame on the next enhanced path hop. In order to ensure that the path of a higher priority frame is no "worse" than that of a lower priority frame, the forwarding decision function should return a smaller value when the priority is higher. On the other hand, the forwarding decision function should return a small value when the STAR hop count is large since the longer the tree path a frame has traversed, it is desirable that the more likely that the frame is going to be forwarded on a next enhanced path hop. Generally speaking, as long as a function satisfies the following conditions, it can be used as the forwarding decision function in this protocol:

- 1) It accepts an integer `priority_value` and an integer `S_hop_count` as arguments
- 2) Given the same `priority_value`, the forwarding decision value of a larger `S_hop_count` should not be larger than the forwarding decision value of a smaller `S_hop_count`
- 3) Given the same `S_hop_count`, the forwarding decision value of a larger `priority_value` should not be smaller than the forwarding decision value of a smaller `priority_value`

Different functions yield different forwarding paths for frames of the same priority. For example, consider the function `priority_value-S_hop_count`. By using this function, an agent bridge will forward frames of priority value 0 to the next enhanced path hop and send frames of other priorities over the next tree hop. In this case, the number of STAR bridges along the tree path a frame has to go through is the same as its priority value. In another example, consider the function `floor(priority_value/4)-S`

station (denoted `t`) with a certain payload (denoted `pld`). The notations used in the pseudocode shown below follow those defined in the original `ESL_Search_Proc` of the STAR bridge protocol described above in detail. The forwarding decision function is denoted as `m()` in the pseudocode. The priority value of a frame is carried in the payload `pld` and `pr(pld)` returns the priority value.

PROCEDURE: `ESL_Search_Proc(s, t, pld)`, FIG. 38

```

Begin
if ESL (self, t) is not found                               /* error */
    Drop the frame
Else if ESL (self, t) = self                                 /*ab(t) = self*/
    Send data frame <s, t, pld> on f(self, t)
Else                                                         /* ab(t) is known and ab(t) ≠ self
*/
    ab_s := ESL (self, s)                                     /*Locate ab(s)*/
    ab_t := ESL (self, t)                                     /*Locate ab(t)*/
    if ab_s = self or FG_R(self, ab_s) = -1                 /* self is ab(s) or an
ancestor of ab(s) */
        if m(pr(pld), S_hop_count (self, ab_s)) ≤ 0       /* send over next
enhanced path hop */
            next_hop := next_t(self, ab_t)
            forwarding_port := F(self, ab_t)
        else
            next_hop := next(self, ab_t)                     /* send over next tree hop */
            forwarding_port := f(self, ab_t)
        endif
    else
        next_hop := next(self, ab_t)                         /* self is not ab(s) and, self is on
another branch or
is a descendant of ab(s) */
        forwarding_port := F(self, ab_t)
    endif
    send data frame encap(<s, t, pld>, addr(self), addr(next_hop)) on
forwarding_port
endif
end

```

`hop_count`, where `floor(x)` represents the largest integral value that does not exceed `x`. In this case, frames of priority values 0 to 3 will be forwarded by their agent enhanced path hops and frames of other priority values will have to travel one tree hop. Table 8 shows the forwarding decision values of different `priority_value` and `S_hop_count` when the forwarding decision function is `priority_value-S_hop_count`. Areas marked with an "*" represent the cases where a frame will be forwarded over a next enhanced path hop.

TABLE 8

	forwarding decision function example (priority_value-S_hop_count)								
	S_hop_count								
	0	1	2	3	4	5	6	7	
Prior-ity_value	0	*0	*-1	*-2	*-3	*-4	*-5	*-6	*-7
	1	1	*0	*-1	*-2	*-3	*-4	*-5	*-6
	2	2	1	*0	*-1	*-2	*-3	*-4	*-5
	3	3	2	1	*0	*-1	*-2	*-3	*-4
	4	4	3	2	1	*0	*-1	*-2	*-3
	5	5	4	3	2	1	*0	*-1	*-2
	6	6	5	4	3	2	1	*0	*-1
	7	7	6	5	4	3	2	1	*0

The following procedure is executed by a STAR bridge (denoted `self`) when it is invoked by its `DF_STAR_Forwarding_Proc` procedure upon receiving a frame that is sent from an end station (denoted `s`) to a destination end

FIG. 35 is an example of a bridged LAN useful for explaining the advantages of the present invention. Open circles (nodes) represent standard bridges while solid circles (nodes) are STAR bridges. Solid lines are the tree lines and dotted lines are crosslinks. End station `s` attaches to STAR bridge `z` and end station `t` attaches to STAR bridge `v'`. Let bridge hop count be the metric being considered. Then, the tree path from STAR bridge `z` to STAR bridge `v'` is of length 5 and the length of the enhanced forwarding path from STAR bridge `z` to STAR bridge `v'` is 1. Let the forwarding decision function be `priority_value-S_hop_count`. Suppose now end station `s` wants to send a frame of priority value 0 to a known end station `t`. STAR bridge `z` is the agent bridge of end station `s`. When STAR bridge `z` receives the frame, as the agent bridge of end station `t` in another branch, it executes `ESL_Search_Proc`. It evaluates function `m`, see FIG. 38, and the function returns 0. Therefore, it encapsulates the frame and sends it to the next enhanced path hop, which is STAR bridge `v'`, the agent bridge of end station `t`. The length of the path is 1 hop and it is an enhanced forwarding path.

Now, suppose that end station `s` sends another frame with priority value 1. Function `m` returns 1 and so STAR bridge `z` encapsulates the frame and sends it to STAR bridge `q`, the next tree hop. When STAR bridge `q` receives the frame, it evaluates function `m` and `m` returns 0. STAR bridge `q` sends the frame to the next enhanced path hop where STAR bridge `x` is located. When STAR bridge `x` receives the frame, since the agent bridge of end station `t` is now on its downstream, it forwards the frame to STAR bridge `v'` through the tree hop. The path that this frame traverses is a hybrid path `z→q→x→v'` and is of length 3, which is longer than the path

of the frame with priority value 0. For a frame of priority value 2 or higher, STAR bridge q will forward it to STAR bridge u, the next tree hop, and the frame will travel to STAR bridge v' using the tree path. The length of the tree path is 5 hops and is longer than the paths used by frames of lower priorities.

FIG. 37 is a flow chart showing the manner in which received frames are processed and the appropriate program is selected and run according to the contents of the received frame. Each process of FIG. 37 is shown in detail in the process Figures previously described.

VIII. CONCLUSION

Many current networking products support the IEEE 802.1D standard, which is widely deployed in today's LAN internetworking environment, for example, DOCSIS-compliant cable data products. DOCSIS specifies the use of the IEEE 802.1D bridge protocol to interconnect cable modem CMTSes over a switched or bridged headend network. In DOCSIS, bridging CMTSes and cable modems intended for commercial use must support the IEEE 802.1D standard. Bridges are also used in CableHome-compliant networks, which are of particular interest for home networking products.

When multiple home network segments are deployed in a home, bridges can be used in the home network environment for interconnecting home network segments to form a transparent extended LAN. It is necessary for home bridges to be installed correctly by those with appropriate skills, or the resulting topology may contain loops, which are known to be detrimental to the performance of the extended bridged LAN. With anticipated proliferation of home devices with a wide spectrum of capabilities, and a plethora of LAN technologies, future home networks must scale not only with bandwidth, but also with connectivity and reliability. To these ends, it is necessary for the topology of the extended bridged LAN in the home to support alternate paths. In other words, a mesh topology will be implemented by design. The IEEE 802.1D bridge protocol can be used not only for ensuring loop-free frame forwarding in the home network, but also for robust network extensions.

The STAR Bridge Protocol, which incorporates frame forwarding over alternate paths that are shorter than their corresponding tree paths and is backward compatible with the IEEE 802.1D standard, can be enabled to support two levels of traffic prioritization. An extension of the STAR Bridge Protocol to support more than two levels of traffic prioritization and improve its load balancing capability enhances the value of the STAR Bridge Protocol in home network products.

We have disclosed a new STAR Bridge Protocol with prioritized path differentiation and load balancing capabilities that is backward compatible with the standard IEEE 802.1D Spanning Tree Bridge Protocol. The protocol described herein selects an enhanced path and tree paths of

varying lengths responsive to a priority number carried by a frame. Incidentally, the standard protocol places a restriction on the maximum bridge diameter, which is the maximum number of bridges between any two points of attachments of end stations. The IEEE 802.1D specification recommends a maximum bridge diameter of 7. Given a set of LANs and bridges, there may not exist any single RST that could be built by the standard protocol to satisfy such a restriction. The STAR Bridge Protocol, on the other hand, may still satisfy the restriction because enhanced forwarding paths may reduce the bridge diameter. With the use of enhanced forwarding paths, frames that would have passed through the root bridge in the standard spanning tree may be diverted over alternate paths such that the load at the root bridge is likely to be alleviated. Providing path differentiation further enhances load balancing.

Being backward compatible with the IEEE 802.1D standard, the proposed bridge protocol offers a smooth migration path to QoS-based bridging. Eventually, all old bridges will be phased out or replaced by STAR bridges. When there is no old bridge in the bridged network, distance vector routing can be used to find all-pair shortest paths. Since distance vectors are always forwarded one hop away, there will be no distant STAR neighbor when all bridges are STAR. Note that the IEEE 802.1D spanning tree is still needed.

APPENDIX

List of Acronyms

BA	Bridge Address
BF	Bridge Forwarding
BPDU	Bridge Protocol Data Unit
CM	Cable Modem
CMTS	Cable Modem Termination System
DF	Data Frame
DHCP	Dynamic Host Configuration Protocol
DLS	Distributed Load Sharing
DOCSIS	Data-Over-Cable Service Interface Specifications
DV	Distance Vector
DVC	Distance Vector Computation
DVCN	Distance Vector Change Notification
ESL	End Station Location
FD	Forwarding Database
GDLS	Generalized Distributed Load Sharing
ID	Identifier
IP	Internet Protocol
LAN	Local Area Network
LLC	Logical Link Control
MAC	Medium Access Control
QoS	Quality of Service
RST	Rooted Spanning Tree
SBPDU	STAR Bridge Protocol Data Unit
SLA	Station Location Announcement
STAR	Spanning Tree Alternate Routing

TABLE A

Notation		
Notation	Definition	Remarks
V	Set of bridges representing all bridges	
B	Set of bridges representing STAR bridges	$B \subseteq V$
N(x)	Set of all direct neighbors of bridge x	$N(x) \subseteq V \setminus \{x\}$
N _B (x)	Set of direct STAR neighbors of x	$N_B(x) \subseteq N(x), N_B(x) \subseteq B$
N' _B (x)	Set of distant STAR neighbors of x	$N'_B(x) \subseteq B$
P(x)	Set of all ports of bridge x	
P _T (x)	Set of tree ports of bridge x	$P_T(x) \subseteq P(x)$
p _r (x)	Root port of bridge x	
p(x, y)	Port of bridge x leading to neighbor bridge y	$p(x, y) \subseteq P(x)$

TABLE A-continued

<u>Notation</u>		
Notation	Definition	Remarks
$c(x, y)$	Weight of link between bridge x and bridge y	$c(x, y) = 0$ if $x = y$ $c(x, y) > 0$ if $x \neq y$
$d_r(x)$	Root path distance of bridge x for x for $x \in V \setminus \{r\}$	Root bridge r is given
$treepath(x, y)$	Tree path from x to y	$x \in V, y \in V$
$d_r(x, y)$	Tree path distance between bridge x and bridge y	$d_r(x, y) = 0$ if $x = y$ $d_r(x, y) > 0$ if $x \neq y$
$d(x, y)$	Current estimated distance from x to y , for $x \in B$ and $y \in B \setminus \{x\}$	$d(x, y) = 0$ if $x = y$ $d(x, y) > 0$ if $x \neq y$
$F(x, y)$	Forwarding port from bridge x to bridge y for $x \in B$ and $y \in B \setminus \{x\}$	$F(x, y) \in P(x)$
$f(x, s)$	Forwarding port from bridge x to end station s for $x \in V$ and $s \in M$	$f(x, s) \in P(x)$
M	Set of all end stations	
$ab(s)$	Agent bridge of end station s	$s \in M, ab(s) \in B$
$db(s)$	Designated bridge of end station s	$db(s) \in V$
$S(x)$	Set of end stations in the FD of $x, x \in V$	$S(x) \subseteq M$
$H(x)$	Set of end stations in ESL Table of $x, x \in B$	$H(x) \subseteq M$
$nca(x, y)$	Nearest common ancestor of bridge x and bridge y	$nca(x, y) \in V$
$path(x, y)$	The forwarding path, from x to y , used by the proposed protocol	$x \in V, y \in V$
$len(x, y)$	Length of path(x, y)	$len(x, y) = 0$ if $x = y$ $len(x, y) > 0$ if $x \neq y$
$dsan(n)$	Distant STAR ancestor neighbor of bridge $n \in B$	$dsan(n) \in B$
$dsanc(n)$	Child of distant STAR ancestor neighbor of bridge $n \in B$	$dsanc(n) \in V$
$dcc(n)$	Doubly counted cost between $dsan(n) \in B$ and $dsanc(n) \in V$	$dcc(n) > 0$
$cb(k, n)$	Child of bridge $k \in B$ on a tree path leading from k to $n \in B$	$cb(k, n) \in V$

While the principles of the invention have been described above in connection with specific apparatus, it is to be clearly understood that this description is made only by way of example and not as a limitation on the scope of the invention.

What is claimed is:

1. Apparatus for determining paths for forwarding frames among end stations in a system of interconnected local area networks, comprised of:

first and second groups of frame forwarding devices comprising bridges, which are used to interconnect local area networks;

said first and second groups of bridges each having first means for creating tree paths for loop-free frame forwarding to an end station;

said second group of bridges each having second means cooperating with said first group of bridges and all other bridges of said second group of bridges to create an enhanced path for loop-free forwarding of a frame to an end station, wherein said enhanced path is an alternate path that is shorter than any tree path to said end station; and

said second group of bridges each having third means for selecting one of a tree path and an alternate path responsive to at least one of a parameter contained in a received frame and a number of hops a frame has made over a tree path between an agent bridge and the bridge of the second group of bridges receiving the frame.

2. The apparatus of claim 1 wherein said parameter value is a priority value identifying a priority assigned to a frame.

3. The apparatus of claim 2 wherein said third means includes means for selecting a path whereby a higher priority of a frame, a smaller number of hops over a tree path a frame makes before being forwarded over an enhanced path.

4. The apparatus of claim 2 wherein said third means includes means for selecting a path whereby a greater number of hops made by a frame over a tree path, a greater is the likelihood that the frame will be forwarded over an enhanced path.

5. The apparatus of claim 1 wherein said second group of bridges are backward compatible with said first group of bridges, wherein

the operation of said first group of bridges is unaltered by the presence of said second group of bridges, and said second group of bridges inter-operate with said first group of bridges when at least one bridge of the latter group is present in the system, and said second group of bridges can perform their functions in the absence of said first group of bridges in order to operate.

6. The apparatus of claim 1 wherein said first group of bridges include means for performing a spanning tree bridge protocol (STBP) to determine a loop-free tree path for forwarding a frame to an end station;

said second group of bridges include means for performing a spanning tree alternate routing bridge protocol (STAR BP) for determining said alternate path; and said second group of bridges include means for determining said number of hops the frame has made over the tree path.

7. The apparatus of claim 6 wherein said means for performing said spanning tree alternate routing bridge protocol includes means for performing said spanning tree bridge protocol.

8. The apparatus of claim 1 wherein said end stations are distributed among said local area networks that are interconnected; and

said first and second groups of bridges selectively forward frames from end stations in one of said local area

networks to end stations in another one of said local area networks.

9. The apparatus of claim 1 wherein each bridge of the second group of bridges further comprises means for selecting a tree path when said means for selecting receives a frame originated by a source bridge and having a priority value falling within a first range of priority values including a lowest priority.

10. The apparatus of claim 9 wherein said first range of priority values decreases with said number of hops.

11. The apparatus of claim 1 wherein selected ones of bridges of said second group are arranged along different tree paths and are joined by cross-links;

said means for determining said alternate paths including means for determining if a path including one of said cross-links provides a path shorter than a tree path; and said selecting means selecting said one of said cross-links responsive to receipt of a frame having a given priority value and said number of hops the frame has made over the tree path.

12. The apparatus of claim 1 wherein each bridge of said first and second groups include:

means for assigning a weighted metric to each port of the bridge;

means for exchanging said weighted metrics with other bridges in the system;

means for electing a root bridge responsive to the exchanged information;

means for determining a distance between each bridge and said root bridge;

means for building a unique spanning tree rooted at the root bridge in a distributed manner;

means for learning and remembering which port to forward a frame over said spanning tree to a given end station; and

means for storing said data.

13. The apparatus of claim 1 wherein each bridge of said second group of bridges includes means for sending an announcement frame to every other bridge of said second group of bridges over said tree path for identifying a tree forwarding port, a next hop bridge of said second group of bridges and a number of bridges of said second group between a source and the bridge sending the announcement.

14. The apparatus of claim 12 wherein each bridge of said second group of bridges is further provided with means for conveying to other bridges of said second group information representing topology known to the conveying bridge;

means for conveying to other bridges of the second group information identifying bridges of the second group which are a source for frames and a destination for frames;

means for conveying to other bridges of the second group topology information known to bridges of the second group other than the conveying bridge; and

means for storing information received from other bridges of said second group for determining, if possible, an alternate path shorter than a corresponding tree path.

15. The apparatus of claim 1 wherein each bridge of said second group is further provided with means for determining if another bridge of said second group is directly connected thereto by a cross-link which is a link that is not used to form any tree path;

said means for determining an alternate path further including means for examining valid cross-links for use as segments of an alternate path, wherein a valid cross-link is one whose end terminals are connected to bridges of said second group; and

one whose end terminals are joined to bridges of said second group wherein a tree path of one of the bridges of said second group connected to one end terminal of a cross-link is not a segment of a tree path of the other bridge of the second group connected to another end terminal said cross-link, and wherein the tree path of the other bridge of the second group is not a segment of the tree path of said one bridge of said second group.

16. The apparatus of claim 1 wherein each bridge of said second group of bridges further includes;

means for storing an end station location (ESL) table associating each end station with bridges in the second group near each end station;

means for storing a bridge address (BA) table associating each bridge of said second group with a medium access control (MAC) address;

means for storing a bridge forwarding (BF) table for indicating a path of the associated bridge that leads to a next hop along a best path found for forwarding a frame; and

means responsive to a received frame and information stored in said BA, ESL and BF tables for selecting a path to an end station identified in destination information incorporated as part of said received frame.

17. Apparatus for determining a path for forwarding a frame received at a bridge in a system comprised of a plurality of interconnected local area networks (LANs) each having a plurality of end stations, and a spanning tree incorporating a plurality of bridges of first and second groups for loop-free forwarding a frame from a source end station to a destination end station wherein said source and destination end stations may reside in different LANs, said bridges of said first and second groups having means for determining a tree path for loop-free forwarding of said frame; and

said second group of bridges further having means for determining if an alternate path exists for loop-free forwarding of said frame which has at least one non-tree path segment; and

means for utilizing one of said alternate path and tree path responsive to a priority number carried by said received frame topological criteria including: a shorter physical path; a less costly path; a path having less delay and a smaller number of hops between source and destination.

18. The apparatus of claim 17 including means for determining said alternate path is according to a group of topological criteria including:

a shorter path;

a less costly path;

a path having less delay; and

a smaller number of hops between a source and a destination.

19. A method for determining a path for forwarding a frame in a system having a plurality of end stations interconnected by first and second types of bridges which create a loop-free spanning tree for transferring frames between end stations, comprising:

a) storing, at each of the second type of bridges, data representing cross-links between bridges of said second type and a number of hops between each second type bridge and every other second type bridge;

b) transferring, at a second type of bridge, a received frame over a first path having a smallest number of hops to a destination end station responsive to a priority number of a given value and a destination end station carried by the received frame.

20. The method of claim 19 further comprising:

- c) transferring, at a second type of bridge, a received frame over a path having a greater number of hops than said first path responsive to a priority number of lesser priority and a destination end station carried by the received frame.

21. A method for prioritizing the forwarding of frames among end stations in a system of local area networks interconnected by a first and second types of bridges, comprising:

- a) creating a loop-free tree path for the transfer of frames employing the first and second types of bridges;
- b) forming cross-links between bridges of said second type to provide enhanced forwarding paths;
- c) storing tree path information at said first and second types of bridges;
- d) storing information of the cross-links at each of the second types of bridges; and
- e) forwarding a frame received by a second type of bridge along a path having a smallest number of hops of said second type bridges between a second type bridge receiving the frame and an end station designated to receive the frame and carried by said frame, responsive to a priority number carried by said frame, said path being determined based on the data stored at steps (c) and (d).

22. The method of claim 21 further comprising forwarding a frame received by said second type bridge along a path having a larger number of hops of the second type bridge between the said second type bridge receiving the frame and an end station designated to receive the frame according to data carried by said frame, responsive to a priority number of lower importance than said first mentioned priority number, employing the data stored at steps (c) and (d).

23. A method for establishing a protocol for prioritizing forwarding of frames among end stations in a system of local area networks interconnected by first and second types of bridges, comprising:

- a) every second type of bridge broadcasting an announcement frame over a loop-free tree path to other bridges of said second type during a path finding process, each announcement frame carrying an identity of the second type of bridge that originates the frame and a count value, which represents a number of bridges of the second type traversed by the frame; and
- b) at least another one of said second type of bridges receiving an announcement from another one of said second type of bridges at a given port thereof incrementing the count value by one, and forwarding a modified announcement frame over all tree ports except said given port.

24. A method of claim 23 further comprising providing in an announcement frame an identity of a second type of bridge that last transferred the announcement frame enabling a second type of bridge to know the next tree hop leading to a second type of bridge that originates the announcement frame.

25. The method of claim 23 wherein each of the second type of bridges decides only the next forwarding hop in place of an entire forwarding path from an agent bridge of the source end station to an agent bridge of the destination end station.

26. The method of claim 23 wherein a second type of bridge receiving a frame from another second type of bridge decides only the next forwarding hop in place of an entire forwarding path from an agent bridge of the source end station to an agent bridge of the destination end station, independently of a decision made by the said another second type of bridge transferring said frame.

27. The method of claim 23 wherein a second type of a bridge receives a frame that is being forwarded downstream on a tree path, said second type of bridge selecting a tree path thereafter.

28. The method of claim 23 wherein a second type of a bridge receives a frame that is being forwarded upstream on a tree path, said second type of bridge selecting either an enhanced path, or a tree path responsive to a priority number carried by the received frame and a number of hops the frame has made over the tree path.

29. The method of claim 23 wherein a second type of a bridge receives a frame on a tree an enhanced path, said second type of bridge selecting either an enhanced path, or a tree path responsive to an availability of an enhanced path thereafter.

30. A method for determining paths for forwarding frames among end stations in a system of interconnected local area networks, comprising first and second groups of frame forwarding devices comprising bridges, for interconnecting local area networks, comprising:

- a) said first and second groups of bridges determining tree paths for loop-free frame forwarding to an end station;
- b) said second group of bridges determining an alternate path for loop-free forwarding of a frame to an end station, which alternate path is shorter than corresponding tree path to said end station; and
- c) said second group of bridges selecting one of a tree path and an alternate path responsive to at least one of a parameter contained in a received frame and a number of hops a frame has made over a tree path between an agent bridge and the bridge of the second group of bridges receiving the frame.

31. The method of claim 30 comprising providing a parameter value which is a priority value identifying a priority assigned to a frame.

32. The method of claim 31 wherein said second group of bridges select a path such that a higher priority of a frame, a smaller a number of hops over a tree path a frame makes before being forwarded over an enhanced path.

33. The method of claim 31 wherein said second groups of devices select a path such that a greater a number of hops made by a frame over a tree path, a greater is the likelihood that the frame will be forwarded over an enhanced path.