13. The Multiple Spanning Tree Protocol (MSTP)

The MSTP algorithm and protocol provides simple and full connectivity for frames assigned to any given VLAN throughout a Bridged Local Area Network comprising arbitrarily interconnected Bridges, each operating MSTP, STP (Clause 8 of IEEE Std 802.1D, 1998 Edition), or RSTP (Clause 17 of IEEE Std 802.1D, ~~1998~~ 2004 Edition). MSTP allows frames assigned to different VLANs to follow separate paths, each based on an independent Multiple Spanning Tree Instance (MSTI), within Multiple Spanning Tree (MST) Regions composed of LANs and or MST Bridges. These Regions and the other Bridges and LANs are connected into a single Common Spanning Tree (CST).

~~NOTE 1—The specification of RSTP is contained in IEEE Std 802.1w-2001, which is an amendment to IEEE Std 802.1D, 1998 Edition. References to IEEE Std 802.1D, 1998 Edition in this standard assume that this amendment, along with IEEE Std 802.1t-2001 (which documents technical and editorial corrections to IEEE Std 802.1D, 1998 Edition), have been applied to the base text of IEEE Std 802.1D, 1998 Edition.~~

MSTP connects all Bridges and LANs with a single Common and Internal Spanning Tree (CIST). The CIST supports the automatic determination of each MST Region, choosing its maximum possible extent. The connectivity calculated for the CIST provides the CST for interconnecting these Regions, and an Internal Spanning Tree (IST) within each Region. MSTP ensures that frames with a given VID are assigned to one and only one of the MSTIs or the IST within the Region, that the assignment is consistent amongst all the Bridges within the region, and that the stable connectivity of each MSTI and the IST at the boundary of the Region matches that of the CST. The stable active topology of the Bridged Local Area Network with respect to frames consistently classified as belonging to any given VLAN thus simply and fully connects all LANs and Bridges throughout the network, though frames belonging to different VLANs can take different paths within any MST Region.

NOTE 1—Readers of this specification are urged to begin by familiarizing themselves with the referenced specification of RSTP.

NOTE 2—Although the active topology determined by STP, RSTP, and MSTP fully connects the components of a Bridged Local Area Network, filtering (GVRP etc.) can restrict frames to a subset of the active topology where some VLANs are not present throughout.

13.1 Protocol Design Requirements

The Spanning Tree Algorithm and its associated Bridge Protocol operate in Bridged Local Area Networks of arbitrary physical topology comprising MSTP, RSTP, or STP Bridges connecting shared media or point to point LANs, so as to support, preserve, and maintain the quality of the MAC Service in all its aspects as specified by Clause 7. In order to do this the algorithm meets the following requirements:

a) It will configure the active topology into a single spanning tree for any given VLAN, such that there is at most one data route between any two end stations for frames consistently allocated to a given VID by Bridges conforming to this standard, eliminating data loops.
b) It will provide for fault tolerance by automatic reconfiguration of the spanning tree topology as a result of Bridge failure or a breakdown in a data path, within the confines of the available Bridged Local Area Network components, and for the automatic accommodation of any Bridge or Bridge Port added to the network without the formation of transient data loops.
c) The active topology will, with a high probability, stabilize within a short, known bounded interval in order to minimize the time for which the service is unavailable for communication between any pair of end stations.
d) The active topology will be predictable and reproducible, and may be selected by management of the parameters of the algorithm, thus allowing the application of Configuration Management, following traffic analysis, to meet the goals of Performance Management.

e) It will operate transparently to the end stations, such that they are unaware of their attachment to a single LAN or a Bridged LAN when using the MAC Service.
f) The communications bandwidth consumed by the Bridges in establishing and maintaining a spanning tree on any particular LAN will be a small percentage of the total available bandwidth and independent of the total traffic supported by the Bridged LAN regardless of the total number of Bridges or LANs.
g) It allows frames assigned to different VLANs to follow different data routes within administratively established regions of the network.
h) It will, with a high probability, continue to provide simple and full connectivity for frames even in the presence of administrative errors in the allocation of VLANs to spanning trees.

Additionally, the algorithm and protocol meet the following goals, which limit the complexity of Bridges and their configuration:

i) The memory requirements associated with each Bridge Port are independent of the number of Bridges and LANs in the Bridged LAN.
j) Bridges do not have to be individually configured before being added to the Bridged LAN, other than having their MAC Addresses assigned through normal procedures.

13.2 Protocol Support Requirements

MSTP does not require any additional configuration mechanisms beyond those specified in 17.1 of IEEE Std 802.1D, ~~1998~~ 2004 Edition in order to support the MAC Service. However to realize the improved throughput and associated frame loss and transit delay performance improvements made possible by the use of multiple spanning trees the following are required:

a) A means of consistently assigning VIDs to MSTIDs within each potential MST Region.
b) Administrative agreement on the Configuration Name and Revision Level used to represent the assignments of VIDs to MSTIDs.
c) A means of assessing the probable distribution of traffic between sets of communicating end stations.
d) A choice of performance goals or the establishment of goals that the quality of service characteristics of some set of communications shall be independent of some other sets.
e) Choices of configuration parameters for the spanning trees that support these goals given the available physical components.

13.3 MSTP Overview

The Multiple Spanning Tree Protocol specifies:

a) ~~An~~ A MST Configuration Identifier (13.7). This allows each Bridge to advertise its configuration for allocating frames with given VIDs to any of a number of distinct, fully and simply connected Multiple Spanning Tree Instances (MSTIs).
b) A priority vector (13.9) that comprises bridge identifier and path cost information for constructing a deterministic and manageable single spanning tree active topology, the CIST, that:
   1) Fully and simply connects all the Bridges and LANs in a Bridged Local Area Network.
   2) Permits the construction and identification of Regions of Bridges and LANs that are guaranteed fully connected by the Bridges and LANs within each Region.
   3) Ensures that paths within each Region are always preferred to paths outside the Region.
c) ~~An~~ A MSTI priority vector (13.9), comprising information for constructing a deterministic and independently manageable active topology for any given MSTI within each Region.
d) Comparisons and calculations performed by each Bridge in support of the distributed spanning tree algorithm (13.10). These select a CIST priority vector for each Port, based on the priority vectors

and MST Configuration Identifiers received from other Bridges and on an incremental Path Cost associated with each receiving Port. The resulting priority vectors are such that in a stable network:

1) One Bridge is selected to be the CIST Root of the Bridged Local Area Network as a whole.
2) A minimum cost path to the CIST Root is selected for each Bridge and LAN, thus preventing loops while ensuring full connectivity.
3) The one Bridge in each Region whose minimum cost path to the Root is not through another Bridge using the same MST Configuration Identifier is identified as its Region's CIST Regional Root.
4) Conversely, each Bridge whose minimum cost path to the Root is through a Bridge using the same MST Configuration Identifier is identified as being in the same MST Region as that Bridge.

e) Priority vector comparisons and calculations performed by each Bridge for each MSTI (13.11). In a stable network:
   1) One Bridge is independently selected for each MSTI to be the MSTI Regional Root.
   2) A minimum cost path to the MSTI Regional Root that lies wholly within the Region is selected for each Bridge and LAN.
f) CIST Port Roles (13.12) that identify the role in the CIST active topology played by each Port on a Bridge.
   1) The Root Port provides the minimum cost path from the Bridge to the CIST Root (if the Bridge is not the CIST Root) through the Regional Root (if the Bridge is not a Regional Root).
   2) A Designated Port provides the least cost path from the attached LAN through the Bridge to the CIST Root.
   3) Alternate or Backup Ports provide connectivity if other Bridges, Bridge Ports, or LANs fail or are removed.
g) MSTI Port Roles (13.12) that identify the role played by each Port on a Bridge for each MSTI's active topology within and at the boundaries of a Region.
   1) The Root Port provides the minimum cost path from the Bridge to the MSTI Regional Root (if the Bridge is not the Regional Root for this MSTI).
   2) A Designated Port provides the least cost path from the attached LAN though the Bridge to the Regional Root.
   3) A Master Port provides connectivity from the Region to a CIST Root that lies outside the Region. The Bridge Port that is the CIST Root Port for the CIST Regional Root is the Master Port for all MSTIs.
   4) Alternate or Backup Ports provide connectivity if other Bridges, Bridge Ports, or LANs fail or are removed.
h) State machines and state variables associated with each spanning tree (CIST or MSTI), port, and port role, to select and change the Port State (8.4, 13.19 of this standard, 17.~~5~~ 10 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) that controls the processing and forwarding of frames allocated to that tree by a MAC Relay Entity (7.3).

13.3.1 Example Topologies

The examples shown in this clause make use of the diagrammatic conventions shown in Figure 13-2.

Figure 13-3 is an example Bridged Local Area Network, chosen to illustrate MSTP calculations rather than as an example of a common or desirable physical topology. Figure 13-6 is the same network showing Bridges and LANs with better CIST spanning tree priorities higher on the page, and including CIST priority vectors, port roles, and MST Regions. In this example:

a) Bridge 0.42 has been chosen as the CIST Root because it has the best (numerically the lowest) Bridge Identifier of all the bridges in the network.
b) Bridges 0.57 and 2.83 are in the same MST Region (1) as 0.42, because they have the same MST Configuration Identifier as the latter. Because they are in the same MST Region as the CIST Root, their External Root Path Cost is 0, and their CIST Regional Root is the CIST Root.

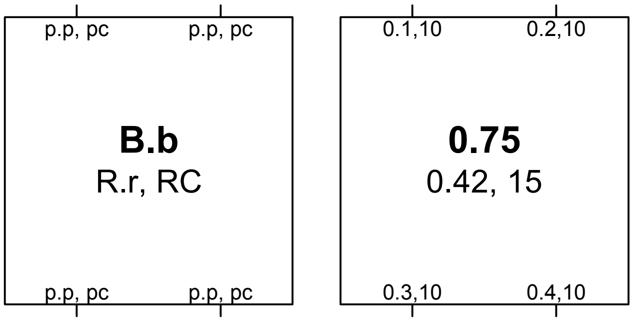

A template for and example of an STP Bridge. **B.b** is the Bridge Identifier (including the manageable priority component **B.**). **R.r** and **RC** are the Root Identifier, Root Path Cost, and the Designated Bridge Identifier, for the Root Port. **p.p, pc** are the Port Identifier (with manageable priority **p.**) and the Port Path Cost for a Bridge Port.

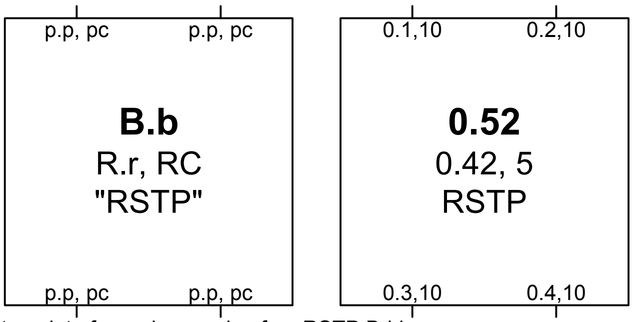

A template for and example of an RSTP Bridge.

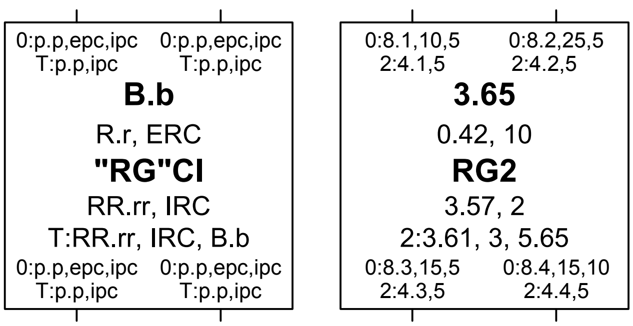

A template for and an example of an MSTP Bridge. **B.b** is the CIST Bridge Identifier. **R.r**, **ERC**, **RR.rr** are the CIST Root Identifier, External Root Path Cost, and Regional Root Identifier. **CI** identifies the Configuration Identifier for the Bridge. **RR.rr, IRC** the CIST Regional Root Identifier and the Internal Root Path Cost. **T:RR.rr, IRC, B.b**, is the Regional Root Identifier, Internal Root Path Cost IRC, and Bridge Identifier for the MSTI with MSTID **T**.
-**p.p, epc, ipc** are the CIST Port Identifier, External Port Path Cost, and Internal Port Path Cost for a Bridge Port. **T:p.p, ipc** are the Port Identifiers and their Regional Costs for MSTI **T**.
Any of the above information may be selectively omitted if deemed irrelevant for the purposes of a diagram.

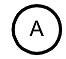

A LAN

Connections between Bridges and LANs indicate the Port Role and Port State by means of their end point symbols, and in some examples, may show the transmission of BPDUs from a Port onto a LAN by means of arrowheads, as shown in the following table.

Port Role	Port State	Legend
Designated	Discarding Learning Forwarding	
& operEdge	Forwarding	
Root Port or Master Port	Discarding Learning Forwarding	
Alternate	Discarding Learning Forwarding	
Backup	Discarding Learning Forwarding	
Disabled	-	
Transmitted Bpdus		
Designated Designated Proposal Root Root Agreement		

NOTE—These diagrammatic conventions allow the representation of Alternate and Backup Ports that are in Learning or Forwarding states; this can happen as a transitory condition due to implementation-dependent delays in switching off Learning and/or Forwarding on a Port that changes role from Designated or Root to Alternate or Backup.

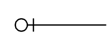

Figure 13-1—Diagrammatic conventions

c) LANs A, B, C, and D are in Region 1 because a Region 1 MST Bridge is the CIST Designated Bridge for those LANs and there are no attached STP Bridges. LAN E is not in a Region (or is in a

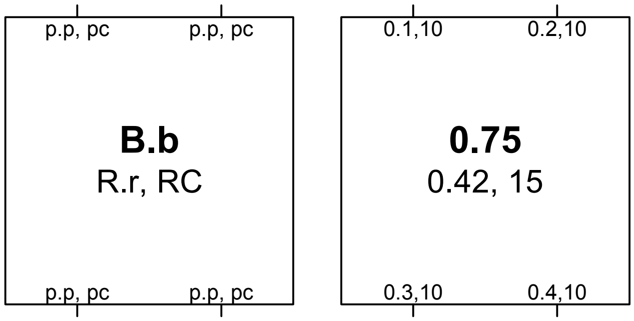

A template for and example of an STP Bridge. **B.b** is the Bridge Identifier (including the manageable priority component **B.**). **R.r** and **RC** are the Root Identifier, Root Path Cost, and the Designated Bridge Identifier, for the Root Port. **p.p, pc** are the Port Identifier (with manageable priority **p.**) and the Port Path Cost for a Bridge Port.

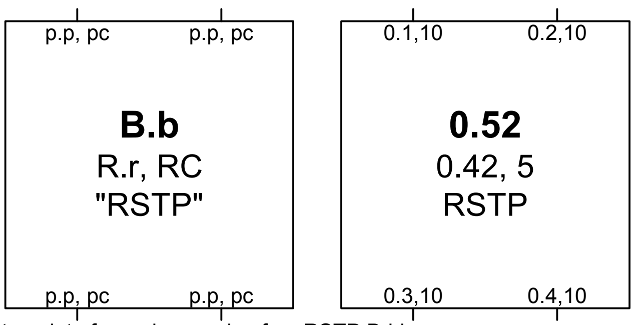

A template for and example of an RSTP Bridge.

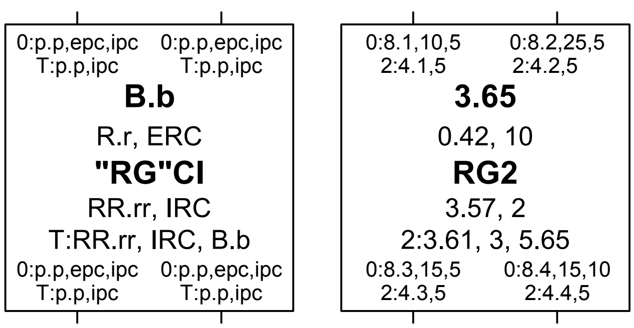

A template for and an example of an MSTP Bridge. **B.b** is the CIST Bridge Identifier. **R.r**, **ERC**, **RR.rr** are the CIST Root Identifier, External Root Path Cost, and Regional Root Identifier. **CI** identifies the Configuration Identifier for the Bridge. **RR.rr, IRC** the CIST Regional Root Identifier and the Internal Root Path Cost. **T:RR.rr, IRC, B.b**, is the Regional Root Identifier, Internal Root Path Cost IRC, and Bridge Identifier for the MSTI with MSTID **T**.
-**p.p, epc, ipc** are the CIST Port Identifier, External Port Path Cost, and Internal Port Path Cost for a Bridge Port. **T:p.p, ipc** are the Port Identifiers and their Regional Costs for MSTI **T**.
Any of the above information may be selectively omitted if deemed irrelevant for the purposes of a diagram.

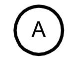

A LAN

Connections between Bridges and LANs indicate the Port Role and Port State by means of their end point symbols, and in some examples, may show the transmission of BPDUs from a Port onto a LAN by means of arrowheads, as shown in the following table.

Port Role	Port State	Legend
Designated	Discarding Learning Forwarding	
& operEdge	Forwarding	
Root Port or Master Port	Discarding Learning Forwarding	
Alternate	Discarding Learning Forwarding	
Backup	Discarding Learning Forwarding	
Disabled	-	
Transmitted Bpdus		
Designated Designated Proposal Root Root Agreement		

NOTE—These diagrammatic conventions allow the representation of Alternate and Backup Ports that are in Learning or Forwarding states; this can happen as a transitory condition due to implementation-dependent delays in switching off Learning and/or Forwarding on a Port that changes role from Designated or Root to Alternate or Backup.

Figure 13-2—Diagrammatic conventions

Region by itself, which is an equivalent view) because it is attached to Bridge 0.53, which is not ~~an~~ a MST Bridge.

d) Bridges 0.77, 0.65, 0.97, 0.86, 3.84, and 3.72 are in the same MST Region (2) because they all have the same MST Configuration Identifier, and are all interconnected by LANs for which one of them is the CIST Designated Bridge.
e) Bridge 0.86 is the CIST Regional Root for Region 2 because it is has the lowest External Root Path Cost through a Boundary Port.
f) LAN N is in Region 2 because its CIST Designated Bridge is in Region 2. Frames assigned to different MSTIDs may reach N from Bridge 0.86 (for example) by either Bridge 0.65 or Bridge 3.72, even though Bridges 0.94 and 0.69 with MST Configuration Identifiers that differ from those for the Bridges in Region 2 are attached to this shared LAN.
g) Bridges 0.94 and 0.69 are in different Regions, even though they have the same MST Configuration Identifier, because the LAN that connects them (N) is in a different Region.

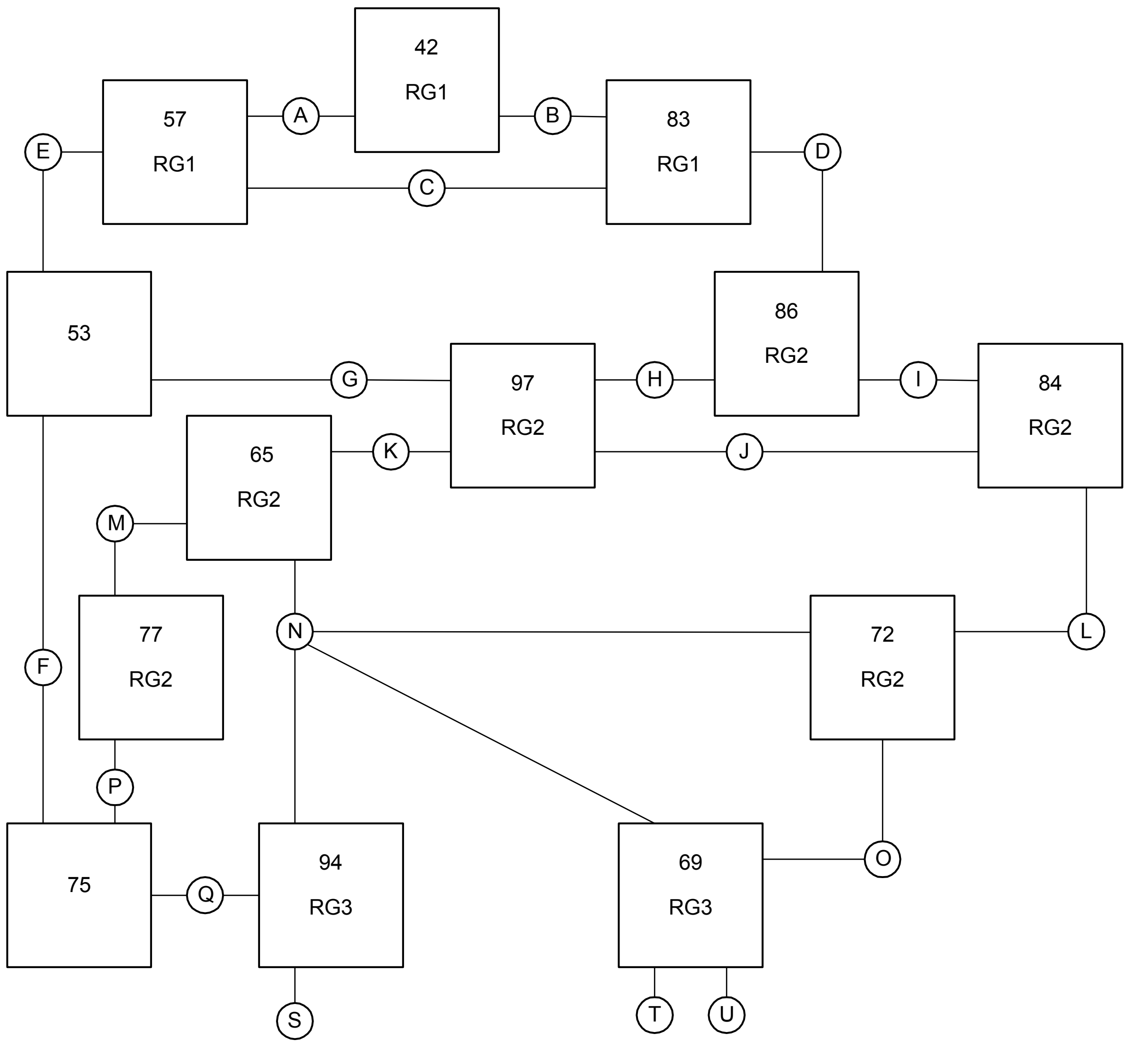

Figure 13-3—An Example Network

Figure 13-8 shows a possible active topology of MSTI 2 within Region 2.

h) Bridge 0.65 has been chosen as the MSTI Regional Root because it has the best (numerically the lowest) Bridge Identifier of all the bridges in the Region for this MSTI.
i) The connectivity between the whole of Region 2 and Region 1 is provided through a single Bridge Port, the Master Port on Bridge 0.86. This port was selected for this role because it is the CIST Root Port on the CIST Regional Root for the Region (see Figure 13-3).

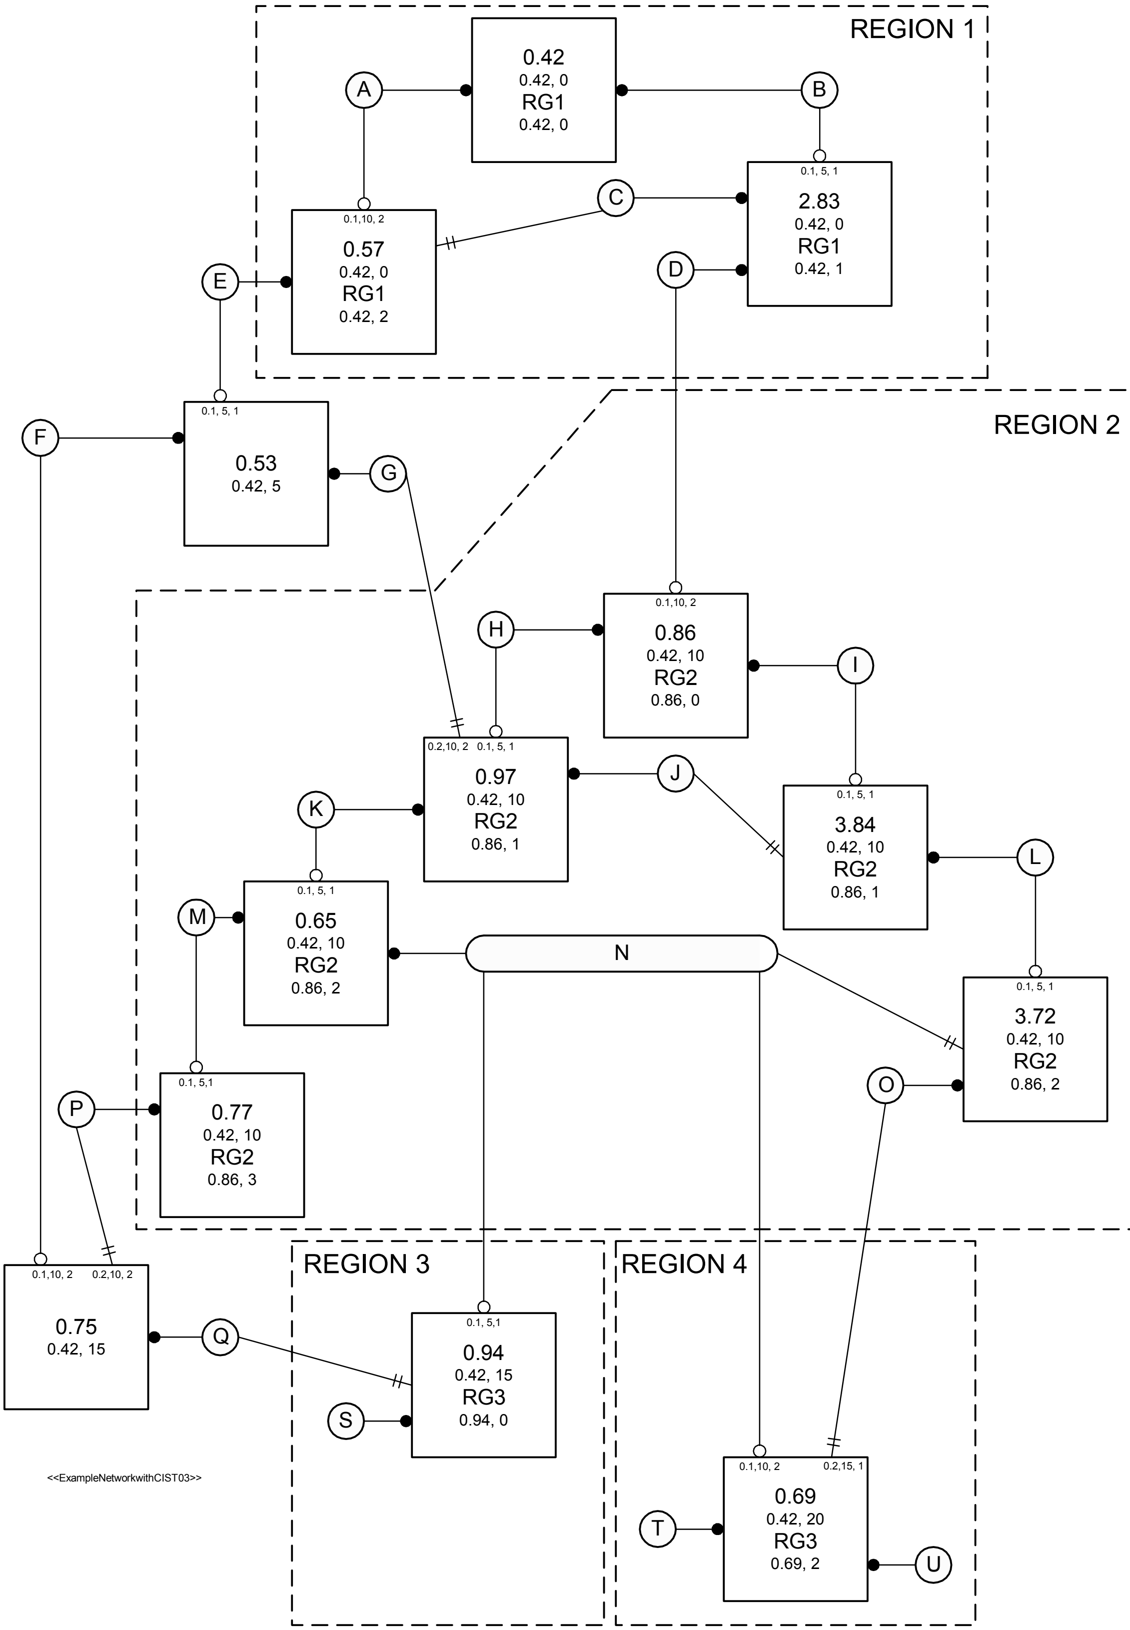

Figure 13-4—Example Network with CIST Priority Vectors, Port Roles, and MST Regions

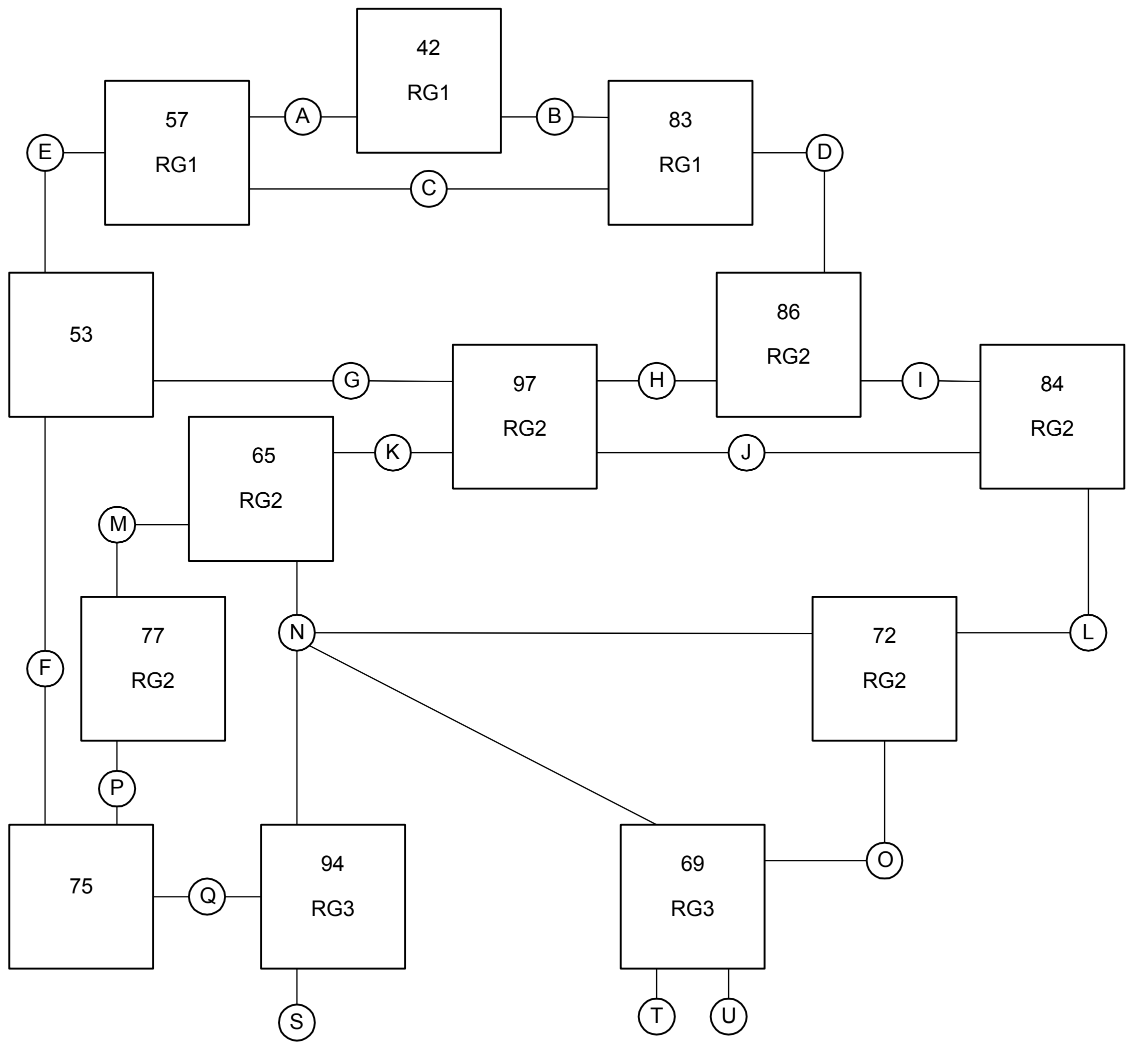

Figure 13-5—An Example Network

j) The connectivity between the whole of Region 2 and LANs and Bridges outside the Region for the MSTI is the same as that for the CIST. This connectivity is similar to that which might result by replacing the entire Region by a single SST Bridge. The Region has a single Root Port (this port is the Master Port for each MSTI) and a number of Designated Ports.

<< Editor's note:

a) The Internal Root Path Cost for MSTID 2 on MST bridge 3.72 (bottom right hand corner of Figure 13-8) indicates 3: it should be 2 instead as connectivity to the MSTI Regional Root is provided by port 0.2 w/ a cost of 2?

b) Why is the Alternate Port on bridge 97 on Figure 13-6 moved to bridge 86 on Figure 13-14?>>

13.4 Relationship of MSTP to RSTP

The design of the Multiple Spanning Tree Protocol is based on that of the Rapid Spanning Tree Protocol (Clause 17 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) extended to provide the capability for frames assigned to different VLANs to be transmitted along different paths within MST Regions.

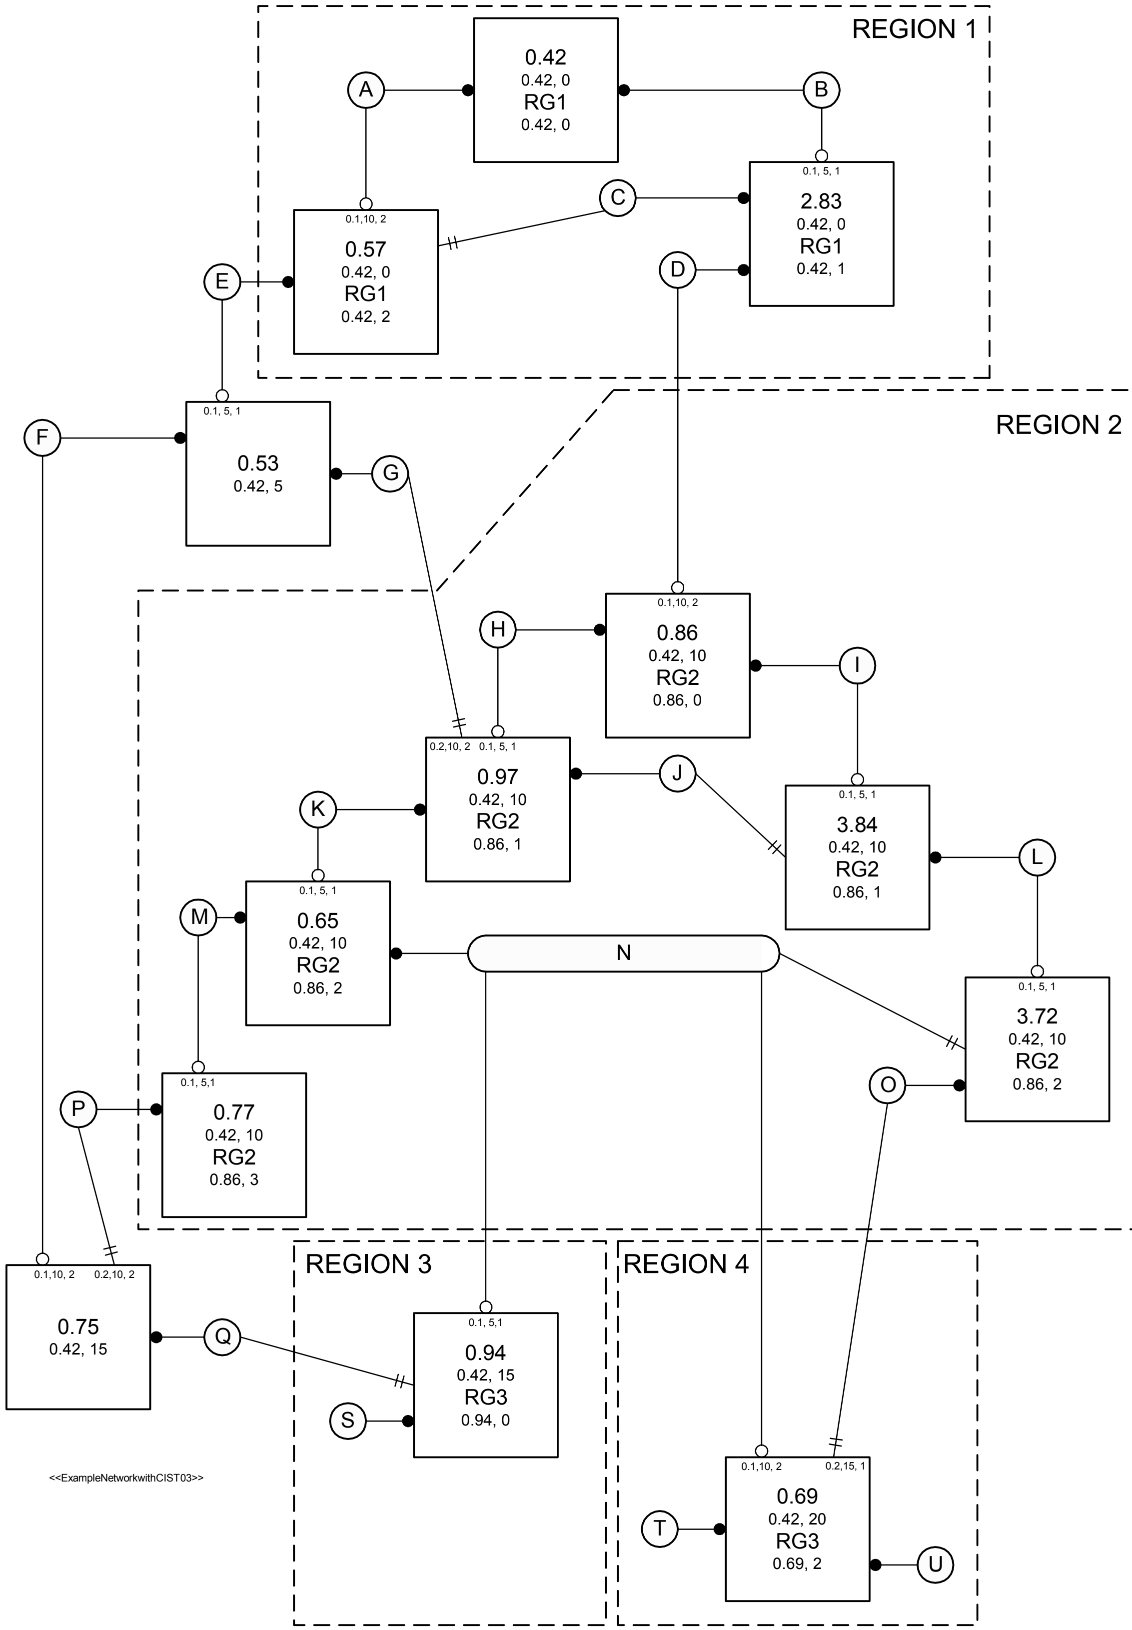

Figure 13-6—Example Network with CIST Priority Vectors, Port Roles, and MST Regions

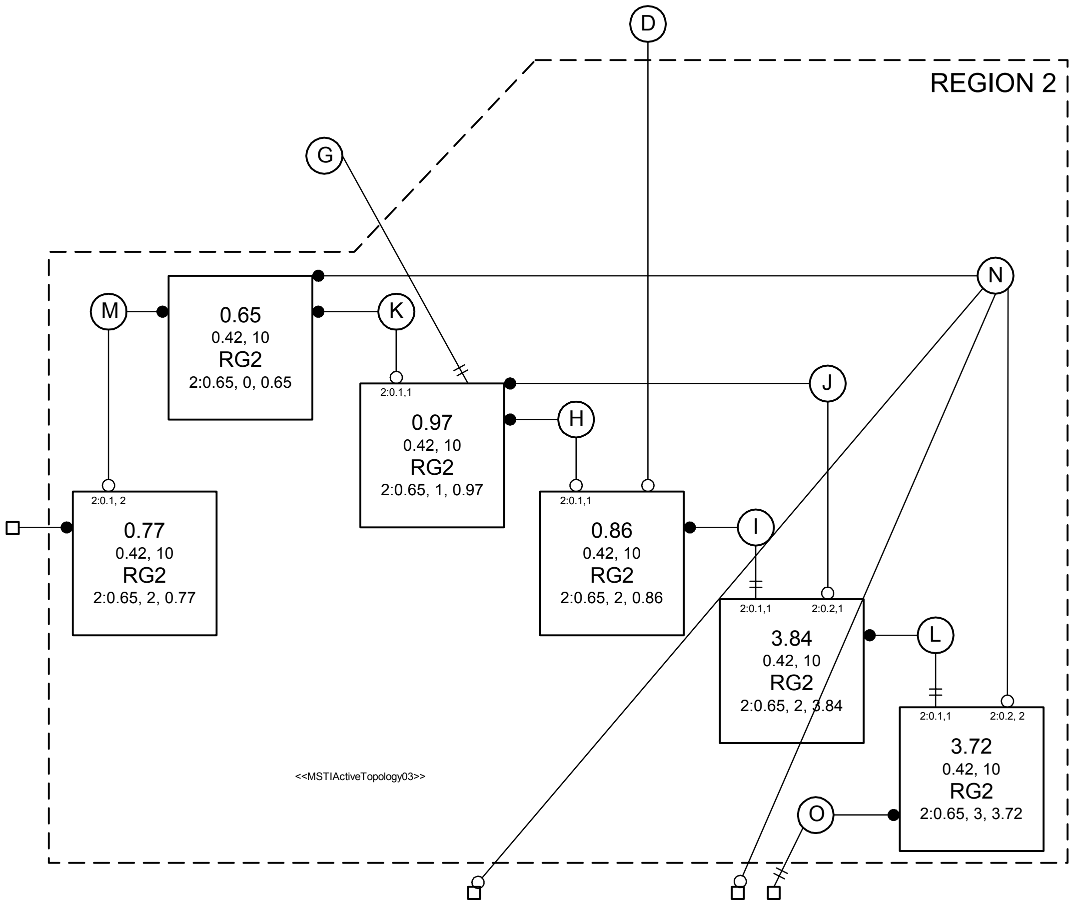

Connections to other MST regions

Figure 13-7—MSTI Active Topology in Region 2 of the example network

a) The selection of the CIST Root Bridge and the computation of port roles for the CIST uses the same fundamental algorithm (17.~~4~~3.1 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) but extended priority vector components and calculations (13.9, 13.10) within MST Regions as compared to RSTP (17.~~4~~5, 17.~~2~~ 6 of IEEE Std 802.1D, ~~1998~~ 2004 Edition). The effect of these extensions is to cause each region to resemble a single bridge from the point of view of the CST as calculated by STP or RSTP.

b) MST Configuration Identification is specific to MSTP.

c) The selection of the MSTI Regional Root Bridge and computation of port roles for each MSTI also uses the same fundamental spanning tree algorithm but modified priority vector components (13.11).

d) Different Bridges may be selected as the Regional Root for different MSTIs by modifying the manageable priority component of the Bridge Identifier differently for the MSTIs.

e) The port roles used by the CIST (Root, Designated, Alternate, Backup or Disabled Port) are the same as those of STP and RSTP (17.~~4~~ 3.1 of IEEE Std 802.1D, ~~1998~~ 2004 Edition). The MSTIs use the additional port role Master Port. The Port States associated with each spanning tree and bridge port are the same as those of RSTP (~~17~~7.~~5~~ 4 of IEEE Std 802.1D, ~~1998~~ 2004 Edition).

f) The state variables associated with each port for each spanning tree and for the tree itself are those specified for RSTP as per bridge port and per bridge (17.~~13~~15, 17.~~15~~17, 17.~~17~~18, 17.~~18~~ 19 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) with few exceptions, additions, and enhancements.

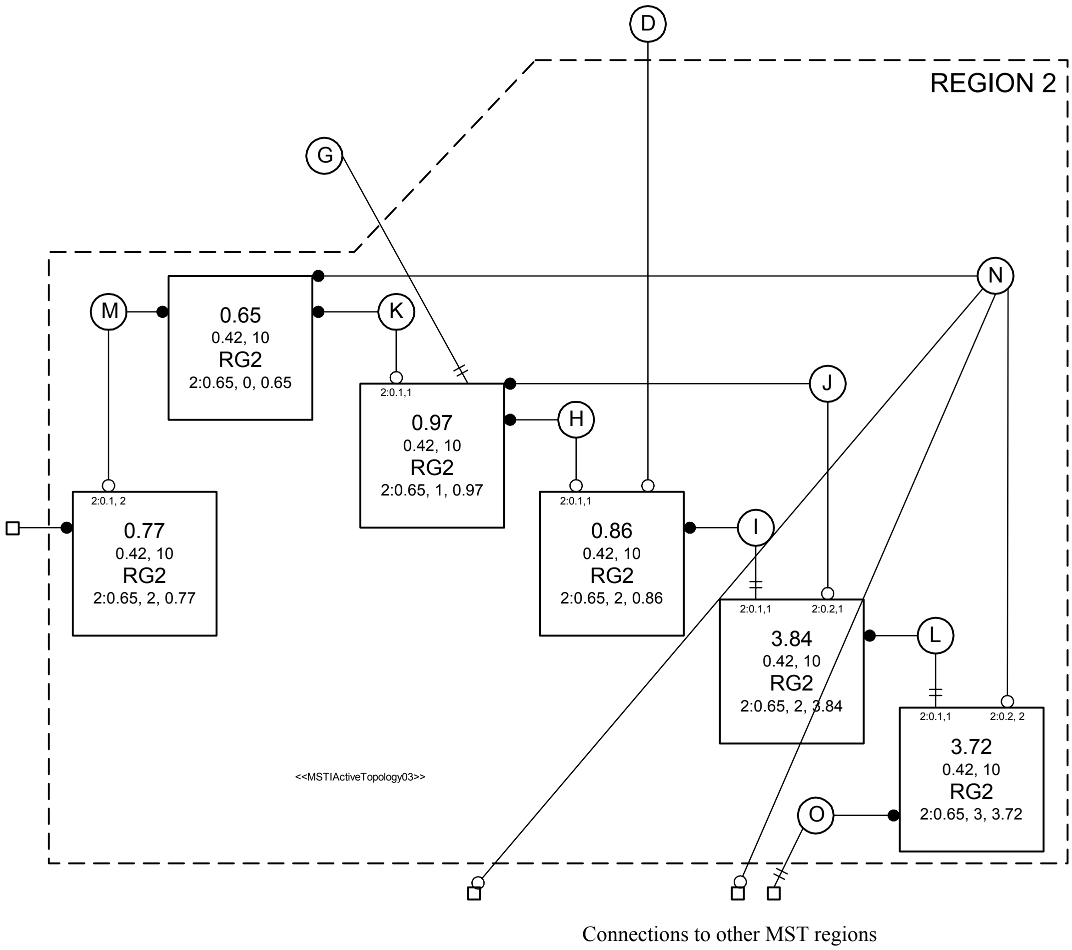

Figure 13-8—MSTI Active Topology in Region 2 of the example network

g) The state machine performance parameters specified for RSTP (17.~~6~~ 13 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) apply equally to the CIST. A simplified set of performance parameters apply to the MSTIs.

h) The state machine procedures of RSTP are used (17.~~9~~ 21 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) with detailed changes.

MSTP, like RSTP:

i) Cannot protect against temporary loops caused by the inter-connection of two LAN segments by devices other than Bridges (e.g., LAN repeaters) that operate invisibly with respect to support of the Bridges’ MAC Internal Sublayer Service.

j) Provides for rapid recovery of connectivity following the failure of a Bridge, Bridge Port, or a LAN. The timers used define worst case delays, only used to backup the normal operation of the protocol.

k) Provides a Force Protocol Version parameter, controlled by management and applicable to all Ports and trees supported by ~~an~~ a MST bridge, to instruct MSTP to emulate aspects of early versions of spanning tree protocol. In particular the Force Protocol Version parameter allows rapid transitions to be disabled. This reduces the risk of an increase, as compared to STP, in the rates of frame duplication and misordering in the Bridged LAN, as discussed in ~~F.2.4~~ Annex K of IEEE Std 802.1D, ~~1998~~ 2004 Edition.

l) Allows Bridge Ports to be configured such that they can transition directly to the Forwarding Port State on re-initialization of the Bridge. This may be appropriate where a specific Bridge Port is known to be connected to a LAN segment that is not connected to further Bridges. The per port operational control, operEdge, that supports this behavior applies equally to all the spanning trees of ~~an~~ a MST Bridge.

13.5 Modelling ~~an~~ a MST Region as a single RSTP Bridge

The specification of MSTP is such that the nominal replacement of an entire MST Region by a single RSTP Bridge leads to little change in the behavior of the remainder of the bridged local area network. This design is intended to assist those familiar with the STP and RSTP specifications to comprehend and verify MSTP, and to administer bridged local area networks using the MSTP specification.

In a network comprising STP Bridges, RSTP Bridges, and multiple MST Regions, treating the MST Regions as single Bridges provides the network administrator with a natural hierarchy. The internal management of MST Regions can be largely separated from the management of the active topology of the ~~bridge~~ bridged local area network as a whole.

The portion of the active topology of the network that connects any two bridges in the same MST Region traverses only MST Bridges and LANs in that region, and never Bridges of any kind outside the region, in other words connectivity within the region is independent of external connectivity. This is because the protocol parameters that determine the active topology of the network as a whole, the Root Identifier and Root Path Cost (known in the MSTP specification as the CIST Root Identifier and CIST External Root Path Cost) are carried unchanged throughout and across the MST Region, so bridges within the region will always prefer spanning tree information that has been propagated within the region to information that has exited the region and is attempting to re-enter it.

NOTE 1—No LAN can be in more than one Region at a time, so two Bridges (0.11 and 0.22 say) that would otherwise be in the same MST Region by virtue of having the same MST Configuration and of being directly connected by a LAN, may be in distinct regions if that is a shared LAN with other Bridges attached (having a different MST Configuration) and no other connectivity between 0.11 and 0.22 and lying wholly within their Region is available. The Region that the shared LAN belongs to may be dynamically determined. No such dynamic partitioning concerns arise with single Bridges. Obviously the sharing of LANs between administrative regions militates against the partitioning of concerns, and should only be done following careful analysis.

The Port Path Cost (MSTP's External Port Path Cost) is added to the Root Path Cost just once at the Root Port of the CIST Regional Root, the closest Bridge in the Region to the Root Bridge of the entire network. The Message Age used by STP and RSTP is also only incremented at this Port. If the CIST Root is within ~~an~~ a MST Region it also acts as the Regional Root, and the Root Path Cost and Message Age advertised are zero, just as for a single Bridge.

Within ~~an~~ a MST Region, each MSTI operates in much the same way as an independent instance of RSTP with dedicated Regional Root Identifier, Internal Root Path Cost, and Internal Port Path Cost parameters. Moreover the overall spanning tree (the CIST) includes a fragment (the IST) within each MST Region that can be viewed as operating in the same way as ~~an~~ a MSTI with the Regional Root as its root.

NOTE 2—Since ~~an~~ a MST Region behaves like a single Bridge and does not partition (except in the unusual configuration involving shared LANs noted above) it has a single Root Port in the CST active topology. Partitioning a network into two or more Regions can therefore force non-optimal blocking of Bridge Ports at the boundaries rather than internal to those Regions.

13.6 STP and RSTP compatibility

MSTP is designed to be STP and RSTP compatible and interoperable without additional operational management practice.

13.6.1 Designated Port Selection

Correct operation of the spanning tree protocols requires that all Bridge Ports attached to any given LAN agree on a single CIST Designated Port after a short interval sufficient for any Bridge Port to receive a configuration message from that Designated Port.

A unique spanning tree priority (13.9) is required for each Bridge Port for STP, which has no other way of communicating port roles. Since port numbers on different bridges are not guaranteed to be unique, this necessitates the inclusion of the transmitting Bridge's Bridge Identifier in the STP BPDU. RSTP and MSTP's Port Protocol Migration state machines (13.30) ensure that all Bridges attached to any LAN with an attached STP Bridge send and receive STP BPDUs exclusively.

NOTE1 —This behavior satisfies the requirement for unique, agreed Designated Port for LANs with attached STP Bridges, but means that ~~an~~ a MST Region cannot completely emulate a single Bridge since the transmitted Designated Bridge Identifier can differ on Bridge Ports at the Region's boundary.

MSTP transmits and receives the Regional Root Identifier and not the Designated Bridge Identifier in the BPDU fields recognized by RSTP (14.6) to allow both the MSTP and RSTP Bridges potentially connected to a single LAN to perform comparisons (13.9, 13.10) between all the spanning tree priority vectors transmitted that yield a single conclusion as to which RSTP Bridge or MST Region includes the Designated Port. MST and RST BPDUs convey the transmitting port's CIST Port Role. This is checked on receipt by RSTP when receiving messages from a Designated Bridge (17.~~19~~21.8 of IEEE Std 802.1D, ~~1998~~ 2004 Edition), thus ensuring that ~~an~~ a RSTP Bridge does not incorrectly identify one MST Bridge Port as being Designated rather than another, even while omitting the competing Bridge Ports' Designated Bridge Identifiers from comparisons.

NOTE 2—This ability of MSTP Bridges to communicate the full set of MSTP information on shared LANs to which RSTP Bridges are attached avoids the need for the Port Protocol Migration machines to detect RSTP Bridges. Two or more MSTP and one or more RSTP Bridges may be connected to a shared LAN, with full MSTP operation. This includes the possibility of different MSTI Designated Ports (see 13.3.1).

13.6.2 Force Protocol Version

A Force Protocol Version parameter, controlled by management, instructs MSTP to emulate additional aspects of the behavior of earlier versions of spanning tree protocol that are not strictly required for interoperability. The value of this parameter applies to all Ports of the Bridge.

a) ST BPDUs, rather than MST BPDUs, are transmitted if Force Protocol Version is 0. ~~RST~~ ST BPDUs omit the MST Configuration Identifier and all MSTI Information.
b) RST BPDUs, rather than MST BPDUs, are transmitted if Force Protocol Version is 2. RST BPDUs omit the MST Configuration Identifier and all MSTI Information.
c) All received BPDUs are treated as being from a different MST Region if Force Protocol Version is 0 or 2.
d) The MSTP state machines disable rapid transitions if Force Protocol Version is 0. This ~~allow~~ allows MSTP Bridges to support applications and protocols that may be sensitive to the increased rates of frame duplication and misordering that can arise under some circumstances, as discussed in Annex K of IEEE Std 802.~~1D F~~1D, 2004 Edition.~~2.4~~
e) The MSTP state machines allow full MSTP behavior if Force Protocol Version is 3 or more.

NOTE 1—Allowing for the case of a Force Protocol Version parameter value greater than 3 can simplify management of Bridges with different protocol versions.

NOTE 2—The Force Protocol Version parameter does not support multiple spanning trees with rapid transitions disabled.

13.7 MST Configuration Identification

It is essential that all Bridges within ~~an~~ a MST Region agree on the allocation of VIDs to specific spanning trees. If the allocation differs, frames for some VIDs may be duplicated or not delivered to some LANs at all. MST Bridges check that they are allocating VIDs to the same spanning trees as their neighboring MST Bridges in the same Region by transmitting and receiving MST Configuration Identifiers along with the spanning tree information. These MST Configuration Identifiers, while compact, are designed so that two matching identifiers have a very high probability of denoting the same configuration even in the absence of any supporting management practice for identifier allocation.

NOTE 1—Suitable management practices for the deployment of equipment and the choice of Configuration Names and Revision Levels (see below) can be used to guarantee that the MST Configuration Identifiers will differ if the VID to spanning tree allocation differs within a single administrative domain.

Each MST Configuration Identifier contains the following components:

1) A Configuration Identifier Format Selector, the value 0 encoded in a fixed field of one octet to indicate the use of the following components as specified in this Standard.
2) The Configuration Name, a variable length text string encoded within a fixed field of 32 octets, conforming to RFC 2271's definition of SnmpAdminString.
3) The Revision Level, an unsigned integer encoded within a fixed field of 2 octets.
4) The Configuration Digest, a 16 octet signature of type HMAC-MD5 (see IETF RFC 2104) created from the MST Configuration Table (~~3.26~~3.26, ~~8.11~~8.11). For the purposes of calculating the Configuration Digest, the MST Configuration Table is considered to contain 4096 consecutive two octet elements, where each element of the table (with the exception of the first and last) contains ~~an~~ a MSTID value encoded as a binary number, with the first octet being most significant. The first element of the table contains the value 0, the second element the MSTID value corresponding to VID 1, the third element the MSTID value corresponding to VID 2, and so on, with the next to last element of the table containing the MSTID value corresponding to VID 4094, and the last element containing the value 0. The key used to generate the signature consists of the 16 octet string specified in Table .

Table 13-1—Configuration Digest Signature Key

Parameter	Mandatory value
Configuration Digest Signature Key	0x13AC06A62E47FD51F95D2BA243CD0346

NOTE 2—The formulation of the signature as described above does not imply that a separate VID to MSTID translation table has to be maintained by the implementation; rather that it should be possible for the implementation to derive the logical contents of such a table, and the signature value as specified above, from the other configuration information maintained by the implementation, as described in Clause ~~12~~12.

The Configuration Digests of some VID to MSTID translations are shown in Table 13-2 to help verify implementations of this specification.

It is recommended that MST Bridge implementations provide an easily selectable or default configuration comprising a Configuration Name of the Bridge Address as a text string using the Hexadecimal Representation specified in IEEE Std 802~~-2001~~.1D, 2004 Edition, a Revision Level of 0, and a Configuration Digest representing a VID to MSTID translation table containing the value 0 for every element. Such a table represents the mapping of all VLANs to the CIST. Since the Bridge Address is unique to each MST Bridge, no two MST Bridges using this default configuration will be identified as belonging to the same MST Region.

Table 13-2—Sample Configuration Digest Signature Keys

VID to MSTID translation	Configuration Digest
All VIDs map to the CIST, no VID mapped to any MSTI	0xAC36177F50283CD4B83821D8AB26DE62
All VIDs map to MSTID 1	0xE13A80F11ED0856ACD4EE3476941C73B
Every VID maps to the MSTID equal to (VID modulo 32) + 1	0x9D145C267DBE9FB5D893441BE3BA08CE

<< Editor's note: the reference to IEEE Std 802.1D, 2004 Edition in the last paragraph replaces one to IEEE Std 802.1-2001: both references are unclear. >>

13.8 MST Regions

~~An~~ A MST Region comprises one or more MST Bridges with the same MST Configuration Identifiers, using the same MSTIs, interconnected by and including LANs for which one of those Bridges is the Designated Bridge for the CIST and which have no Bridges attached that cannot receive and transmit ~~RSTP~~ RST BPDUs.

13.9 Spanning Tree Priority Vectors

All Bridges, whether they use STP, RSTP, or MSTP, send information to each other, in Configuration Messages (13.14 of this standard, 17.~~7~~ 8 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) to assign Port roles that determine each Port's participation in a fully and simply connected active topology based on one or more spanning trees. The information communicated is known as a *spanning tree priority vector*. Spanning tree priority vectors provide the basis for a concise specification of each protocol's computation of the active topology, in terms of both the entire Bridged LAN and the operation of individual Bridges in support of the distributed algorithm.

CIST priority vectors comprise the following components:

a) CIST Root Identifier, the Bridge Identifier of the CIST Root;
b) CIST External Root Path Cost, the path cost between MST Regions from the transmitting Bridge to the CIST Root;
c) CIST Regional Root Identifier, the Bridge Identifier of the single bridge in a Region whose CIST Root Port is a Boundary Port, or the Bridge Identifier of the CIST Root if that is within the Region;
d) CIST Internal Root Path Cost, the path cost to the CIST Regional Root;
e) CIST Designated Bridge Identifier, the Bridge Identifier for the transmitting bridge for the CIST;
f) CIST Designated Port Identifier, the Port Identifier for the transmitting port for the CIST;
g) CIST Receiving Port Identifier (not conveyed in Configuration Messages, used as tie-breaker between otherwise equal priority vectors within a receiving Bridge).

The CIST External Root Path Cost is significant throughout the Bridged LAN. It is propagated along each path from the CIST Root, and is added to at Bridge Ports that receive the priority vector from a Bridge in a different MST Region. The External Path Cost transmitted by a Bridge thus represents costs accumulated at the Root Ports of Bridges that are either not MST Bridges or are CIST Regional Roots, and is constant within a Region. The CIST Internal Root Path Cost is only significant and explicitly defined within a Region.

NOTE 1—The path to the CIST Root from a bridge with a CIST Root Port within a region always goes to or through the CIST Regional Root.

NOTE 2—The STP and RSTP specifications refer to the CIST Root Identifier and CIST External Root Path Cost simply as the Root Bridge Identifier and Root Path Cost respectively and omit the CIST Internal Root Path Cost. MSTP encodes the CIST Regional Root Identifier in the (External Designated) Bridge Identifier BPDU field used by RSTP to convey the Designated Bridge Identifier (14.3.3), so an entire MST Region appears to ~~an~~ a RSTP capable Bridge as a single Bridge. However, this is not possible for STP, as the latter lacks the fields necessary for MST Bridges to communicate the Designated Bridge Identifier to resolve a potential priority vector tie, and MSTP BPDUs are not sent on a LAN to which ~~an~~ a STP Bridge is attached.

MSTI priority vectors comprise the following components:

h) MSTI Regional Root Identifier, the Bridge Identifier of the MSTI Regional Root for this particular MSTI in this MST Region;
i) MSTI Internal Root Path Cost, the path cost to the MSTI Regional Root for this particular MSTI in this MST Region;
j) MSTI Designated Bridge Identifier, the Bridge Identifier for the transmitting bridge for this MSTI;
k) MSTI Designated Port Identifier, the Port Identifier for the transmitting port for this MSTI;
l) MSTI Receiving Port Identifier (not conveyed in Configuration Messages).

The set of priority vectors for a given MSTI is only defined within ~~an~~ a MST Region. Within each Region they are totally and uniquely ordered. A CIST Root Identifier, CIST External Root Path Cost, and CIST Regional Root Identifier tuple defines the connection of the Region to the external CST and is required to be associated with the source of the MSTI priority vector information when assessing the agreement of information for rapid transitions to forwarding, but plays no part in priority vector calculations.

As each Bridge and Bridge Port receives priority vector information from other Bridges and Ports closer to the Root, priority vector calculations and comparisons are made to decide which priority information to record, and what information to be passed on. Decisions about a given Port's role are made by comparing the priority vector components that could be transmitted with that received by the Port. For all components a lesser numerical value is better, and earlier components in the above lists are more significant. As each Bridge Port receives priority vector information from Ports closer to the Root, additions are made to one or more priority vector components to yield a worse priority vector for potential transmission through other ports of the same Bridge.

NOTE 3—The consistent use of lower numerical values to indicate better information is deliberate as the Designated Port that is closest to the Root Bridge, i.e. has a numerically lowest path cost component, is selected from amongst potential alternatives for any given LAN (13.9). Adopting the conventions that lower numerical values indicate better information, that where possible more significant priority components are encoded earlier in the octet sequence of a BPDU (14.3), and that earlier octets in the encoding of individual components are more significant (14.2) allow concatenated octets that compose a priority vector to be compared as if they were a multiple octet encoding of a single number, without regard to the boundaries between the encoded components. To reduce the confusion that naturally arises from having the lesser of two numerical values represent the better of the two, i.e. the one to be chosen all other factors being equal, this clause uses the following consistent terminology. Relative numeric values are described as "least", "lesser", "equal", and "greater", and their comparisons as "less than", "equal to", ~~or~~ and "greater than", while relative Spanning Tree priorities are described as "best", "better", "the same", "different", and "worse" and their comparisons as "better than", "the same as", "different from", and "worse than". The operators "<" and "=" represent less than and equal to respectively. The terms "superior" and "inferior" are used for comparisons that are not simply based on priority but include the fact that a priority vector can replace an earlier vector transmitted by the same Bridge Port. All these terms are defined for priority vectors in terms of the numeric comparison of components below (13.10, 13.11).

NOTE 4—To ensure that the CIST and each MSTI's view of the boundaries of each MST region remain in synchronization at all times, each BPDU carries priority vector information for the CIST as well as for MSTIs. Associating the CIST Root Identifier, External Path Cost, and Regional Root Identifier with the priority vector information for each MSTI does not therefore raise a requirement to transmit these components separately. A single bit per MSTI vector, the Agreement flag, satisfies the requirement to indicate that the vector beginning with the MSTI Regional Root Identifier for that specific MSTI has always been associated with the single CIST Root Identifier etc. transmitted in the BPDU.

To allow the active topology to be managed for each tree through adjusting the relative priority of different Bridges and Bridge Ports for selection as the CIST Root, a CSTI or MSTI Regional Root, Designated

Bridge, or Designated Port, the priority component of the Bridge's Bridge Identifier can be independently chosen for the CIST and for each MSTI. The priority component used by the CIST for its CIST Regional Root Identifier can also be chosen independently of that used for the CIST Root Identifier. Independent configuration of Port Path Cost and Port Priority values for the CIST and for each MSTI can also be used to control selection of the various roles for the CIST and for each MSTI.

13.10 CIST Priority Vector Calculations

The *port priority vector* is the priority vector held for the port when the reception of BPDUs and any pending update of information has been completed:

port priority vector = *{RootID : ExtRootPathCost : RRootID : IntRootPathCost : DesignatedBridgeID : DesignatedPortID : RcvPortID}*

The *message priority vector* is the priority vector conveyed in a received Configuration Message. For a Bridge with Bridge Identifier B receiving a Configuration Message on a Port P_B from a Designated Port P_D on Bridge D claiming a CIST Root Identifier of R_D, a CIST External Root Path Cost of ERC_D, a CIST Regional Root Identifier of RR_D, and a CIST Internal Root Path Cost of IRC_D:

message priority vector = $\{R_D : ERC_D : RR_D : IRC_D : D : P_D : P_B\}$

If B is not in the same MST Region as D, the Internal Root Path Cost is decoded as 0, as it has no meaning to B.

NOTE—If a Configuration Message is received in ~~an RSTP~~ a RST or ~~STP~~ ST BPDU, both the Regional Root Identifier and the Designated Bridge Identifier are decoded from the single BPDU field used for the Designated Bridge Parameter (the MST BPDU field in this position encodes the CIST Regional Root Identifier). ~~An~~ A STP or RSTP Bridge is always treated by MSTP as being in ~~an~~ a MST Region of its own, so the Internal Root Path Cost is decoded as zero, and the tests below become the familiar checks used by STP and RSTP.

The received CIST message priority vector is the same as B's port priority vector if:

$$(R_D == RootID)\ \&\&\ (ERC_D == ExtRootPathCost)\ \&\&\ (RR_D == RRootID)\ \&\&$$
$$(IRC_D == IntRootPathCost)\ \&\&\ (D == DesignatedBridgeID)\ \&\&\ (P_D == DesignatedPortID)$$

and is better if:

$$((R_D < RootID))\ ||$$
$$((R_D == RootID)\ \&\&\ (ERC_D < ExtRootPathCost))\ ||$$
$$((R_D == RootID)\ \&\&\ (ERC_D == ExtRootPathCost)\ \&\&\ (RR_D < RRootID))\ ||$$
$$((R_D == RootID)\ \&\&\ (ERC_D == ExtRootPathCost)\ \&\&\ (RR_D == RRootID)$$
$$\&\&\ (IRC_D < IntRootPathCost))\ ||$$
$$((R_D == RootID)\ \&\&\ (ERC_D == ExtRootPathCost)\ \&\&\ (RR_D == RRootID)$$
$$\&\&\ (IRC_D == IntRootPathCost)\ \&\&\ (D < DesignatedBridgeID))\ ||$$
$$((R_D == RootID)\ \&\&\ (ERC_D == ExtRootPathCost)\ \&\&\ (RR_D == RRootID)$$
$$\&\&\ (IRC_D == IntRootPathCost)\ \&\&\ (D == DesignatedBridgeID)$$
$$\&\&\ (P_D < DesignatedPortID))$$

A received CIST message priority vector is superior to the port priority vector if, and only if, the message priority vector is better than the port priority vector, or the Designated Bridge Identifier and Designated Port Identifier components are the same in which case the message has been transmitted from the same Designated Port as a previously received superior message, i.e. if:

$$
\begin{aligned}
&(\{R_D : ERC_D : RR_D : IRC_D : D : P_D : P_B\} \\
&\qquad \textit{is better than} \\
&\{RootID : ExtRootPathCost : RRootID : IntRootPathCost : \\
&\qquad DesignatedBridgeID : DesignatedPortID : RcvPortID\} \\
&)\,||\,(\quad (\cancel{D}\ \underline{D.BridgeAddress} == DesignatedBridgeID\underline{.BridgeAddress})\ \underline{\&}\&\cancel{\&} \\
&\qquad (\cancel{P_D}\ \underline{P_D.PortNumber} == DesignatedPortID\underline{.PortNumber}))
\end{aligned}
$$

If the message priority vector received in a Configuration Message from a Designated Port is superior it will replace the current port priority vector.

A *root path priority vector* for a Port can be calculated from a port priority vector that contains information from a message priority vector, as follows.

If the port priority vector was received from a bridge in a different MST Region (13.27.5), the External Port Path Cost EPC_{PB} is added to the External Root Path Cost component, and the Regional Root Identifier is set to the value of the Bridge Identifier for the receiving Bridge. The Internal Root Path Cost component will have been set to zero on reception.

$$\textit{root path priority vector} = \{R_D : ERC_D + EPC_{PB} : B : 0 : D : P_D : P_B)$$

If the port priority vector was received from a bridge in the same MST Region (13.27.5), the Internal Port Path Cost IPC_{PB} is added to the Internal Root Path Cost component.

$$\textit{root path priority vector} = \{R_D : ERC_D : RR_D : IRC_D + IPC_{PB} : D : P_D : P_B)$$

The *bridge priority vector* for a Bridge *B* is the priority vector that would, with the Designated Port Identifier set equal to the transmitting Port Identifier, be used as the message priority vector in Configuration Messages transmitted on Bridge *B's* Designated Ports if *B* was selected as the Root Bridge of the CIST.

$$\textit{bridge priority vector} = \{B : 0 : B : 0 : B : 0 : 0\}$$

The *root priority vector* for Bridge *B* is the best priority vector of the set of priority vectors comprising the bridge priority vector plus all root path priority vectors whose Designated Bridge Identifier *D* is not equal to *B*. If the bridge priority vector is the best of this set of priority vectors, Bridge *B* has been selected as the Root of the tree.

The *designated priority vector* for a port *Q* on Bridge *B* is the root priority vector with *B's* Bridge Identifier *B* substituted for the *DesignatedBridgeID* and *Q's* Port Identifier Q_B substituted for the *DesignatedPortID* and *RcvPortID* components. If *Q* is attached to a LAN which has one or more STP Bridges attached (as determined by the Port Protocol Migration state machine), *B's* Bridge Identifier *B* is also substituted for the the *RRootID* component.

If the designated priority vector is better than the port priority vector, the Port will be the Designated Port for the attached LAN and the current port priority vector will be updated. The message priority vector in Configuration Messages transmitted by a Port always comprises the components of the port priority vector of the Port, even if the Port is a Root Port.

13.11 MST Priority Vector Calculations

The *port priority vector* is the priority vector held for the port when the reception of BPDUs and any pending update of information has been completed:

port priority vector = $\{RRootID : IntRootPathCost : DesignatedBridgeID : DesignatedPortID : RcvPortID\}$

The *message priority vector* is the priority vector conveyed in a received Configuration Message. For a Bridge with Bridge Identifier B receiving a Configuration Message on a Regional Port P_B from a Designated Port P_D on Bridge D belonging to the same MST Region and claiming an Internal Root Path Cost of IRC_D:

message priority vector = $\{RR_D : IRC_D : D : P_D : P_B\}$

~~An~~ A MSTI message priority vector received from a Bridge that does not belong to the same MST Region is discarded.

~~An~~ A MSTI message priority vector received from a bridge port internal to the region is the same as the port priority vector if:

$$((RR_D == RRootID)\ \&\&\ (IRC_D == IntRootPathCost)\ \&\&\ (D == DesignatedBridgeID)\ \&\&\ (P_D == DesignatedPortID))$$

and is better if:

$$\begin{aligned}&((RR_D < RRootID))\ \|\\ &((RR_D == RRootID)\ \&\&\ (IRC_D < IntRootPathCost))\ \|\\ &((RR_D == RRootID)\ \&\&\ (IRC_D == IntRootPathCost)\ \&\&\ (D < DesignatedBridgeID))\ \|\\ &((RR_D == RRootID)\ \&\&\ (IRC_D == IntRootPathCost)\ \&\&\ (D == DesignatedBridgeID)\\ &\qquad \&\&\ (P_D < DesignatedPortID))\end{aligned}$$

~~An~~ A MSTI message priority vector is superior to the port priority vector if, and only if, the message priority vector is better than the port priority vector, or the Designated Bridge Identifier and Designated Port Identifier components are the same in which case the message has been transmitted from the same Designated Port as a previously received superior message, i.e. if:

$$\begin{aligned}&(\{RR_D : IRC_D : D : P_D : P_B\}\\ &\qquad \text{is better than}\\ &\{RRootID : IntRootPathCost : DesignatedBridgeID : DesignatedPortID : RcvPortID\}\\ &)\ \|\ (\quad (\cancel{D}\ D.BridgeAddress == DesignatedBridgeID.BridgeAddress)\ \&\&\ \cancel{\&}\\ &\qquad (\cancel{P_D}\ P_D.PortNumber == DesignatedPortID.PortNumber))\end{aligned}$$

If the message priority vector received in a Configuration Message from a Designated Port for the MSTI is superior it will replace the current port priority vector.

NOTE 1—the agreed flag (13.24.2) for the Port and this MSTI will be cleared if the CIST Root Identifier, CIST External Root Path Cost, and CIST Regional Root Identifier) in the received BPDU are not the same as those for the CIST designated priority vector for the port following processing of the received BPDU.

A *root path priority vector* for a given MSTI can be calculated for a Port that has received a port priority vector from a bridge in the same region by adding the Internal Port Path Cost IPC_{PB} to the Internal Root Path Cost component.

root path priority vector = $\{RR_D : IRC_D + IPC_{PB} : D : P_D : P_B)$

NOTE 2—Internal Port Path Costs are independently manageable for each MSTI, as are the priority components of the Bridge and Port Identifiers. This permits topology management of each MSTI independent of other MSTIs. The ability to independently manage MSTIs in this way without explicitly transmitting individual Port Path Costs is a key reason for retaining the use of a Distance Vector protocol for constructing MSTIs. A simple Link State Protocol requires transmission (or apriori sharing) of all Port Costs for all links.

The *bridge priority vector* for a Bridge *B* is the priority vector that would, with the Designated Port Identifier set equal to the transmitting Port Identifier, be used as the message priority vector in Configuration Messages transmitted on Bridge *B's* Designated Ports if *B* was selected as the Root Bridge of a given tree.

bridge priority vector = {B : 0 : B : 0}

The *root priority vector* for Bridge *B* is the best priority vector of the set of priority vectors comprising the bridge priority vector plus all root path priority vectors whose Designated Bridge Identifier *D* is not equal to *B*. If the bridge priority vector is the best of this set of priority vectors, Bridge *B* has been selected as the Root of the tree.

The *designated priority vector* for a port *Q* on Bridge *B* is the root priority vector with *B's* Bridge Identifier *B* substituted for the *DesignatedBridgeID* and *Q's* Port Identifier Q_B substituted for the *DesignatedPortID* and *RcvPortID* components.

If the designated priority vector is better than the port priority vector, the Port will be the Designated Port for the attached LAN and the current port priority vector will be updated. The message priority vector in MSTP BPDUs transmitted by a Port always comprises the components of the port priority vector of the Port, even if the Port is a Root Port.

Figure 13-8 shows the priority vectors and the active topology calculated for ~~an~~ a MSTI in a Region of the example network of Figure 13-3.

13.12 Port Role Assignments

Port Role assignments for Bridge Ports that are enabled are determined by each bridge in the Bridged Local Area Network (13.12) according to the source and relative priority of the spanning tree port priority vectors (13.9) selected for each Port following priority vector calculations (13.10, 13.11).

Each MST Bridge Port that is enabled is assigned a Port Role for each spanning tree. First one of the following roles: Root Port, Designated Port, Alternate Port, or Backup Port, is assigned for the CIST.

a) If the Bridge is not the CIST Root, the Port that is the source of the root priority vector is the CIST Root Port.
b) Each Port whose port priority vector is the designated priority vector derived from the root priority vector is a CIST Designated Port.
c) Each Port, other than the Root Port, that has a port priority vector that has been received from another Bridge is a CIST Alternate Port.
d) Each Port that has a port priority vector that has been received from another Port on this Bridge is a CIST Backup Port.

Then one of these roles, or the additional role of Master Port, is assigned for each MSTI.

e) If the Port is the CIST Root Port and the CIST port priority vector was received from a Bridge in another MST Region, the Port is the MSTI Master Port.
f) If the Bridge is not the MSTI Regional Root, the Port that is the source of the MSTI root priority vector is the MSTI Root Port.
g) Each Port whose port priority vector is the designated priority vector derived from the root priority vector is a MSTI Designated Port.
h) Each Port, other than a Master Port or Root Port, that has a port priority vector that has been received from another Bridge or has a CIST port priority vector that has been received from a Bridge in a different region, is ~~an~~ a MSTI Alternate Port.

i) Each Port that has a port priority vector that has been received from another Port on this Bridge is ~~an~~ a MSTI Backup Port.

If the Port is not enabled, it is assigned the Disabled Port role for the CIST and all MSTIs, to identify it as having no part in the operation of any of the spanning trees or the active topology of the network.

<< Editor's note: shall we refer to 17.7 in IEEE Std 802.1D, 2004 Edition or make explicit the conditions for which a Port is disabled (e.g. MAC_Operational FALSE, network access port, AuthControlledPortStatus Unauthorized, etc.)? >>

13.13 Stable Connectivity

This clause provides an analysis to show that MSTP meets its goal of providing full and simple connectivity for frames allocated to any given VLAN in a stable network, i.e. where the physical topology has remained constant for long enough that the spanning tree information communicated and processed by Bridges is not changing on any Bridge Port.

Each MST Region independently can allocate such frames to the IST or any given MSTI. Root Ports, Designated Ports, and Master Ports forward data frames, and Alternate, Backup, and Disabled Ports do not.

NOTE—The term Common Spanning Tree (CST) refers to the CIST connectivity between Regions, and the term Internal Spanning Tree (IST) to the CIST connectivity within each Region.

The CIST interconnects both individual LANs and Bridges, and complete MST Regions into a single spanning tree, each Region being part of the CIST as a whole. Frames with VIDs consistently allocated to the CIST within every MST Region follow an active topology determined by the minimum path costs to each Bridge and LAN provided by that single tree throughout the network, and thus enjoy full and simple connectivity.

Frames otherwise allocated follow the CIST outside and ~~an~~ a MSTI within ~~an~~ a MST Region. Simple and, in the absence of continual changes in physical connectivity, full connectivity of this composite active topology is ensured as follows:

a) Each Bridge or LAN is in one and only one Region.
(SST Bridges, LANs connected to STP Bridges, and LANs whose Designated Bridge is an SST Bridge, are all conveniently regarded as being in a Region of their own.)
b) Each and every frame is associated with one and only one VID.
c) Frames with any given VID are allocated either to the IST or to a given MSTI within any given Region, i.e. all frames are allocated to some tree and no frames are allocated to more than one tree.
d) The IST and each MSTI provides full and simple connectivity between all LANs and Bridges in ~~an~~ a MST Region for frames allocated to the IST or that MSTI.

Hence full and simple connectivity is provided for all frames from any Bridge or LAN within ~~an~~ a MST Region to any other within the Region.

Further:

e) All Bridges within ~~an~~ a MST Region with ports connected to a given LAN reach a consistent agreement as to whether each of those ports is or is not a Boundary Port (i.e. attaches a Bridge to a LAN that is not in the same Region) prior to forwarding frames.
(MST Bridges make the determination on the basis of the CIST Designated Port for the LAN or the selection of the protocol by the Protocol Migration machines, both are necessarily complete prior to frame forwarding. SST Bridges being unaware of MST Regions behave as if each LAN is in a different Region to the Bridge.)

f) At a Boundary Port frames allocated to the CIST and all MSTIs are forwarded or not forwarded alike. This is because Port Role assignments are such that if the CIST Port Role is Root Port the MSTI Port Role will be Master Port, and if the CIST Port Role is Designated Port, Alternate Port, Backup Port, or Disabled Port, each MSTI's Port Role will be the same.
g) The CIST provides full and simple connectivity between all LANs and Bridges in the network, including the LANs and Bridges attached to the Boundary Ports of any MST Region.

Hence full and simple connectivity is provided for all frames between Bridges and LANs outside the MST Region since those frames will be carried across the MST Region if necessary, just as if they were allocated to the CIST whichever tree they are allocated to within the Region.

Similarly full and simple connectivity is provided for all frames between a Bridge or LAN inside the Region and a Bridge or LAN outside the region since the connectivity provided from within the Region by ~~an~~ a MSTI to that outer Bridge or LAN is the same as that provided by the CIST.

Figure 13-10 illustrates the above connectivity with the simple example of Region 1 from the example network of ~~Figure 13-3~~Figure 13-5 and Figure 13-6. Bridge 0.42 has been selected as the CIST Root and Regional Root, Bridge 0.57 as the Regional Root for MSTI 1, and Bridge 2.83 for MSTI 2 by management of the per MSTI Bridge Identifier priority component. The potential loop through the three bridges in the Region is blocked at different Bridge Ports for the CIST, and each MSTI, but the connectivity across the Region and from each LAN and Bridge in the region through the boundaries of the Region is the same in all cases.

13.14 Communicating Spanning Tree Information

Bridges transmit and receive MAC frames, each containing a Bridge Protocol Data Unit (BPDU) (Clauses 9 and 14 of IEEE Std 802.1D, ~~1998~~ 2004 Edition), to communicate Spanning Tree messages. A MAC frame conveying a BPDU carries the Bridge Group Address in the destination address field and is received by all the Bridges connected to the LAN on which the frame is transmitted. The Bridge Group Address is one of a small number of addresses that identify frames that are not directly forwarded by Bridges (7.12.6), but the information contained in the BPDU can be used by a Bridge in calculating its own BPDUs to transmit, and can stimulate that transmission.

BPDUs are used to convey three types of Spanning Tree message:

a) Configuration Messages;
b) Topology Change Notification (TCN) Messages;
c) MST Configuration Identifiers.

A Configuration Message for the CIST can be encoded and transmitted in ~~an~~ a STP Configuration BPDU (9.3.1), ~~an~~ a RST BPDU (9.3.3), or ~~an~~ a MST BPDU (14). A TCN Message for the CIST can be encoded in ~~an~~ a STP Topology Change Notification BPDU (9.3.2), ~~an~~ a RST BPDU with the TC flag set, or ~~an~~ a MST BPDU. Configuration and TCN Messages for the CIST and for all MSTIs in ~~an~~ a MST Region are encoded in a single MST BPDU, as is the MST Configuration Identifier. No more than 64 MSTI Configuration Messages may be encoded in ~~an~~ a MST BPDU, and no more than 64 MSTIs may be supported by ~~an~~ a MST Bridge.

Configuration and Topology Change Notification BPDUs are distinguished from each other and from RST and MST BPDUs by their BPDU Type (Clause 9 of IEEE Std 802.1D, ~~1998~~ 2004 Edition). RST and MST BPDUs share the same BPDU Type and are distinguished by their version identifiers. Bridges implementing STP (Clause 8 of IEEE Std 802.1D, 1998 Edition) transmit and decode Configuration and Topology Change

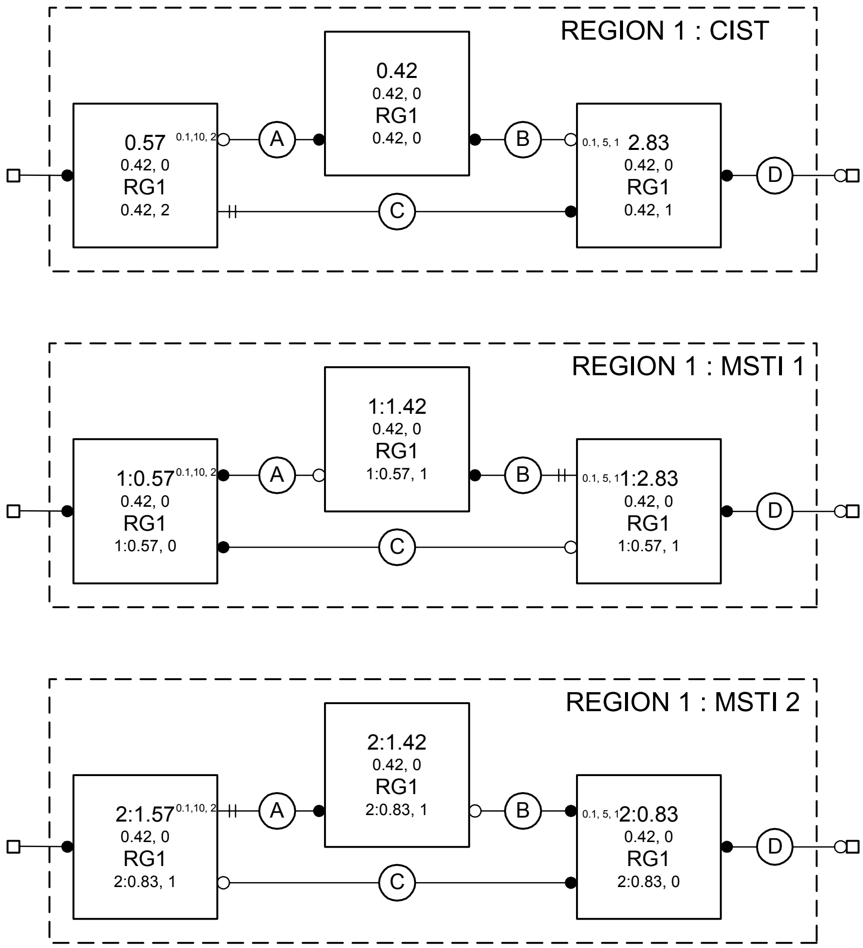

Figure 13-9—CIST and MSTI active topologies in Region 1 of the example network

Notification BPDUs, and ignore RST and MST BPDUs on receipt. This ensures that connection of a Bridge Port of such a Bridge to a LAN that is also attached to by a Bridge implementing RSTP or MSTP is detected, as transmission of RSTP or MSTP BPDUs does not suppress regular transmissions by the STP Bridge. This functionality is provided by the Port Protocol Migration state machine for RSTP (17.~~26~~ 24 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) and MSTP (~~17.26~~13.30). The Port Protocol Migration state machines selects the BPDU types used to encode Spanning Tree messages so that all Bridges attached to the same LAN participate in a spanning tree protocol, while maximizing the available functionality. If one or more attached Bridges only implement STP, only Configuration and Topology Change Notification BPDUs will be used and the functionality provided by the protocol will be constrained.

Each Configuration Message contains, among other parameters, a message priority vector (~~13.4.2~~13.10, ~~13.4.3~~13.11). This allows a receiving Bridge to determine the Port Role (~~13.4.4~~13.12) including that of Designated Port. Configuration Messages are transmitted if the information to be transmitted by a Designated Port changes, or if a ~~Root~~ Root, Master, Alternate or Backup Port has an Agreement to convey. In addition Designated Ports transmit Configuration Messages at regular intervals to guard against loss and to assist in the detection of failed components (LANs, Bridges, or Bridge Ports). In both cases, message transmission is subject to a maximum transmission rate (see ~~Transmission Limit~~ 17.26 in ~~13~~IEEE Std 802.~~28.2~~1D, 2004 Edition).

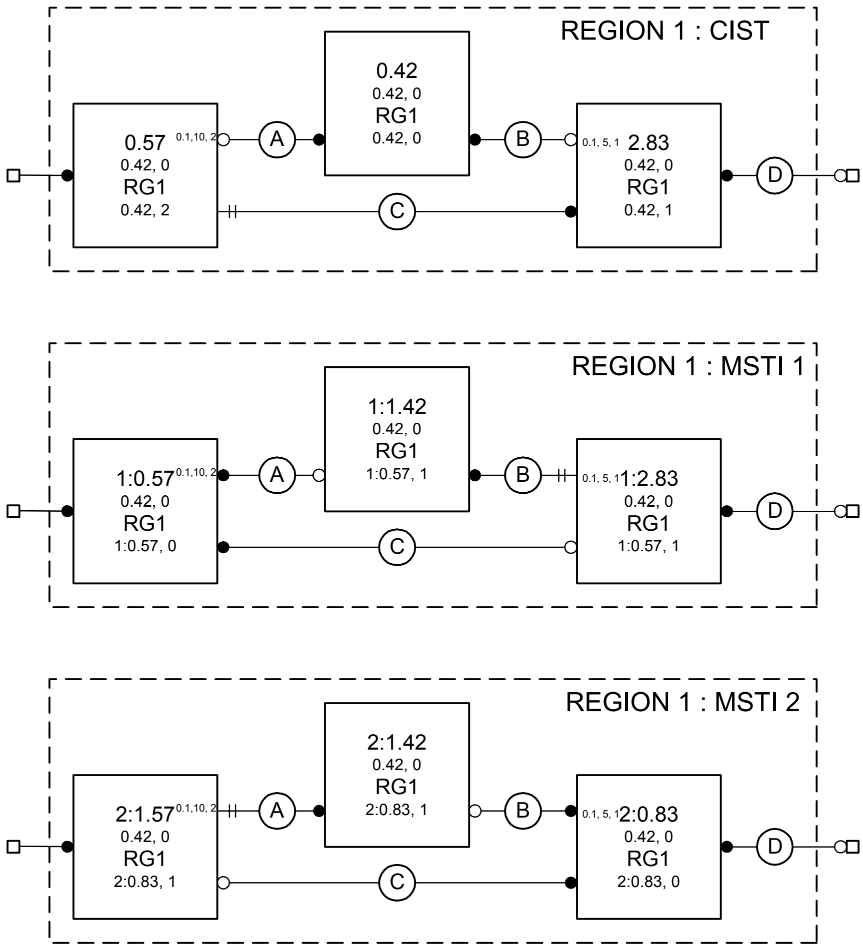

Figure 13-10—CIST and MSTI active topologies in Region 1 of the example network

13.15 Changing Spanning Tree Information

Addition, removal, failure, or management of the parameters of Bridges and LAN connectivity can change spanning tree information and require Port Role changes in all or part of the network (for the CIST) or all or part of ~~an~~ a MST Region (for ~~an~~ a MSTI). Received information for a spanning tree is considered superior to, and will replace, that recorded in the receiving Port's port priority vector if its message priority vector is better, or if it was transmitted by the same Designated Bridge and Designated Port and the message priority vector, timer, or hop count information differ from those recorded.

The new information will be propagated rapidly from Bridge to Bridge, superseding prior information and stimulating further transmissions until it reaches either Designated Ports that have already received the new information through redundant paths in the network or the leaves of the Spanning Tree, as defined by the new configuration. Configuration Message transmissions will then once more occur at regular intervals from Ports selected as Designated Ports.

To ensure that old information does not endlessly circulate through redundant paths in the network, preventing the effective propagation of the new information, MSTP associates a hop count with the information for each spanning tree. The hop count is assigned by the CIST Regional Root or the MSTI Regional Root and decremented by each receiving Port. Received information is discarded and the Port made a Designated Port if the hop count reaches zero.

If a Bridge Port's MAC_Operational parameter becomes FALSE, the Port becomes a Disabled Port and received spanning tree information is immediately discarded. Spanning tree information for the tree can be recomputed, the Bridge's Port Roles changed, and new spanning tree information transmitted if necessary. Not all component failure conditions can be detected in this way, so each Designated Port transmits spanning tree information at regular intervals and a receiving Port will discard information and become a Designated Port if two transmissions are missed.

The Spanning Tree Protocol (STP, Clause 8 of IEEE Std 802.1D, 1998 Edition) and the Rapid Spanning Tree Protocol (RSTP, Clause 17 of IEEE Std 802.1D, 1998 Edition) do not use a hop count and detect both circulating aged information and loss of connectivity to a neighboring bridge by means of Message Age and Max Age (maximum message age) parameters. To ensure compatibility MSTP increments Message Age for information received at the boundary of ~~an~~ a MST Region, discarding the information if necessary.

NOTE 1—MSTP's use of a separate hop count and message loss detection timer provides superior reconfiguration performance compared to STP and RSTP's use of Message Age and Max Age. Detection of loss of connectivity to a neighboring Bridge is not compromised by the need to allow for the overall diameter of the network, nor does the time allowed extend the number of hops permitted to aged recirculating information. Management calculation of the necessary parameters for custom topologies is also facilitated, as no allowance needs to be made for relative timer jitter and accuracy in different Bridges.

NOTE 2—The hop count applies to the spanning tree information for each tree. Message loss detection applies to all information transmitted by a given Bridge. The separate hop count can be more compactly encoded than the Message Age and Max Age timer values, and thus provides some per tree encoding efficiency in MST BPDUs.

<< Editor's note: in the last two notes and paragraph, MSTP is compared to RSTP as specified in the 1998 edition of IEEE Std 802.1D. The 2004 edition of the latter has changed the use of Max Age to reflect MSTP's hop count. The state machines and attendant definitions found later in this specification have been drafted to maintain Max Age even though it is now redundant with MSTP's hop count. Shall we keep the last two notes and paragraph and add a third note explaining the difference between the 1998 (actually 2001) and 2004 editions of RSTP wrt. Max Age? >>

13.16 Changing Port States

The Port State for each Bridge Port and spanning tree (CIST and MSTIs) is controlled by state machines whose goal is to maximize connectivity without introducing temporary data loops in the network. Root Ports, Master Ports, and Designated Ports are transitioned to the Forwarding Port State, and Alternate Ports and Backup Ports to the Discarding Port State, as rapidly as possible.

Transitions to the Discarding Port State can be simply effected without the risk of data loops. This clause describes the analysis used to determine the conditions for transitioning the Port State for a given spanning tree to Forwarding.

Starting with the assumption that any connected fragment of a network is composed of Bridges, Bridge Ports, and connected LANs that form a subtree of a spanning tree, this clause derives the conditions for transitioning ports with Root Port, Master Port, or Designated Port roles, such that the newly enlarged fragment continues to form either a subtree or the whole of the spanning tree. Since these conditions are applied every time a fragment is enlarged it is possible to trace the growth of a fragment from a single Bridge, which is clearly a consistent, if small, subtree of a spanning tree, to any sized fragment—thus justifying the initial assumption.

The requirement for consistent Port States in two subtrees, each bounded by Ports that either are not forwarding or are attached to LANs not attached to any other Bridge Port, can be met by waiting sufficient time for the priority vector information used to assign the Port Roles to reach all Bridges in the network. This ensures that these fragments of the potential active topology are not and are not about to be joined by other Forwarding Ports. However it can be shown that a newly selected Root Port can forward frames just as soon as prior recent root ports on the same bridge cease to do so, without further communication from other bridges. Rapid transitions of Designated Ports and Master Ports do require an explicit signal from the bridges and bridge ports in the connected subtrees. The Agreement mechanism is described, together with a Proposal mechanism that forces satisfaction of the conditions if they have not already been met by blocking Designated Ports connecting lower subtrees that are not yet in agreement. The same agreement mechanism is then used to transition the newly blocked ports back to forwarding, advancing the temporary cut in the active topology towards the edge of the network.

13.16.1 Subtree connectivity and priority vectors

Any given Bridge B, the LANs connected through its Forwarding Designated Ports, the further Bridges connected to those LANs through their Root Ports, the LANs connected to their Forwarding Designated Ports, and so on recursively, comprise a subtree S_B. Any LAN L that is part of S_B will be connected to B through a Forwarding Designated Port P_{CL} on a Bridge C also in S_B. L cannot be directly connected to any Port P_B on Bridge B unless B and C are one and the same, since the message priority vector for P_B is better than that of any Port of any other Bridge in S_B, and prior to Forwarding P_{CL} will have advertised its spanning port priority vector for long enough for it to receive any better message priority vector (within the design probabilities of protocol failure due to repeated BPDU loss) or will have engaged in an explicit confirmed exchange (see below) with all other Bridge Ports attached to that LAN.

13.16.2 Root Port transition to Forwarding

It follows from the above that B's Root Port can be transitioned to Forwarding immediately whether it is attached to a LAN in S_B or in the rest of the network, provided that all prior recent Root Ports on B (that might be similarly arbitrarily attached) have been transitioned to Discarding and the Root Port was not a Backup Port recently (B and C the same above).

13.16.3 Designated Port transition to Forwarding

On any given Bridge A, the Designated Port P_{AM} connected to a LAN M can be transitioned to Forwarding immediately provided that the message priority advertised by the Designated Port P_{CL} on any LAN L in any subtree S_{M1}, S_{M2},etc. connected to M is worse than that advertised by P_{AM}, that any bridge D attached to L has agreed that P_{CL} is the Designated Port, and only the Root Port and Designated Ports on D are Forwarding. A sufficient condition for P_{AM} to transition to Forwarding is that M is a point-to-point link attached to the Root Port P_{BM} of a Bridge B, that the port priority of P_{BM} is same as or worse than that of P_{AM}, and any port P_{BN} on B is Discarding or similarly attached to a Bridge C. P_{BM} signals this condition to P_{AM} by setting the Agreement flag in a Configuration Message carrying P_{BM}'s ~~port~~ designated priority and Port Role.

Figure 13-12 illustrates the generation of an Agreement at a Bridge's Root Port from an Agreement received or a Port State of Discarding at each of its Designated Ports, and a Port State of Discarding at each of its Alternate and Backup Ports. To solicit an Agreement each Designated Port that has been set to discard frames sends a Proposal. A Bridge receiving a Proposal transitions any Designated Port not already synchronized to Discarding, and solicits an Agreement by sending a Proposal in its turn.

NOTE 1—Agreements can be generated without prior receipt of a Proposal as soon as the conditions for the Agreement have been met. In that case subsequent receipt of a Proposal serves to elicit a further Agreement.

NOTE 2—If all Designated Ports have already been synchronized and the spanning priority vector received with the proposal does not convey worse information the synchronization is maintained and there is no need to transition Designated Ports to Discarding once more, or to transmit further Proposals.

13.16.4 Master Port transition to Forwarding

While the connectivity of the CIST from the CIST Regional Root through the Region to the rest of the CIST comprises a subtree rooted in the CIST Regional Root, the connectivity of the MSTI from the Master Port

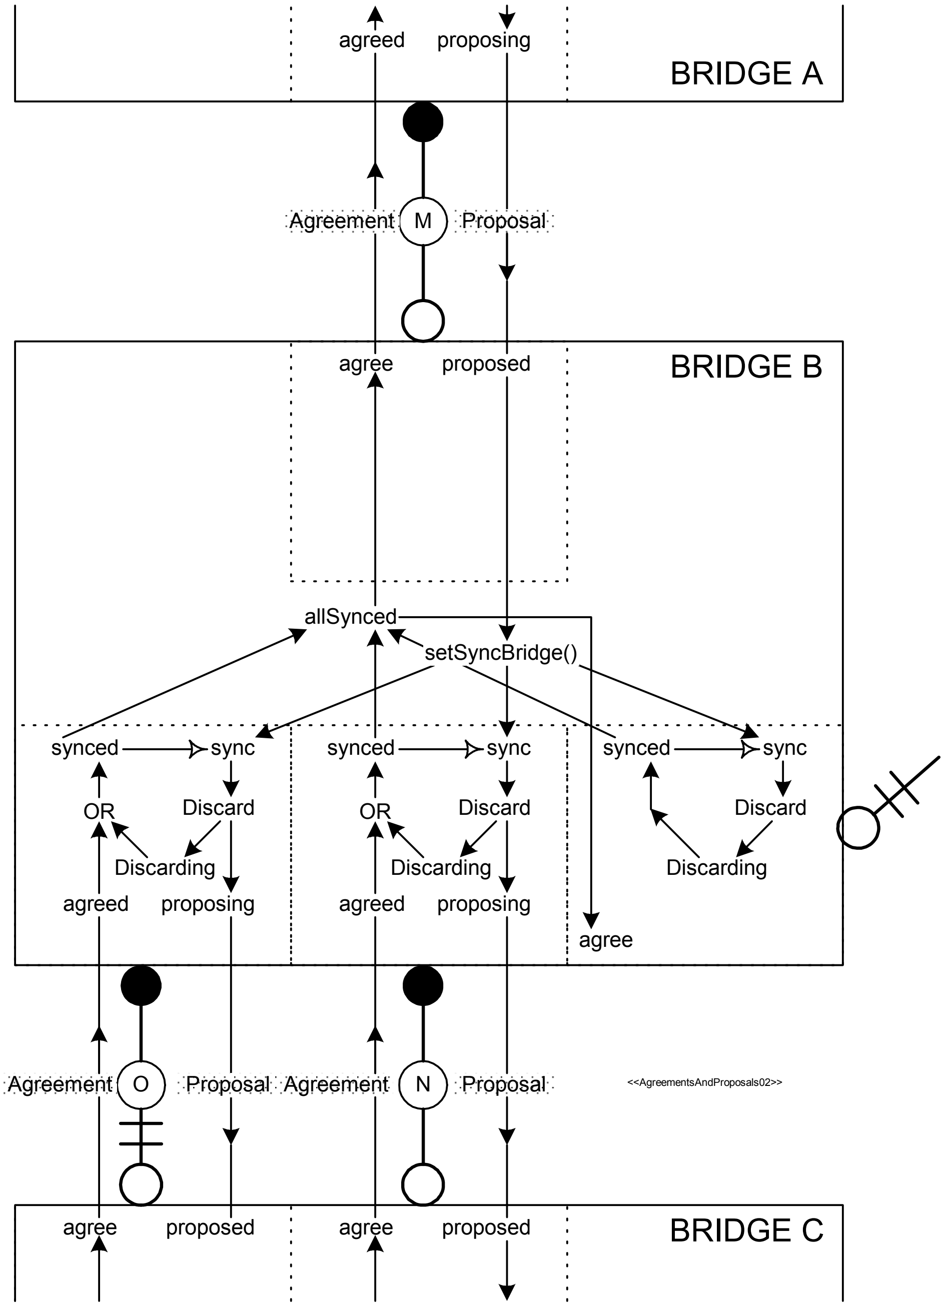

Figure 13-11—Agreements and Proposals

includes both a subtree below the CIST Regional Root and a subtree rooted in the MSTI Regional Root and connected to the CIST Regional Root by ~~an~~ a MSTI Root Port. Figure 13-14 illustrates this connectivity for both part of the CIST and ~~an~~ a MSTI through a Region in the example network of Figure 13-3. (In the example this latter subtree provides connectivity from the Master Port through LAN N to the subtree of the CIST outside the Region). Prior to the Master Port's transition to Forwarding it is possible that either MSTI subtree is providing connectivity to a prior Master Port. Before the Master Port can transition the connectivity of both subtrees has to agree with the new CIST Regional Root.

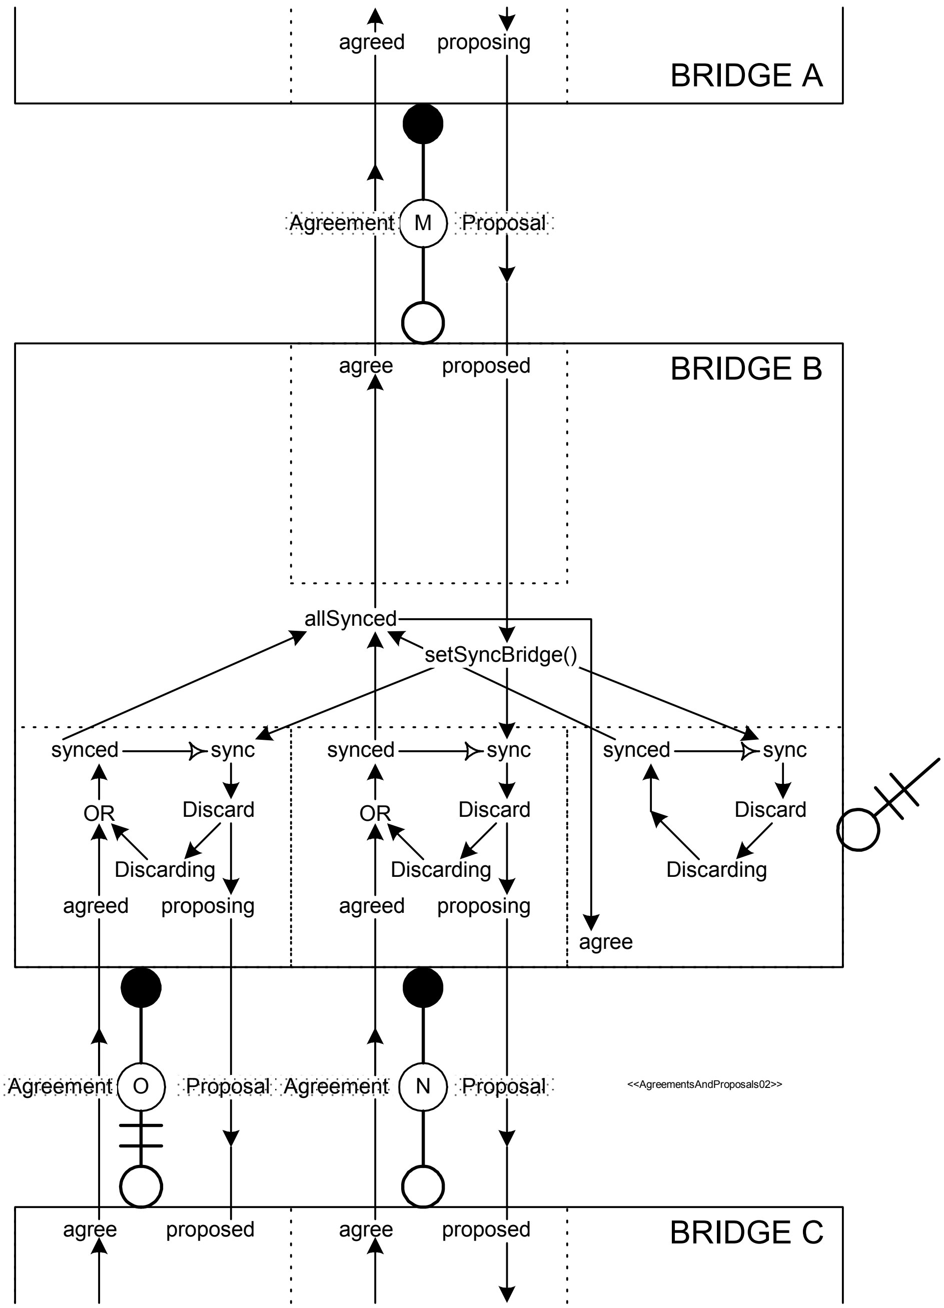

Figure 13-12—Agreements and Proposals

NOTE 1—The physical layout shown in the two halves of Figure 13-14 differs in order to reflect the different priorities and logical topologies for the two spanning tree instances. The layout convention used is that designated Ports are shown as horizontal lines, root Ports as vertical lines, and alternate Ports as diagonal lines.

Figure 13-16 illustrates the extension of the Agreement mechanism to signal from Designated Ports to Root Ports as well as vice versa. To ensure that ~~an~~ a MSTI does not connect alternate Master Ports, an Agreement is only recognized at ~~an~~ a MSTI Port when the CIST Regional Root associated with the information matches that selected by the receiving port. Proposals, eliciting Agreements, necessarily flow from Designated Ports to Root Ports with the propagation of spanning tree information so a new CIST Regional Root cannot

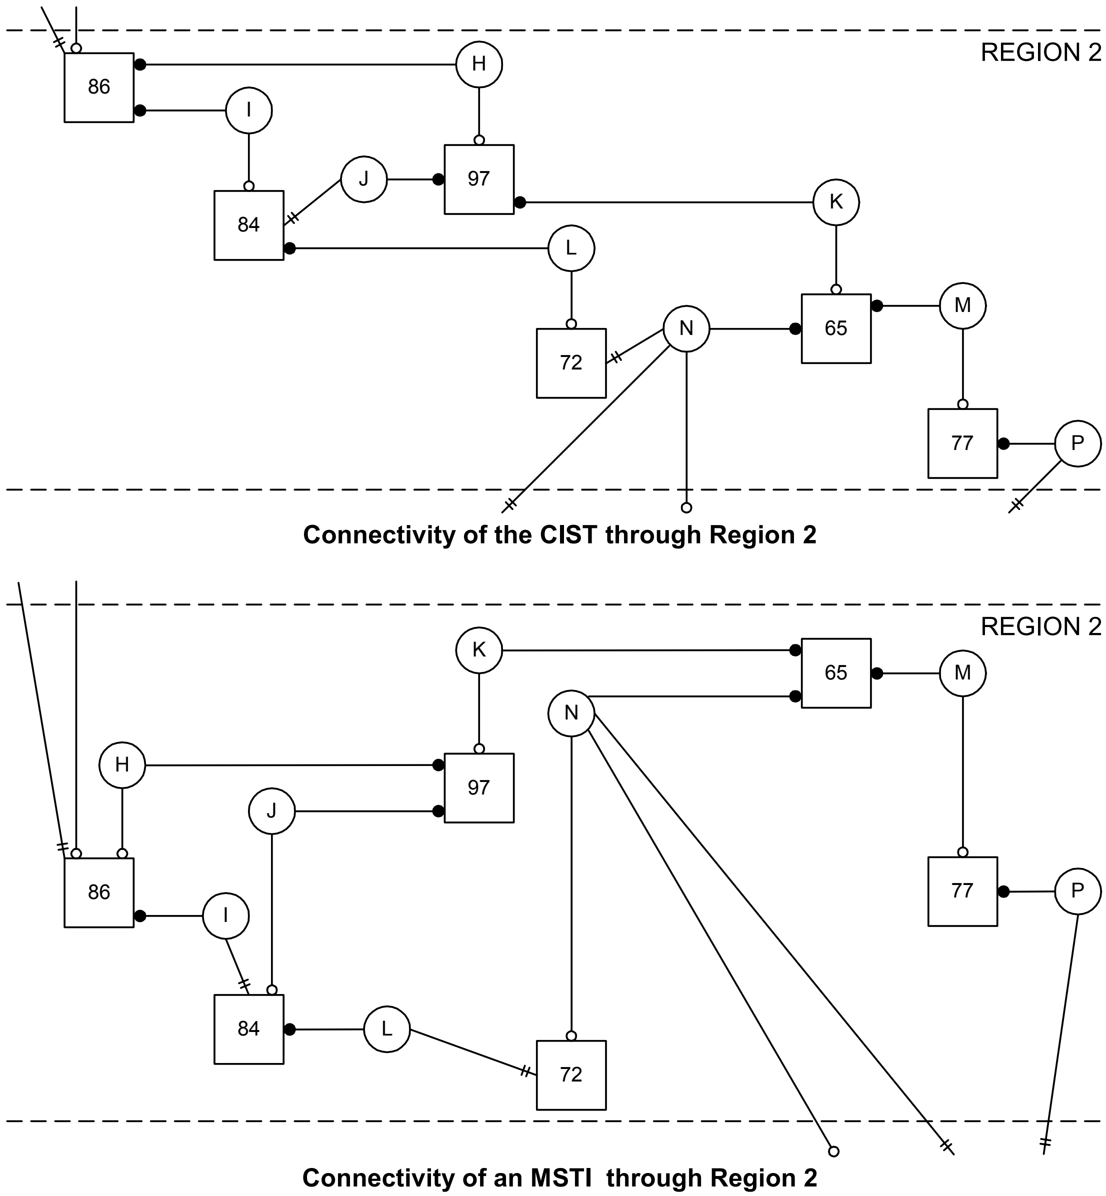

Figure 13-13—CSTI and MSTI Active Topologies in a Region

transmit a Proposal directly on its MSTI Root Ports. However updating of a CIST Designated Port's port priority vector with a new Regional Root Identifier forces the port to discard frames for all MSTIs, thus initiating the Proposal from the first Bridge nearer the MSTI Regional Root that learns of the new Regional Root.

When an Agreement, A_{MR}, is sent by a Root Port P_{MR} on a Regional Root M it attests that the CIST Root Identifier and External Root Path Cost components of the message priority advertised on all LANs connected to the CIST by P_{MR} through M are the same as or worse than those accompanying A_{MR}. The connectivity provided by each MSTI can be independent of that provided by the CIST within the MST Region, and can therefore connect P_{MR} and one or more CIST Root Ports external to but attached at the boundary of the region even as CIST connectivity within the region is interrupted in order to satisfy the conditions for generating A_{MR}. The Agreement cannot therefore be generated unless all MSTI subtrees as

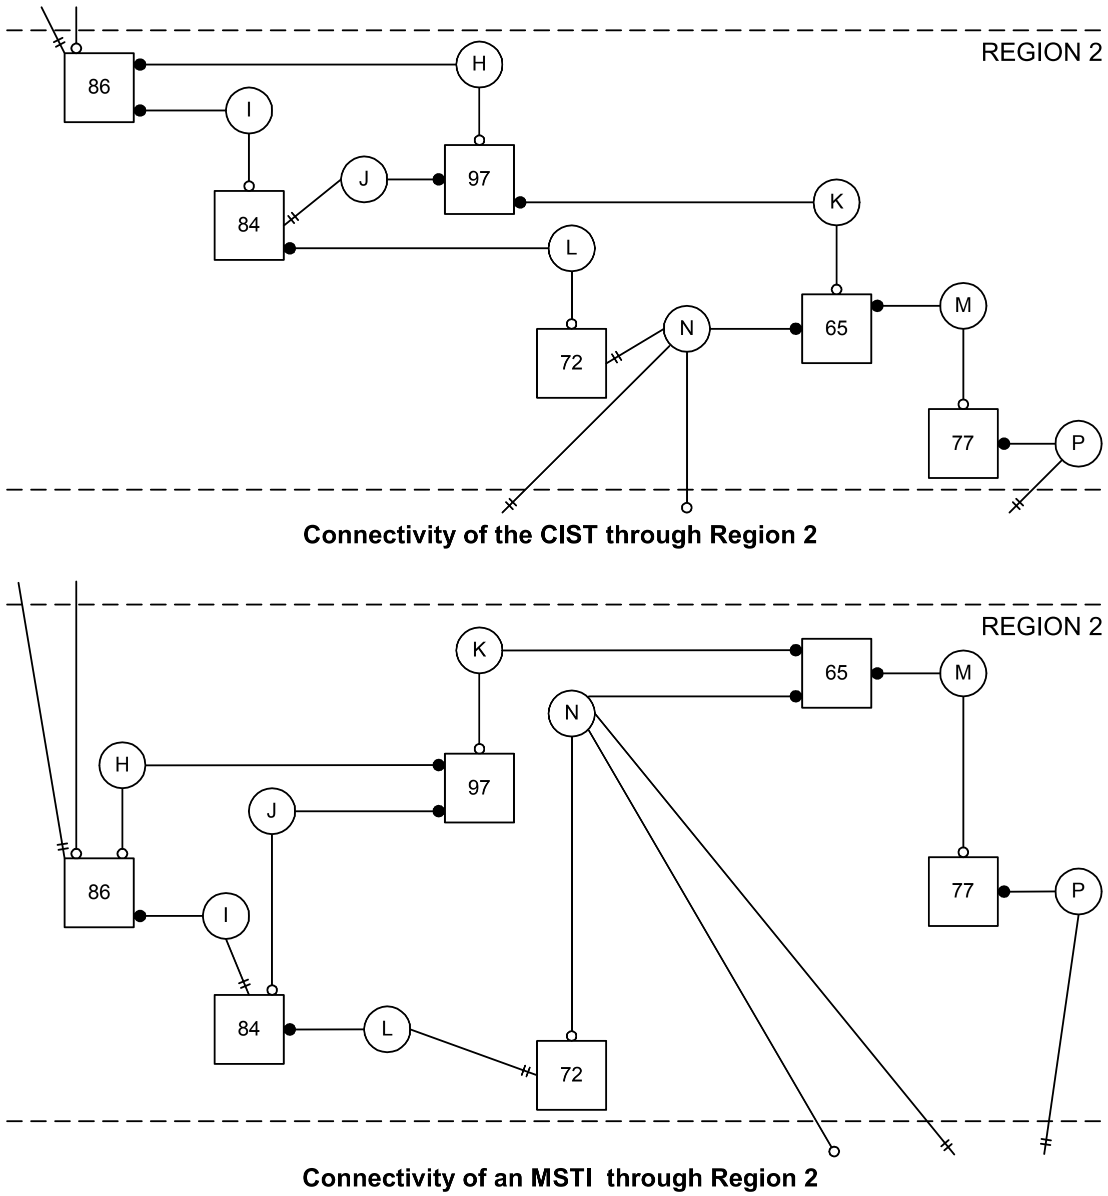

Figure 13-14—CSTI and MSTI Active Topologies in a Region

well as the CIST subtree internal to the Region are in Agreement. To ensure that ~~an~~ a MSTI does not connect to a CIST subtree external to the Region that does not meet the constraints on the CST priority vector components, an Agreement received at ~~an~~ a MSTI Designated Port from a Bridge Port not internal to the Region is only recognized if the CIST Root Identifier and External Root Path Cost of the CIST root priority vector selected by the transmitting Bridge Port are equal to or worse than those selected by the receiver. Updating of a CIST Designated Port's port priority vector with a worse CIST Root Identifier and External Root Path Cost forces the port to discard frames for all MSTIs, thus initiating a Proposal that will elicit agreement.

NOTE 2—MSTI Designated Ports are forced to discard frames, as required above, through the following state machine mechanisms. The CIST Port Information machine sets the 'sync' variable for all MSTIs on a transition into the UPDATE state if updating the port priority with the designated priority changes the Regional Root Identifier or replaces the CIST Root Identifier or External Path Cost with a worse tuple.The Port Role Transition machine acts on the 'sync', instructing

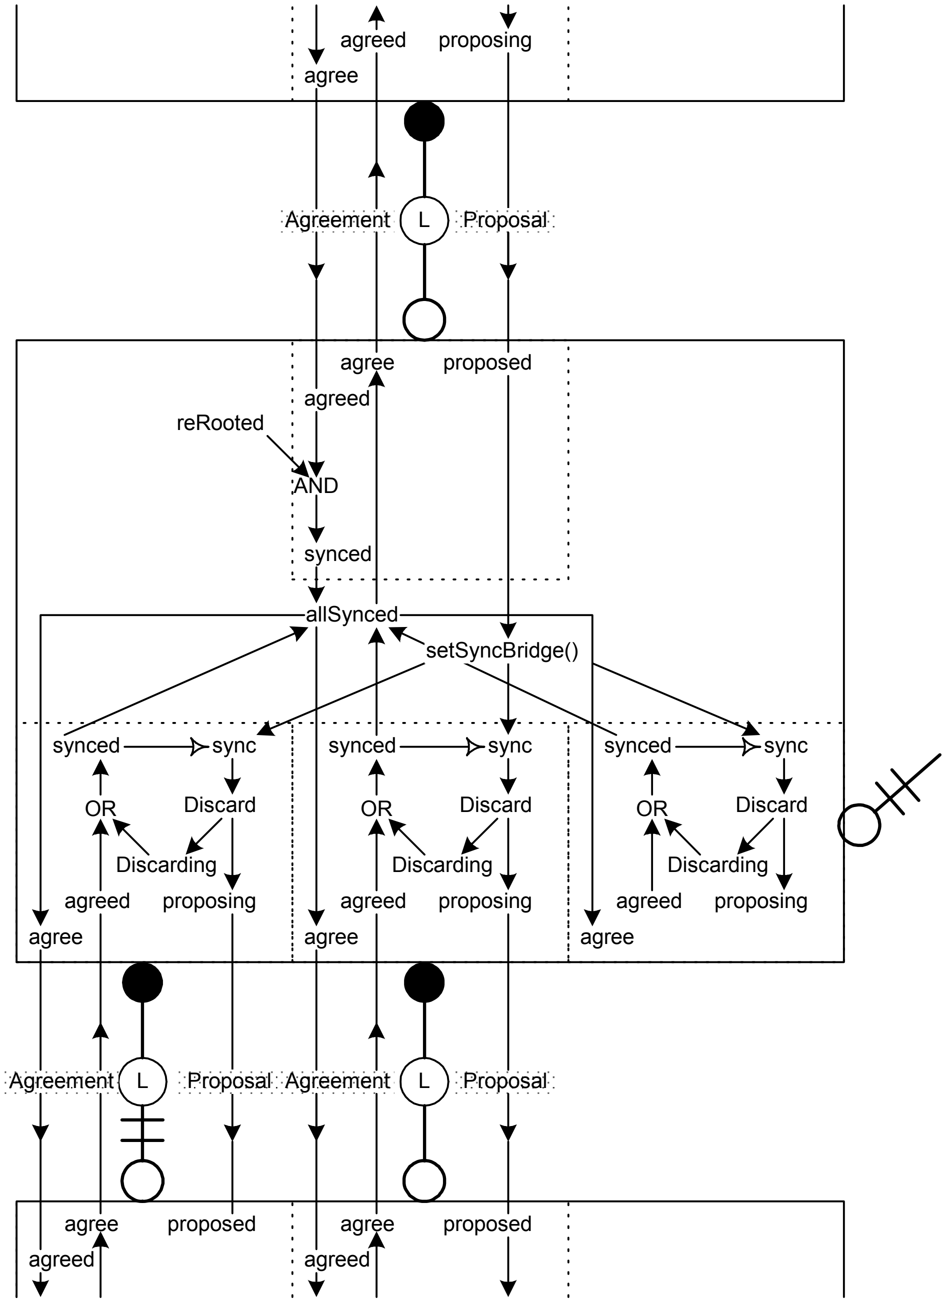

Figure 13-15—Enhanced Agreements

the port to discard frames, and setting ‘synced’ and cancelling ‘sync’ when the port is discarding or an agreement is received.

NOTE 3—A ‘cut’ in ~~an~~ a MSTI can be transferred to the CST, either at a Designated Port attached to the same LAN as ~~an~~ a STP Bridge, or at the Root Port of a Bridge in an adjacent Region. However if the CST priority components have already been ‘synced’, as they mostly likely will have if the original cut was caused by changes in physical topology within the Region, the cut will terminate there. Otherwise the transferred cut precedes a cut in the CIST, and the synced port may terminate the latter. In that way cuts in the CST will proceed through ~~an~~ a MST Region by the quickest tree that will carry them.

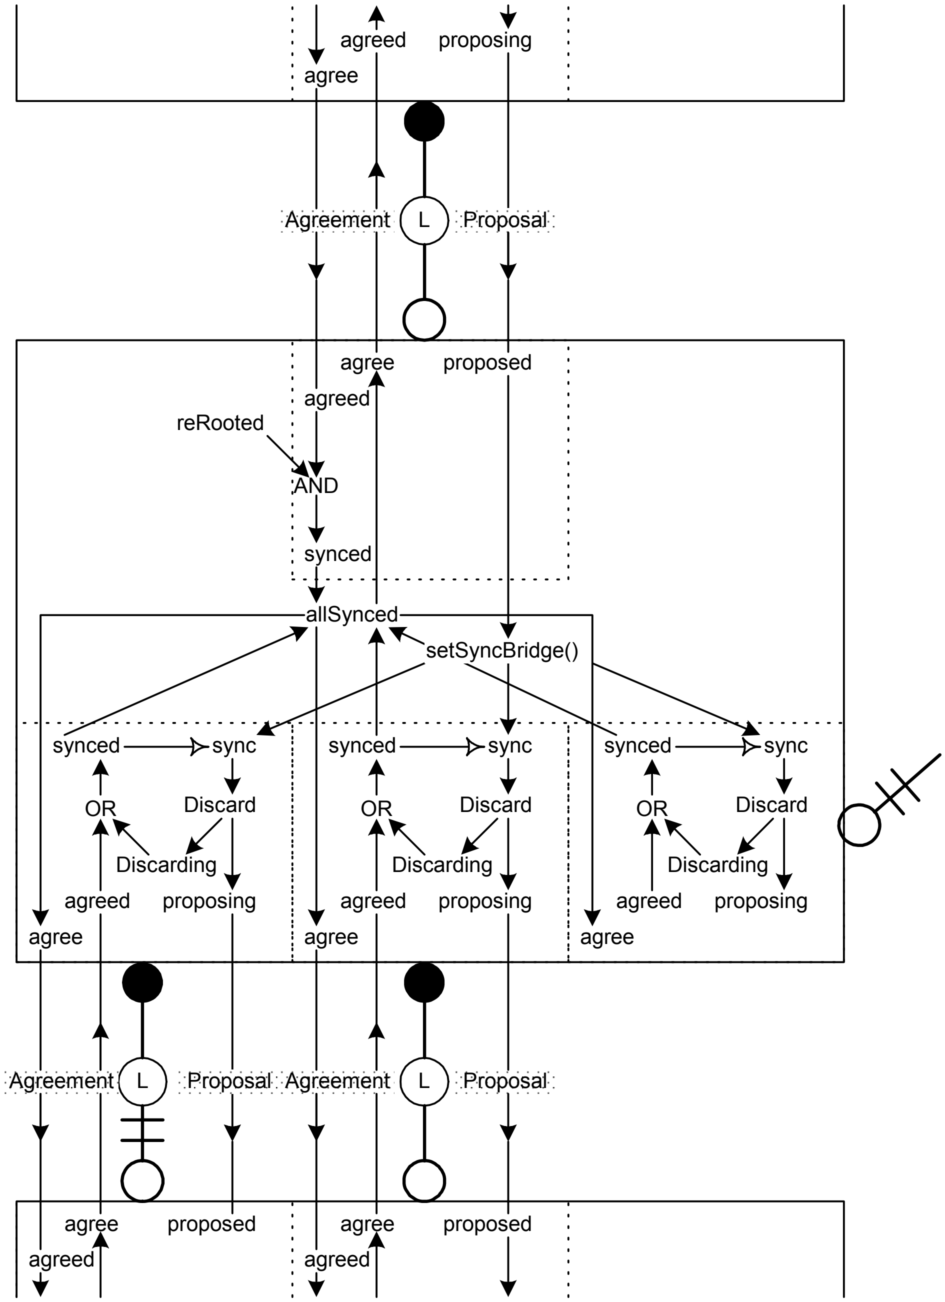

Figure 13-16—Enhanced Agreements

NOTE 4—In the important topology where the CIST Root Bridge is chosen to be within ~~an~~ a MST Region, cuts are not transferred from the CIST to any MSTI in that Region. Thus the propagation of cuts in the CIST will not disrupt MSTI connectivity in the Region.

NOTE 5—Topology change detection by MSTI state machines is based on changes from Root, Master, and Designated Port Roles to Alternate, Backup, Disabled Port Roles, not on changes in Port State. Propagating 'cuts' designed to prevent temporary loops through a Region does not therefore require unnecessary changes to Filtering Databases with attendant temporary flooding of frames.

<< Editor's note: section to be updated. >>

13.17 Updating Learned Station Location Information

A spanning tree reconfiguration can cause end stations to appear to move from the point of view of any given Bridge, even if that Bridge's Port States do not change, and is signaled from the Bridge whose Port Roles have changed to others using TCN messages. MST BPDUs encode separate TCN messages for the CIST and each MSTI, and MSTP supports the optimizations specified for RSTP (Clause 17.~~10~~ 11 of IEEE Std 802.1D, ~~1998~~ 2004 Edition). Together these facilitate removal of entries for the minimum set of Ports from the Filtering Databases associated with the spanning trees whose active topology has changed. In addition MSTP only detects topology changes following a change of Port Role to Root Port, Master Port, or Designated Port. Temporary cuts in the active topology, introduced to ensure that rapid Port State transitions to Forwarding do not cause loops, and do not therefore cause Filtering Database entries to be flushed throughout the network, unless they are accompanied by Port Role changes.

<< Editor's note: regarding the last two sentences, the first one should instead state that MSTP provides for Master Ports the same topology change detection mechanism that RSTP (as specified in 17.11 of IEEE Std 802.1D, 2004 Edition) provides for Root and Designated Ports. The second sentence is gramatically incorrect and might actually be replaced to reference Clause 17.11 of IEEE Std 802.1D, 2004 Edition so as to state that MSTP offers the same capabilities as RSTP wrt. updating learned station location information. >>

Changes in the active topology of any MSTI do not change end station locations for the CIST or any other MSTI, unless the underlying changes in the physical topology that gave rise to the reconfiguration also cause those trees to reconfigure. Changes to the CST, i.e. the connectivity provided by the CIST between MST Regions, can cause end station location changes for all trees. Changes to an IST can cause CST end station location changes, but do not affect MSTIs in that Region unless those trees also reconfigure.

NOTE 1—The shorthand terms "end station locations for a given tree", "the CST", and "an IST", are used to mean "the apparent location of end stations as recorded by filtering databases associated with the given tree", "the connectivity provided by the CIST between and not internal to MST Regions", and "the connectivity provided by the CIST internal to a given MST Region" respectively.

On receipt of a CIST TCN Message from a Bridge Port not internal to the Region, or on a change in Port Role for a Bridge Port not internal to the Region, TCN Messages are transmitted through each of the other Ports of the receiving Bridge for each MSTI, and the Filtering Databases for those ports are flushed.

NOTE 2—TCN Messages for the CIST are always encoded in the same way, irrespective of whether they are perceived to have originated from topology changes internal to the Region or outside it. This allows RSTP Bridges whose Root Ports attach to a LAN within ~~an~~ a MST Region to receive these TCN Messages correctly.

NOTE 3—The Port receiving a CIST TCN Message from another Bridge Port external to the Region can be a Master Port, a Designated Port attached to the same LAN as ~~an~~ a STP Bridge, or a Designated Port attached to a LAN that is within the Region but is attached to by the Root Ports of Bridges in other Regions.

<< Editor's note: section to be updated. >>

13.18 MSTP and point-to-point links

MSTP uses the adminPointToPointMAC and operPointToPointMAC parameters (6.4.3 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) to allow the point-to-point status of LANs to be manipulated administratively, and the operational state to be signalled to the MSTP state machines. This use, paralleling that of RSTP (17.~~11~~ 12 of IEEE Std 802.1D, ~~1998~~ 2004 Edition), facilitates use of the Agreement mechanism (13.16) to enable rapid Forwarding Port State transitions.

13.19 Multiple Spanning Tree State Machines

The operation of the Multiple Spanning Tree Protocol is represented by the following common set of state machines:

a) A Port Timers state machine for the Bridge (13.28)
b) A Port ~~Timers~~ Protocol Migration state machine for ~~the Bridge~~ each Port (~~13.28~~13.30)
c) A Port ~~Protocol Migration~~ Receive state machine for each Port (~~13.30~~13.29)
d) A Port ~~Receive~~ Transmit state machine for each Port (~~13.29~~13.32)
e) A ~~Port Transmit~~ Bridge Detection state machine for each Port (~~13.32~~13.31)

with the following set for the CIST and each MSTI:

f) A Port Information state machine for each Port (13.33)
g) A Port Role Selection state machine for the Bridge (13.34)
h) A Port Role Transitions state machine for each Port (13.35)
i) A Port State Transition state machine for each Port (13.36)
j) A Topology Change state machine for each Port (13.37)

The operation of each state machine and its associated variable and procedural definitions is specified in detail below. Each is modeled on the corresponding state machine for RSTP as described in Clause 17 of IEEE Std 802.1D, ~~1998~~ 2004 Edition. Modifications:

a) support both the CIST and the MSTIs
b) use the extended spanning tree priority vector definition and calculations for the CIST (13.10)
c) provide communication between the CIST state machines and the MSTI state machines as required by 13.16, 13.19, and 13.17.
d) enhance the operation of the topology change state machine to avoid unnecessary removal of end station location information following port state transitions without changes in port roles.

All references to named variables or procedures in the specification of the state machines are to those corresponding to the instance of the state machine using the function unless explicit reference is made to the CIST or given MSTIs.

<< Editor's note: The previous statement should add that, if no explicit reference is made to the CIST or given MSTIs, then references are to variables or procedures corresponding to the CIST. This would suppress ambiguity in txConfig(), for example, where the value of tcWhile has to be used without any reference to the appropriate tree. (txConfig() is used from PTX, see http://www.ieee802.org/1/private/email/msg01392.html.) >>

For further economy of specification, the names of CIST and MSTI conditions ~~or procedures~~ providing the same general functionality but differing in detail incorporate the prefix or suffix "Cist" or "Msti". These are substituted for names beginning or ending in "Xst" in a single state machine description common to the CIST and any given MSTI.

~~NOTE 1—The specification of RSTP does not use a distinct Port Receive state machine, but combines the functionality into the Port Information state machine. Since separate instances of the latter are required for each tree, each operating on the messages contained in a single BPDU, it is convenient to separate out the aspects of reception that are purely per Port and per BPDU in this specification into a separate Port Receive state machine.~~

NOTE ~~2~~1—Individual MSTIs do not implement their own Port Transmit state machine, but signal the need to transmit to a single Port Transmit State machine by setting the newInfoMsti variable. The procedures used by this machine transmit information for the CIST and all MSTIs in a single BPDU.

~~NOTE 3—A detailed list of the differences between the RSTP state machines specified in IEEE Std 802.1D, 1998 Edition as amended by IEEE Std 802.1D, 1998 Edition and the state machines specified in this standard is presented in Annex G.~~

<< Editor's note: shall we have more text here to explain that procedures know by implementation (the same way that "cist" in 13.26.2 is "TRUE only for CIST state machines, i.e. FALSE for MSTI state machine instances") on which tree (CIST or any given MSTI) they apply? Please note that the definition of "cist" in previous edition of this text does not rely on the definition of a "this" variable that would take either of two enumerated values, 'cist' or 'msti'. >>

Figure 13-18 illustrates the state machines, their state variables and communication between state machines. This overview diagram is not itself a state machine, but serves to illustrate the principal variables that are used to communicate between the individual machines and the variables local to each machine Figure 13-20 describes its notation.

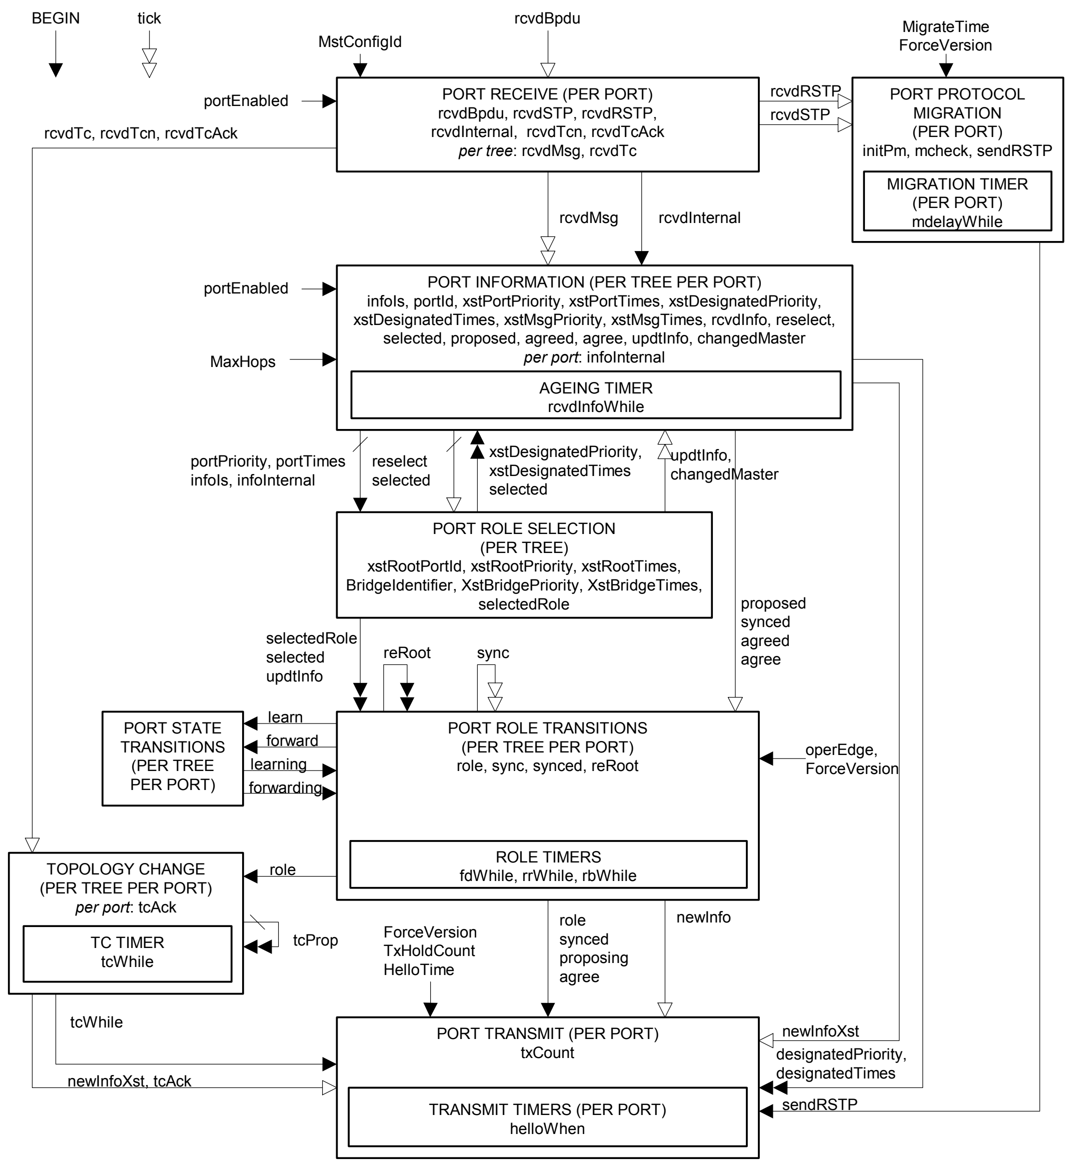

NOTE: For convenience all timers are collected together into one state machine.

Figure 13-17—MSTP state machines—overview and relationships

<< Editor's note: following diagram to be updated. >>

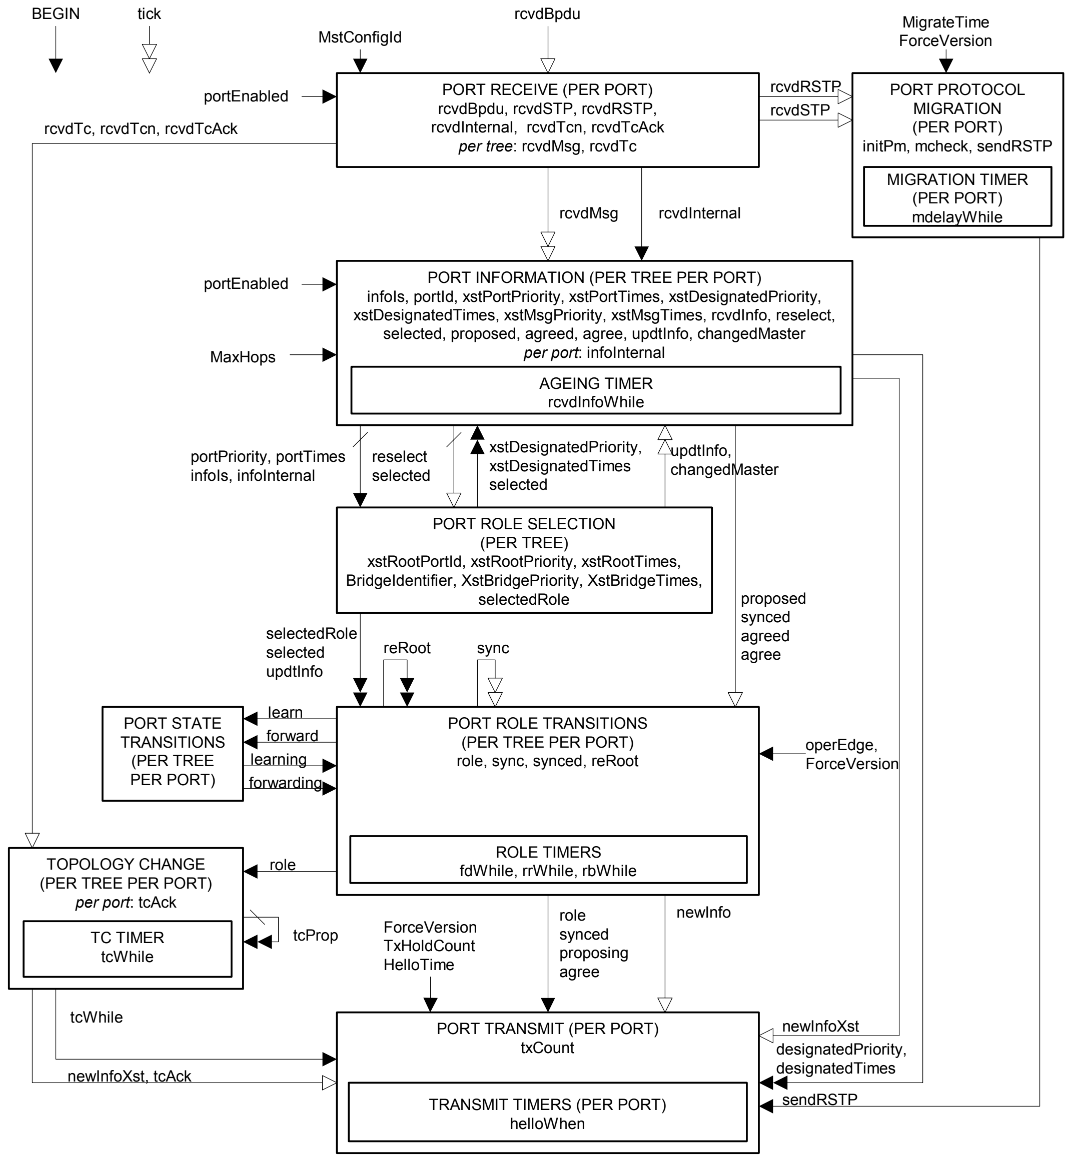

NOTE: For convenience all timers are collected together into one state machine.

Figure 13-18—MSTP state machines—overview and relationships

13.20 Notational Conventions used in State Diagrams

The notational conventions used in the specification of MSTP are identical to those used in the specification of RSTP and defined in 17.~~14~~ 16 of IEEE Std 802.1D, ~~1998~~ 2004 Edition.

13.21 State Machine Timers

Each of the state machine timers are as specified in 17.~~15~~ 17 of IEEE Std 802.1D, ~~1998~~ 2004 Edition.

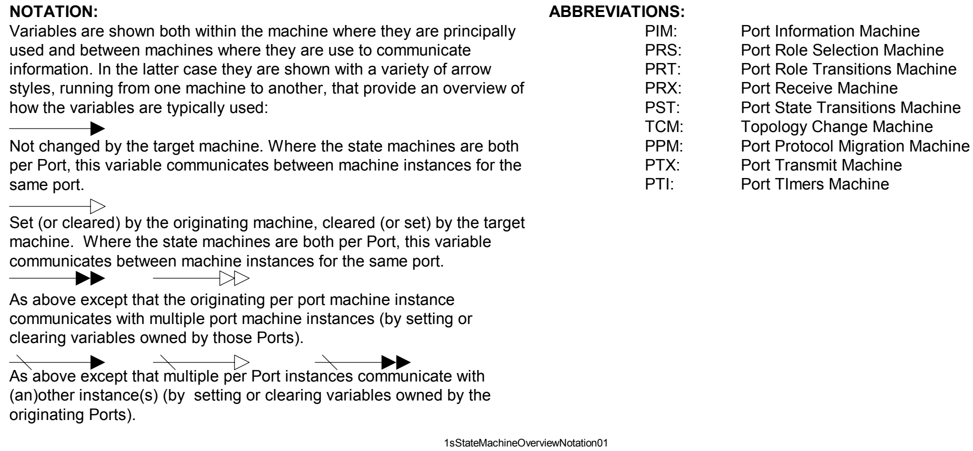

Figure 13-19—MSTP overview notation

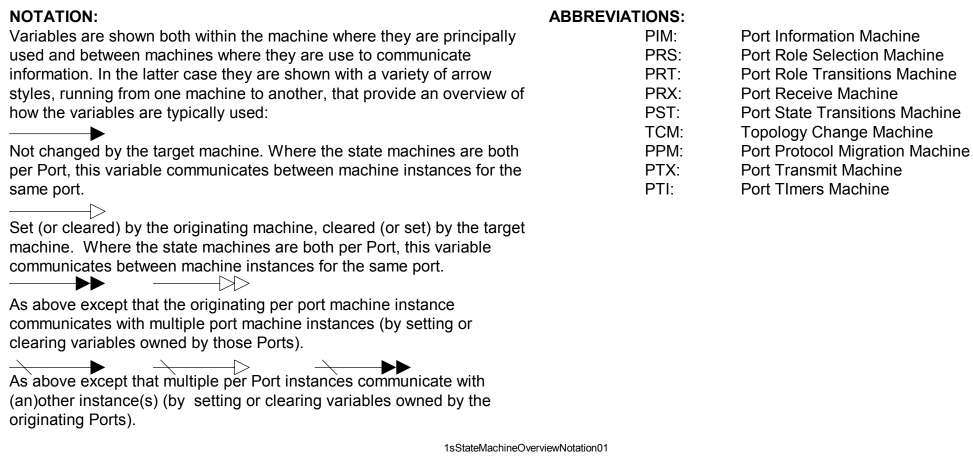

Figure 13-20—MSTP overview notation

One instance of the following shall be implemented per-Port:

a) mdelayWhile
b) helloWhen
c) edgeDelayWhile

One instance per-Port of the following shall be implemented for the CIST and one per-Port for each MSTI:

d) fdWhile
e) rrWhile
f) rbWhile
g) tcWhile
h) rcvdInfoWhile

13.22 ~~State Machine~~ MSTP Performance Parameters

~~These parameters are treated as constants by the CIST and MSTI state machines; their values can be modified only by management action.~~

These parameters are not modified by the operation of MSTP, but are treated as constants by the MSTP state machines and the associated variables (13.23, 13.24) and conditions (13.26). They may be modified by management.

The following parameters are as specified in 17.~~16~~ 13 of IEEE Std 802.1D, ~~1998~~ 2004 Edition for RSTP. A single value of each parameter applies to the MST Bridge as a whole, including all Ports and all CIST and MSTI state machines.

~~a) ForceVersion~~
~~b) FwdDelay~~
~~c) TxHoldCount~~
~~d) MigrateTime~~
a) Force Protocol Version
b) Bridge Forward Delay
c) Transmit Hold Count
d) Migrate Time
e) Bridge Max Age

The following parameter is as specified in 17.~~16~~ 13 of IEEE Std 802.1D, ~~1998~~ 2004 Edition for RSTP, but may be managed separately for each Port.

~~f) HelloTime~~
g) Bridge Hello Time
h) Admin Edge Port
i) Ageing Time
j) AutoEdge
k) Port Identifier Priority

The following parameter is additional to those specified for RSTP. A single value applies to all Spanning Trees within ~~an~~ a MST Region (the CIST and all MSTIs) for which the Bridge is the Regional Root.

l) MaxHops

<< Editor's note: MaxHops has no definition in the current specification of MSTP. >>

The following parameter is as specified for RSTP, but has been renamed for clarity in this specification. One value per-Port applies to the CIST.

m) ExternalPortPathCost (specified as PortPathCost in 17.~~16~~ 13 of IEEE Std 802.1D, ~~1998~~ 2004 Edition)

The following parameter is additional to those specified for RSTP, and may be managed separately for the CIST and for each MSTI per-Port.

n) InternalPortPathCost

13.23 Per-Bridge Variables

Per-bridge variable(s) perform the functions described in 17.~~17~~ 18 of IEEE Std 802.1D, ~~1998~~ 2004 Edition, but have enhanced or extended specifications or considerations.

A single instance of each of the following variables applies to the CIST and to all MSTIs.

a) BEGIN (13.23.1)
b) MstConfigId (13.23.8)

NOTE—MstConfigId is not specified in 17.~~17~~ 18 of IEEE Std 802.1D, ~~1998~~ 2004 Edition.

~~There is one instance per-Bridge of each of the following for the CIST, and one for each MSTI.~~

~~c) BridgeIdentifier (13.23.2)~~

~~And~~ There is one instance per-Bridge of each of the following for the CIST, and one for each MSTI.

~~d) CistBridgePriority (13.23.3)~~
~~e) CistBridgeTimes (13.23.4)~~
~~f) cistRootPortId (13.23.5)~~
~~g) cistRootPriority (13.23.6)~~
h) ~~cistRootTimes~~ BridgeIdentifier (~~13.23.7~~13.23.2)

~~And one instance per-Bridge of each of the following for each MSTI.~~

i) ~~MstiBridgePriority~~ BridgePriority (~~13.23.9~~13.23.3)
j) ~~MstiBridgeTimes~~ BridgeTimes (~~13.23.10~~13.23.4)
k) ~~mstiRootPortId~~ rootPortId (~~13.23.11~~13.23.5)
l) ~~mstiRootPriority~~ rootPriority (~~13.23.12~~13.23.6)
m) ~~mstiRootTimes~~ rootTimes (~~13.23.13~~13.23.7)

13.23.1 BEGIN

~~This variable is~~ A boolean controlled by the system initialization process. ~~A value of~~ If TRUE causes all CIST and MSTI state machines, including per Port state machines, to ~~transit to~~ continuously execute their initial state. ~~A value of~~ If FALSE allows all state machines to perform transitions out of their initial state, in accordance with the relevant state machine definitions.

Changes to any of the following parameters cause BEGIN to be asserted for the state machines for the Bridge, for all trees, and for each Port:

a) The MST Configuration Identifier~~.~~,
b) Force Protocol Version.

13.23.2 BridgeIdentifier

The unique Bridge Identifier assigned to this Bridge for this tree (CIST or MSTI).

The 12-bit system ID extension component of a Bridge Identifier (9.2.5 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) shall be set to zero for the CIST, and to the value of the MSTID for ~~an~~ a MSTI, thus allocating distinct Bridge Identifiers to the CIST and each MSTI all based on the use of a single Bridge Address component value for the MST Bridge as a whole.

NOTE—This convention is used to convey the MSTID for each MSTI Configuration Message encoded in ~~an~~ a MST BPDU.

The four most significant bits of the Bridge Identifier (the settable Priority component) for the CIST and for each MSTI can be modified independently of the setting of those bits for all other trees, as a part of allowing full and independent configuration control to be exerted over each Spanning Tree instance.

13.23.3 ~~CistBridgePriority~~BridgePriority

For the CIST, the value of the CIST bridge priority vector, as defined in 13.10. The CIST Root Identifier, CIST Regional Root Identifier, and Designated Bridge Identifier components are all equal to the value of the CIST Bridge Identifier. The remaining components (External Root Path Cost, Internal Root Path Cost, Designated Port Identifier) are set to zero.

~~The~~ For a given MSTI, the value of the ~~CIST~~ MSTI bridge priority vector, as defined in ~~13.10~~13.11. The ~~CIST Root Identifier, CIST~~ MSTI Regional Root ~~Identifier,~~ Identifier and Designated Bridge Identifier components are ~~all~~ equal to the value of the ~~CIST~~ MSTI Bridge ~~Identifier~~Identifier (13.23.2). The remaining components (~~External Root Path Cost,~~ MSTI Internal Root Path Cost, Designated Port Identifier) are set to zero.

~~CistBridgePriority~~ BridgePriority is used by ~~updtCistRolesBridge~~updtRolesTree() in determining the value of the ~~cistRootPriority~~ rootPriority variable (see 13.23.6).

13.23.4 ~~CistBridgeTimes~~BridgeTimes

~~CistBridgeTimes comprises:~~

For the CIST, BridgeTimes comprises:

a) The current values of Bridge Forward Delay and Bridge Max Age (see Table 17-~~5~~ 1 of IEEE Std 802.1D, ~~1998~~ 2004 Edition). These parameter values are determined only by management;
b) A Message Age value of zero.
c) The current value of MaxHops (13.22).

For a given MSTI, BridgeTimes is composed of a single component:

a) The current value of Bridge Max Hops (13.22). This parameter value is determined only by management.

~~CistBridgeTimes~~ BridgeTimes is used by ~~updtCistRolesBridge~~updtRolesTree() in determining the value of the ~~cistRootTimes~~ rootTimes variable (13.23.7).

13.23.5 ~~cistRootPortId~~rootPortId

~~The~~ For the CIST, the Port Identifier of the Root Port, and a component of the CIST root priority vector (13.10).

For a given MSTI, the Port Identifier of the Root Port, and a component of the root priority vector (13.11).

13.23.6 ~~cistRootPriority~~rootPriority

~~The~~ For the CIST, the CIST Root Identifier, CIST External Root Path Cost, CIST Regional Root Identifier, CIST Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Bridge's CIST root priority vector (13.10).

For a given MSTI, the MSTI Regional Root Identifier, MSTI Internal Root Path Cost, MSTI Designated Bridge Identifier, and MSTI Designated Port Identifier components of the Bridge's root priority vector (13.11).

13.23.7 ~~cistRootTimes~~rootTimes

~~The~~ For the CIST, the Bridge's timer parameter values (Message Age, Max Age, Forward Delay, and remainingHops). The values of these timers are derived from the values stored in ~~cistPortTimes~~ the CIST's portTimes parameter (17.~~18~~19.~~8~~ 22 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) for the Root Port. Max Age and Forward Delay are set equal to the values held by the Root Port. If the CIST root priority vector was received from a Bridge in a different MST Region (infoInternal FALSE) Message Age is the value held by the Root Port incremented by ~~the greater of (~~1~~/16 Max Age) and 1~~, rounded to the nearest whole second (see 17.~~19~~21.~~21~~ 25 of IEEE Std 802.1D, ~~1998~~ 2004 Edition), and remainingHops is set equal to MaxHops. Otherwise, if the CIST root priority vector was received from a Bridge in the same MST Region or the Bridge is itself the CIST Root (rcvdInternal TRUE), Message Age is the value held by the Root Port (in ~~cistRootTimes~~rootTimes) and remainingHops is the value held by the Root Port minus one.

For a given MSTI, the value of remainingHops as stored in the MSTI's portTimes parameter (13.24.20) for the MSTI Root Port, derived from the value received in the MSTI's msgTimes parameter by subtracting one.

13.23.8 MstConfigId

The value of the MST Configuration Identifier (13.7) corresponding to the Bridge's current MST Region Configuration.

~~13.23.9 MstiBridgePriority~~

~~The value of the MSTI bridge priority vector for a given MSTI (13.11). The MSTI Regional Root Identifier and Designated Bridge Identifier components are equal to the value of the MSTI Bridge Identifier (13.23.2). The remaining components (MSTI Internal Root Path Cost, Designated Port Identifier) are set to zero.~~

~~MstiBridgePriority is used by updtMstiRolesBridge() in determining the value of the mstiRootPriority variable (see 13.23.12).~~

~~13.23.10 MstiBridgeTimes~~

~~MstiBridgeTimes for a given MSTI is composed of a single component:~~

~~a) The current value of Bridge Max Hops (13.22). This parameter value is determined only by management.~~

~~MstiBridgeTimes is used by updtMstiRolesBridge() in determining the value of the mstiRootTimes variable (see 13.23.13).~~

~~13.23.11 mstiRootPortId~~

~~The Port Identifier of the Root Port for a given MSTI, and a component of the root priority vector (13.11).~~

~~13.23.12 mstiRootPriority~~

~~The MSTI Regional Root Identifier, MSTI Internal Root Path Cost, MSTI Designated Bridge Identifier, and MSTI Designated Port Identifier components of the Bridge's root priority vector (13.11) for a given MSTI.~~

~~13.23.13 mstiRootTimes~~

~~The value of remainingHops for a given MSTI as stored in mstiPortTimes (13.24.20) for the MSTI Root Port, derived from the value received in mstiMsgTimes by subtracting one.~~

13.24 Per-Port Variables

The following variables perform the function specified in 17.~~17~~ 19 of IEEE Std 802.1D, ~~1998~~ 2004 Edition. A single per-Port instance applies to the CIST and to all MSTIs.

a) ageingTime
b) fdbFlush
c) operEdge
d) portEnabled
e) tick
f) txCount
~~g) operEdge~~
~~h) portEnabled~~

A single per-Port instance of the following variable(s) not specified in IEEE Std 802.1D, ~~1998~~ 2004 Edition applies to the CIST and or all MSTIs.

i) infoInternal (13.24.10)
j) ~~newInfoCist~~ rcvdInternal (~~13.24.21~~13.24.28)
k) ~~newInfoMsti~~ restrictedRole(~~13.24.22~~13.26.16)
l) ~~rcvdInternal~~ restrictedTcn(~~13.24.28~~13.26.17)

A single per-Port instance of the following variable(s) not specified in IEEE Std 802.1D, 2004 Edition applies to all MSTIs.

m) newInfoMsti (13.24.22)

The following variables perform the function specified in 17.~~17~~ 19 of IEEE Std 802.1D, ~~1998~~ 2004 Edition. A single per-Port instance is used by all state machines.

~~n) initPm~~
o) mcheck
p) rcvdBpdu
q) rcvdRSTP
r) rcvdSTP
s) rcvdTcAck
t) rcvdTcn
u) sendRSTP
v) tcAck

The following variables are as specified in 17.~~17~~ 19 of IEEE Std 802.1D, ~~1998~~ 2004 Edition. There is one instance per-Port of each variable for the CIST, and one per-Port for each MSTI.

w) agree
x) disputed
y) forward
z) forwarding
aa) infoIs
ab) learn
ac) learning
ad) proposed
ae) proposing
af) rcvdInfo
ag) rcvdMsg

ah) rcvdTc
ai) reRoot
aj) reselect
ak) selected
al) tcProp
am) updtInfo

The following variables perform the functions described in 17.~~17~~ 19 of IEEE Std 802.1D, ~~1998~~ 2004 Edition, but have enhanced or extended specifications or considerations. There is one instance per-Port of each variable for the CIST, and one per-Port for each MSTI.

an) agreed (13.24.2)
ao) ~~portId~~ designatedPriority(~~13.24.25~~13.24.4)
~~ap) rcvdInfo (13.24.29) replaces the functionality of the variable previously named rcvdMsg;~~
aq) ~~role~~ designatedTimes (~~13.24.31~~13.24.5)
ar) ~~selectedRole~~ msgPriority (~~13.24.32~~13.24.23)
as) ~~sync~~ msgTimes (~~13.24.33~~13.24.24)
at) ~~synced~~ portId (~~13.24.34~~13.24.25)

~~The following variables perform the related functions described in 17.17 of IEEE Std 802.1D, 1998 Edition, but have extended specifications. There is one instance per-Port of each variable for the CIST.~~

au) ~~cistDesignatedPriority~~portPriority (~~13.24.4~~13.24.26)
av) ~~cistDesignatedTimes~~ portTimes (~~13.24.5~~13.24.27)
aw) ~~cistMsgPriority~~ role (~~13.24.23~~13.24.31)
ax) ~~cistMsgTimes~~ selectedRole (~~13.24.24~~13.24.32)
ay) ~~cistPortPriority~~ sync (~~13.24.26~~13.24.33)
az) ~~cistPortTimes~~ synced (~~13.24.27~~13.24.34)

The following variables perform the related functions described in 17.~~17~~ 19 of IEEE Std 802.1D, ~~1998~~ 2004 Edition, but have ~~enhanced or~~ extended specifications. There is one instance per-Port of each variable for ~~each MSTI~~the CIST.

ba) ~~mstiDesignatedPriority~~newInfo (~~13.24.11~~13.24.21)
~~bb) mstiDesignatedTimes (13.24.12)~~
~~bc) mstiMsgPriority (13.24.17)~~
~~bd) mstiMsgTimes (13.24.18)~~
~~be) mstiPortPriority (13.24.19)~~

The following variable(s) are additional to those specified in 17.19 of IEEE Std 802.1D, 2004 Edition. There is one instance per-Port of each variable for the CIST, and one per-Port for each MSTI.

bf) ~~mstiPortTimes~~ changedMaster (~~13.24.20~~13.24.3)

The following variable(s) are additional to those specified in 17.~~17~~ 19 of IEEE Std 802.1D, ~~1998~~ 2004 Edition. There is one instance per-Port of each variable ~~for the CIST, and one per-Port~~ for each MSTI.

bg) ~~agree~~ master (~~13.24.1~~13.24.14)
bh) ~~changedMaster~~ mastered (~~13.24.3~~13.24.16)
~~bi) rcvdMsg (13.24.30) provides functionality distinct from the previous rcvdMsg variable, now renamed rcvdInfo (13.24.29)~~

~~The following variable(s) specified in 17.17 of IEEE Std 802.1D, 1998 Edition are not required or have been replaced by other variables in this specification.~~

~~bj) newInfo (replaced by newInfoCist and newInfoMsti)~~

~~bk) tc (not required)~~

~~13.24.1 agree~~

~~This variable is used by the Port Transmit state machine to set the value of the Agreement flag in transmitted Configuration Messages for the given tree.~~

13.24.2 agreed

A Boolean value indicating that a Configuration Message has been received from another Bridge attached to the same LAN indicating Agreement that all the Port States for the given tree of all other Bridges attached to the same LAN as this Port are known to be likewise compatible with a loop free active topology determined by this Bridge's priority vectors and, in the absence of further communication with this Bridge, will remain compatible within the design probabilities of protocol failure due to repeated BPDU loss (13.16,13.19).

13.24.3 changedMaster

Set, for all Ports for all MSTIs, by the CIST Port Role Selection state machine using the ~~updtRoleCist~~updtRolesTree() procedure if the root priority vector selected has a different Regional Root Identifier than that previously selected, and has or had a non-zero CIST External Path Cost. changedMaster is always FALSE for the CIST.

NOTE—Changes in Regional Root Identifier will not cause loops if the Regional Root is within ~~an~~ a MST Region, as is the case if and only if the MST Region is the Root of the CST. This important optimization allows the MSTIs to be fully independent of each other in the case where they compose the core of a network.

13.24.4 ~~cistDesignatedPriority~~designatedPriority

~~The~~ For the CIST and a given Port, the CIST Root Identifier, External Root Path Cost, Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Port's CIST designated priority vector, as defined in 13.10.

For a given MSTI and Port, the Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Port's designated priority vector, as defined in 13.11.

13.24.5 ~~cistDesignatedTimes~~designatedTimes

~~The~~ For the CIST and a given Port, the set of timer parameter values (Message Age, Max Age, Forward Delay, and remainingHops) that are used to update Port Times when updtInfo is set. The value of designatedTimes is copied from the ~~cistRootTimes~~ CIST rootTimes Parameter (13.23.7) by the operation of the ~~updtRolesCist~~updtRolesTree() procedure.

~~13.24.6 cistMsgPriority~~

~~The CIST Root Identifier, External Root Path Cost, Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the CIST message priority vector conveyed in a received BPDU, as defined in 13.10.~~

~~**13.24.7 cistMsgTimes**~~

~~The timer parameter values (Message Age, Max Age, Forward Delay, Hello Time, and remainingHops) conveyed in a received BPDU. If the BPDU is an STP or RSTP without MSTP parameters, remainingHops is set to zero.~~

~~**13.24.8 cistPortPriority**~~

~~The CIST Root Identifier, External Root Path Cost, Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Port's port priority vector, as defined in 13.10.~~

~~**13.24.9 cistPortTimes**~~

~~The Port's timer parameter values (Message Age, Max Age, Forward Delay, Hello Time, and remainingHops). These timer values are used in BPDUs transmitted from the Port.~~

For a given MSTI and Port, the value of remainingHops used to update this MSTI's portTimes parameter when updtInfo is set. The value of designatedTimes is copied from this MSTI's rootTimes parameter (13.23.13) by the operation of the updtRolesTree() procedure.

13.24.10 infoInternal

If infoIs is Received, indicating that the port has received current information from the Designated Bridge for the attached LAN, infoInternal is set if that Designated Bridge is in the same MST Region as the receiving Bridge, and ~~clear~~ reset otherwise.

~~**13.24.11 mstiDesignatedPriority**~~

~~The Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Port's designated priority vector, as defined in 13.11.~~

~~**13.24.12 mstiDesignatedTimes**~~

~~The value of remainingHops used to update mstiPortTimes when updtInfo is set. The value of mstiDesignatedTimes is copied from the mstiRootTimes Parameter (13.23.13) by the operation of the updtMstiRolesBridge() procedure.~~

~~**13.24.13 mstiMaster**~~

13.24.14 master

A Boolean variable used to determine the value of the Master flag for this MSTI and Port in transmitted MST BPDUs.

Set TRUE if the Port Role for the MSTI and Port is Root Port or Designated Port, and the Bridge has selected one of its Ports as the Master Port for this MSTI or the ~~mstiMastered~~ mastered flag is set for this MSTI for any other Bridge Port with a Root Port or Designated Port Role. Set FALSE otherwise.

~~13.24.15 mstiMastered~~

13.24.16 mastered

A Boolean variable used to record the value of the Master flag for this MSTI and Port in MST BPDUs received from the attached LAN.

NOTE—~~mstiMaster~~ master and ~~mstiMastered~~ mastered signal the connection of the MSTI to the CST via the Master Port throughout the MSTI. These variables and their supporting procedures do not affect the connectivity provided by this revision of this standard, but permit future enhancements to MSTP providing increased flexibility in the choice of Master Port without abandoning plug and play network migration. They are therefore omitted from the overviews of protocol operation, including Figure 13-36.

~~13.24.17 mstiMsgPriority~~

~~The Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the MSTI message priority vector, as defined in 13.11 and conveyed in a received BPDU for this MSTI.~~

~~13.24.18 mstiMsgTimes~~

~~The value of remainingHops received with message priority components of the mstiMsgPriority for this MSTI.~~

~~13.24.19 mstiPortPriority~~

~~The Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Port's MSTI port priority vector, as defined in 13.11.~~

~~13.24.20 mstiPortTimes~~

~~The value of remainingHops for this MSTI in BPDUs transmitted through the Port.~~

13.24.21 ~~newInfoCist~~newInfo

A Boolean variable set TRUE if a BPDU conveying changed CIST information is to be transmitted. It is set FALSE by the Port Transmit state machine.

13.24.22 newInfoMsti

A Boolean variable set TRUE if a BPDU conveying changed MSTI information is to be transmitted. It is ~~value is~~ set FALSE by the Port Transmit state machine.

13.24.23 msgPriority

For the CIST and a given Port, the CIST Root Identifier, External Root Path Cost, Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the CIST message priority vector conveyed in a received BPDU, as defined in 13.10.

For a given MSTI and Port, the Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the MSTI message priority vector, as defined in 13.11 and conveyed in a received BPDU for this MSTI.

13.24.24 msgTimes

For the CIST and a given Port, the timer parameter values (Message Age, Max Age, Forward Delay, Hello Time, and remainingHops) conveyed in a received BPDU. If the BPDU is a ST or RST BPDU without MSTP parameters, remainingHops is set to zero.

<< Editor's note: isn't this last sentence in contradiction with updtRcvdInfoWhile[Cist]()'s definition where "the value of the remainingHops component of portTimes is set to MaxHops" when information is received from a Bridge external to the MST Region? >>

For a given MSTI and Port, the value of remainingHops received with message priority components of this MSTI's msgPriority parameter.

<< Editor's note: is that correct, i.e. how do we infer a time value from message priority components? >>

13.24.25 portId

The Port Identifier for this Port. This variable forms a component of the port priority and designated priority vectors (13.10,13.11).

The four most significant bits of the Port Identifier (the settable Priority component) for the CIST and for each MSTI can be modified independently of the setting of those bits for all other trees, as a part of allowing full and independent configuration control to be exerted over each Spanning Tree instance.

13.24.26 portPriority

For the CIST and a given Port, the CIST Root Identifier, External Root Path Cost, Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Port's port priority vector, as defined in 13.10.

For a given MSTI and Port, the Regional Root Identifier, Internal Root Path Cost, Designated Bridge Identifier, and Designated Port Identifier components of the Port's MSTI port priority vector, as defined in 13.11.

13.24.27 portTimes

For the CIST and a given Port, the Port's timer parameter values (Message Age, Max Age, Forward Delay, Hello Time, and remainingHops). These timer values are used in BPDUs transmitted from the Port.

For a given MSTI and Port, the value of remainingHops for this MSTI in BPDUs transmitted through the Port.

13.24.28 rcvdInternal

A Boolean variable set TRUE by the Receive Machine if the BPDU received was transmitted by a Bridge in the same MST Region as the receiving Bridge.

~~13.24.29 rcvdInfo~~

~~Set to the result of the rcvInfoCist procedure for the CIST and the rcvInfoMsti procedure for an MSTI. It can take the values SuperiorDesignatedInfo, RepeatedDesignatedInfo, RootInfo, or OtherInfo.~~

~~13.24.30 rcvdMsg~~

~~A Boolean variable set TRUE by the Receive Machine if the BPDU received contains a message for this tree.~~

13.24.31 role

The assigned Port Role. The ~~port~~ port's role is either ~~a~~ DisabledPort, ~~a~~ RootPort, ~~a~~ DesignatedPort, ~~an~~ AlternatePort, ~~a BackupPort,~~ BackupPort or ~~a~~ MasterPort.

NOTE—The role of MasterPort is introduced for MSTIs for a Port where the CIST Port Role is ~~a~~ RootPort and the spanning tree information received is from another MST Region. ~~An~~ A MSTI Master Port forms part of the stable active topology for frames allocated to that MSTI, just as the CIST Root Port forwards frames allocated to the CIST. The Port State for each MSTI may differ for each MSTI as required to suppress temporary loops.

13.24.32 selectedRole

A newly computed role for the Port.

13.24.33 sync

A Boolean value. Set TRUE to force the Port State to be compatible with the loop free active topology determined by the priority vectors held by this Bridge (13.16,13.19) for this tree (CIST, or MSTI), by transitioning the Port State to Discarding and soliciting an Agreement if possible, if the Port is not already synchronized (13.24.34).

13.24.34 synced

A Boolean value. TRUE only if the Port State is compatible with the loop free active topology determined by the priority vectors held by this Bridge for this tree (13.16,13.19).

~~13.25 State Machine Conditions~~

13.26 State Machine Conditions and Parameters

The following boolean variable evaluations are defined for notational convenience in the state machines. These definitions also serve to highlight those cases where a state transition for one tree (CIST or MSTI) depends on the state of the variables of one or more other trees.

The following conditions and parameters are as specified in 17.20 of IEEE Std 802.1D, 2004 Edition:

a) AdminEdge
b) AutoEdge
c) EdgeDelay
d) forwardDelay
e) MigrateTime
f) reRooted
g) rstpVersion
h) stpVersion
i) TxHoldCount

The following conditions and parameters are similar to those specified in 17.20 of IEEE Std 802.1D, 2004 Edition but have enhanced or extended specifications or considerations:

j) allSynced (13.26.1)
k) FwdDelay (13.26.5)
l) HelloTime (13.26.6)
m) MaxAge (13.26.7)

The following conditions and parameters are additional to those described in 17.20 of IEEE Std 802.1D, 2004 Edition:

n) cist (13.26.2)
o) cistRootPort (13.26.3)
p) cistDesignatedPort (13.26.4)
q) mstiRootPort (13.26.9)
r) mstiDesignatedPort (13.26.10)
s) mstiMasterPort (13.26.11)
t) rcvdAnyMsg (13.26.12)
u) rcvdCistMsg (13.26.13)
v) rcvdMstiMsg (13.26.14)
w) restrictedRole (13.26.16)
x) restrictedTcn (13.26.17)
y) updtCistInfo (13.26.18)
z) updtMstiInfo (13.26.19)

13.26.1 allSynced

~~TRUE if and only if synced is TRUE for all Ports for the given Tree (CIST or MSTI).~~

TRUE if and only if, for all Ports for the given Tree, selected is TRUE and the port's role is the same as its selectedRole, updtInfo is FALSE and either

a) synced is true; or
b) The port is the Root Port; or
c) The port is the Master Port.

13.26.2 cist

TRUE only for CIST state machines, i.e. FALSE for MSTI state machine instances.

13.26.3 cistRootPort

TRUE if the CIST role for the given Port is RootPort.

13.26.4 cistDesignatedPort

TRUE if the CIST role for the given Port is DesignatedPort.

13.26.5 FwdDelay

The Forward Delay component of the CIST's designatedTimes parameter (13.24.5).

13.26.6 HelloTime

The Hello Time component of the CIST's designatedTimes parameter (13.24.5).

13.26.7 MaxAge

The Max Age component of the CIST's designatedTimes parameter (13.24.5).

13.26.8 MigrateTime

The Max Age component of the CIST's designatedTimes parameter (13.24.5).

13.26.9 mstiRootPort

TRUE if the role for any MSTI for the given Port is RootPort.

13.26.10 mstiDesignatedPort

TRUE if the role for any MSTI for the given Port is DesignatedPort.

13.26.11 mstiMasterPort

TRUE if the role for any MSTI for the given Port is MasterPort.

13.26.12 rcvdAnyMsg

~~rcvdAnyMsg is~~ TRUE for a given Port if rcvdMsg is TRUE for the CIST or any MSTI for that Port.

13.26.13 ~~rcvdCistInfo~~rcvdCistMsg

~~rcvdCistInfo is~~ TRUE for a given Port if and only if rcvdMsg is TRUE for the CIST for that Port.

13.26.14 ~~rcvdMstiInfo~~rcvdMstiMsg

~~rcvdMstiInfo is~~ TRUE for a given Port and MSTI if and only if rcvdMsg is FALSE for the CIST for that Port and rcvdMsg is TRUE for the MSTI for that Port.

~~13.26.15 reRooted~~

~~TRUE if the rrWhile timer is clear (zero) for all Ports for the given Tree other than the given Port.~~

13.26.16 restrictedRole

A Boolean value set by management. If TRUE causes the Port not to be selected as Root Port for the CIST or any MSTI, even it has the best spanning tree priority vector. Such a Port will be selected as an Alternate Port after the Root Port has been selected. This parameter should be FALSE by default. If set it can cause lack of spanning tree connectivity. It is set by a network administrator to prevent bridges external to a core region of the network influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator.

13.26.17 restrictedTcn

A Boolean value set by management. If TRUE causes the Port not to propgate received topology change notifications and topology changes to other Ports. This parameter should be FALSE by default. If set it can cause temporary loss of connectivity after changes in a spanning trees active topology as a result of persistent incorrectly learnt station location information. It is set by a network administrator to prevent bridges external to a core region of the network causing address flushing in that region, possibly because those bridges are not under the full control of the administrator or MAC_Operational for the attached LANs transitions frequently.

13.26.18 updtCistInfo

~~updtCistInfo is~~ TRUE for a given Port if and only if updtInfo is TRUE for the CIST for that Port.

13.26.19 updtMstiInfo

~~updtMstiInfo is~~ TRUE for a given Port and MSTI if and only if updtInfo is TRUE for the MSTI for that Port or either updtInfo or selected are TRUE for the CIST for that Port.

NOTE—The dependency of ~~revdMstiInfo~~ rcvdMstiMsg and updtMstiInfo on CIST variables for the Port reflects the fact that MSTIs exist in a context of CST parameters. The state machines ensure that the CIST parameters from received BPDUs are processed and updated prior to processing MSTI information.

13.27 State Machine Procedures

The following procedures perform the functions specified in 17.~~17~~ 21 of IEEE Std 802.1D, ~~1998~~ 2004 Edition for the CIST state machines.

a) txTcn()

The following procedures perform the functions specified in 17.~~17~~ 21 of IEEE Std 802.1D, ~~1998~~ 2004 Edition for the CIST or any given MSTI instance.

~~b) disableForwarding()~~
c) ~~disableLearning~~disableForwarding()
d) ~~enableForwarding~~disableLearning()
e) ~~enableLearning~~enableForwarding()
f) ~~flush~~enableLearning()
g) ~~updtBPDUVersion~~recordPriority()

The following procedures perform the general functions described in 17.~~17~~ 21 of IEEE Std 802.1D, ~~1998~~ 2004 Edition for both the CIST and the MSTI state machines or specifically for the CIST or a given MSTI, but have enhanced or extended specifications or considerations.

h) betterorsameInfo(newInfoIs) (13.27.1)
i) clearReselectTree() (13.27.4)
j) ~~clearReselectTree~~newTcWhile() (~~13.27.4~~13.27.6)
k) ~~newTcWhile~~rcvInfo() (~~13.27.6~~13.27.7)
l) ~~rcvInfoCist~~recordAgreement() (~~13.27.7) and rcvInfoMsti() (13.27.8~~13.27.9)
m) ~~recordProposalCist~~recordDispute() (~~13.27.14) and recordProposalMsti() (13.27.15~~13.27.11)
n) ~~setReRootTree~~recordProposal(~~13.27.17~~) (13.27.14)
o) ~~setSelectedTree~~recordTimes() (~~13.27.18~~13.27.15)
p) ~~setSyncTree~~setReRootTree(13.27.17) ~~(13.27.19)~~
q) ~~setTcFlags~~setSelectedTree() (~~13.27.20~~13.27.18)
r) ~~setTcPropTree~~setSyncTree() (~~13.27.21)~~ 13.27.19)
s) ~~txConfig~~setTcFlags() (~~13.27.22~~13.27.20)
t) ~~txMstp~~setTcPropTree() (~~13.27.23)~~13.27.21)

~~NOTE 1—the equivalent procedure to txMstp() in IEEE Std 802.1D, 1998 Edition is called txRstp().~~

u) ~~updtRcvdInfoWhileCist~~txConfig() (~~13.27.25) and updtRcvdInfoWhileMsti() (13.27.26~~13.27.22)

v) ~~updtRolesCist~~txMstp() (~~13.27.28) and updtRolesMsti() (13.27.28~~13.27.23)

NOTE 1—the equivalent procedure to txMstp() in IEEE Std 802.1D, 2004 Edition is called txRstp().

w) ~~updtRolesDisabledTree~~updtBPDUVersion() (~~13.27.29~~13.27.24)

~~NOTE 2—clearReselectBridge() has been replaced by clearReselectTree(); rcvBpdu() by rcvInfoCist() and rcvInfoMsti(); recordProposed() by recordProposalCist() and recordProposalMsti(); setSelectedBridge() by setSelectedTree(); setReRootBridge() by setReRootTree(); setSyncBridge() by setSyncTree(); setTcPropBridge() by setTcPropTree(); updtRcvdInfoWhile() by updtRcvdInfoWhileCist() and updtRcvdInfoWhileMsti(); updtRolesBridge() by updtRolesCist() and updtRolesMsti(); and updtRoleDisabledBridge() by updtRolesDisabledTree().~~

~~The following procedures perform functions additional to those described in 17.17 of IEEE Std 802.1D, 1998 Edition for both the CIST and the MSTI state machines or for the CIST or a given MSTI specifically:~~

x) ~~betterorsameInfoCist~~updtRcvdInfoWhile() (~~13.27.1) and betterorsameInfoMsti() (13.27.2~~13.27.25)
y) ~~clearAllRcvdMsgs~~updtRolesTree() (~~13.27.3~~13.27.28)
z) ~~fromSameRegion~~updtRolesDisabledTree() (~~13.27.5~~13.27.29)

The following procedures perform functions additional to those described in 17.21 of IEEE Std 802.1D, 2004 Edition for both the CIST and the MSTI state machines or for the CIST or a given MSTI specifically:

aa) ~~recordAgreementCist~~clearAllRcvdMsgs() (~~13.27.9) and recordAgreementMsti() (13.27.10~~13.27.3)
ab) ~~recordMasteredCist~~fromSameRegion() (~~13.27.12~~13.27.5)
ac) ~~recordMasteredMSTI~~recordMastered() (~~13.27.13~~13.27.12)
ad) setRcvdMsgs() (13.27.16)

All references to named variables in the specification of procedures are to instances of the variables corresponding to the instance of the state machine using the function, i.e. to the CIST or the given MSTI as appropriate. References to the forwarding and learning functions for a Port apply to all and only those Filtering Databases associated with that specific tree.

13.27.1 ~~betterorsameInfoCist~~betterorsameInfo(newInfoIs)

~~Returns TRUE if the received CIST priority vector is better than or the same as (13.10) the CIST port priority vector.~~

~~13.27.2 betterorsameInfoMsti()~~

~~Returns TRUE if the MSTI priority vector is better than or the same as (13.11) the MSTI port priority vector.~~

Returns TRUE if, for a given Port and Tree (CIST, or MSTI), either

a) The procedure's parameter newInfoIs is Received, and infoIs is Received and the msgPriority vector is better than or the same as (13.10) the portPriority vector; or,
b) The procedure's parameter newInfoIs is Mine, and infoIs is Mine and the designatedPriority vector is better than or the same as (13.10) the portPriority vector.

Returns False otherwise.

NOTE—This procedure is not invoked (in the case of a MSTI) if the received BPDU carrying the MSTI information was received from another MST Region. In that event, the Port Receive Machine (using setRcvdMsgs()) does not set rcvdMsg for any MSTI, and the Port Information Machine's SUPERIOR_DESIGNATED state is not entered.

13.27.3 clearAllRcvdMsgs()

Clears rcvdMsg for the CIST and all MSTIs, for all Ports.

13.27.4 clearReselectTree()

Clears reselect for the tree (the CIST or a given MSTI) for all Ports of the Bridge.

13.27.5 fromSameRegion()

Returns TRUE if rcvdRSTP is TRUE, and the received BPDU conveys ~~an~~ a MST Configuration Identifier that matches that held for the Bridge. Returns FALSE otherwise.

13.27.6 newTcWhile()

~~This procedure sets the value of tcWhile, if and only if it is currently zero, to twice HelloTime on point-to-point links (i.e., links where the operPointToPointMAC parameter is TRUE; see 6.4.3) where the partner bridge port is RSTP capable, and to the sum of the Max Age and Forward Delay components of rootTimes otherwise (non-RSTP capable partners or shared media). The value of HelloTime is taken from cistPortTimes (13.24.27) for this Port.~~

If the value of tcWhile is zero and sendRSTP is TRUE, this procedure sets the value of tcWhile to HelloTime plus one second and sets either newInfoCist TRUE for the CIST, or newInfoMsti TRUE for a given MSTI. The value of HelloTime is taken from the CIST's portTimes parameter (13.24.27) for this Port.

If the value of tcWhile is zero and sendRSTP is FALSE, this procedure sets the value of tcWhile to the sum of the Max Age and Forward Delay components of rootTimes and does not change the value of either newInfoCist or newInfoMsti.

Otherwise the procedure takes no action.

13.27.7 ~~rcvInfoCist~~rcvInfo()

Decodes, for a given Port and Tree (CIST, or MSTI), the message priority and timer values from the received BPDU storing them in the msgPriority and msgTimes variables.

Returns SuperiorDesignatedInfo if, for a given Port and Tree (CIST, or MSTI):

a) The received CIST or MSTI message conveys a Designated Port Role, and
   1) The message priority (msgPriority—13.24.23) is superior (13.10 or 13.11) to the Port's port priority vector, or
   2) The message priority is the same as the Port's port priority vector, and any of the received timer parameter values (msgTimes—13.24.24) differ from those already held for the Port (portTimes—13.24.27).

Returns RepeatedDesignatedInfo if, for a given Port and Tree (CIST, or MSTI):

b) ~~Returns SuperiorDesignatedInfo if the~~ The received CIST or MSTI message conveys a Designated Port Role, and a message priority ~~(cistMsgPriority—13.24.23)~~ vector and timer parameters that ~~is superior (13.10) to~~ are the same as the Port's port priority ~~vector,~~ vector or ~~any of the received~~ timer ~~parameter values (cistMsgTimes—13.24.24) differ from those already held for the Port (cistPortTimes—13.24.27)~~values.

Returns InferiorDesignatedInfo if, for a given Port and Tree (CIST, or MSTI):

c) ~~Returns RepeatedDesignatedInfo if the~~ The received CIST or MSTI message conveys a Designated Port Role, and a message priority vector ~~and timer parameters~~ that ~~are the same as~~ is worse than the Port’s port priority ~~vector or timer values~~vector.

Returns InferiorRootAlternateInfo if, for a given Port and Tree (CIST, or MSTI):

d) ~~Returns RootInfo if the~~ The received CIST or MSTI message conveys a Root Port, Alternate Port, or Backup Port Role and a CIST or MSTI message priority that is the same as or worse than the CIST or MSTI port priority vector.

Otherwise, returns OtherInfo.

~~13.27.8 revInfoMsti()~~

NOTE—A Configuration BPDU implicitly conveys a Designated Port Role.

<< Editor’s note: should we modify this definition as suggested in http://www.ieee802.org/1/private/email/msg02116.html (the reference to Message Age in a)2) below would only apply wrt. the CIST’s msgTimes parameter):

“Decodes, for a given Port and Tree (CIST, or MSTI), the message priority and timer values from the received BPDU storing them in the msgPriority and msgTimes variables.

Returns SuperiorDesignatedInfo if, for a given Port and Tree (CIST, or MSTI):

a) The received CIST or MSTI message conveys a Designated Port Role, and

1) The message priority (msgPriority—13.24.23) is superior (13.10 or 13.11) to the Port’s port priority vector, or

~~Returns SuperiorDesignatedInfo if the received MSTI message conveys a~~ 2) The message priority ~~(mstiMsgPriority—13.24.17) that~~ is ~~superior (13.11) to~~ the same as the Port’s port priority vector, ~~or~~ and any of the received timer parameter values (~~mstiMsgTimes~~msgTimes—~~13.24.18~~13.24.24) **other than Message Age** differ from those already held for the Port (~~mstiPortTimes~~portTimes—~~13.24.20~~13.24.27).

Otherwise, r~~R~~eturns RepeatedDesignatedInfo if, for a given Port and Tree (CIST, or MSTI):

b) The received CIST or MSTI message conveys a Designated Port Role, and**:**

~~a~~**1) A** message priority vector and timer parameters that are the same as the Port’s port priority vector or timer values~~,~~; and

2) infoIs is Mine.

Otherwise, r~~R~~eturns InferiorDesignatedInfo if, for a given Port and Tree (CIST, or MSTI):

~~Returns RepeatedDesignatedInfo if the~~ c) The received CIST or MSTI message conveys a Designated Port Role~~, and~~ ~~**a message priority vector**~~ ~~and timer parameters~~ ~~**that**~~ ~~are~~ ~~**is worse than**~~ ~~**the**~~ ~~same as the~~ ~~**Port’s port priority**~~ ~~vector or timer values~~~~**vector**~~.

Otherwise, r~~R~~eturns InferiorRootAlternateInfo if, for a given Port and Tree (CIST, or MSTI):

~~Returns RootInfo if the~~ d) The received CIST or MSTI message conveys a Root Port, Alternate Port, or Backup Port Role and a CIST or MSTI message priority that is the same as or worse than the CIST or MSTI port priority vector.

Otherwise, returns OtherInfo.

NOTE—A Configuration BPDU implicitly conveys a Designated Port Role.”

>>

13.27.9 ~~recordAgreementCist~~recordAgreement()

~~If~~ For the CIST ~~Message was~~ and a ~~Configuration Message received on a point to point link~~given Port, if rstpVersion is TRUE, operPointToPointMAC (6.4.3) is TRUE, and the received CIST ~~message~~ Message has the Agreement flag set, and conveys either:

1) a Root Port Role with message priority the same as or worse than the port priority vector, or
2) a Designated Port Role with message priority the same as or better than the port priority vector,

~~the CIST agreed flag is set. Otherwise the CIST agreed flag is cleared.~~

the CIST agreed flag is set and the CIST proposing flag is cleared. Otherwise the CIST agreed flag is cleared. Additionally, if the CIST message was received from a Bridge in a different MST Region i.e. the rcvdInternal flag is clear, the agreed flags for this Port for all MSTIs are set or cleared to the same value as the CIST agreed flag. If the CIST message was received from a Bridge in the same MST Region, the MSTI agreed flags are not changed.

~~13.27.10 recordAgreementMsti()~~

~~If the MSTI Message was received on a point to point link and:~~

For a given MSTI and Port, if operPointToPointMAC (6.4.3) is TRUE, and:

a) the message priority vector of the CIST Message accompanying ~~this~~ the received MSTI Message (i.e. received in the same BPDU) has the same CIST Root Identifier, CIST External Root Path Cost, and Regional Root Identifier as the CIST port priority vector, and
b) the received MSTI Message has the Agreement flag set, and conveys either
   1) a Root Port Role with message priority the same as or worse than the MSTI port priority vector, or
   2) a Designated Port Role with message priority the same as or better than the port priority vector,

the MSTI agreed flag is ~~set~~set and the MSTI proposing flag is cleared. Otherwise the MSTI agreed flag is cleared.

NOTE—MSTI Messages received from Bridges external to the MST Region are discarded and not processed by ~~recordAgreeementMsti~~recordAgreeement() or ~~recordProposalMsti~~recordProposal().

13.27.11 ~~recordMasteredCist~~recordDispute()

If a RST BPDU or a MST BPDU with the learning flag set has been received:

a) the disputed variable is set; and
b) the agreed variable is cleared.

13.27.12 recordMastered()

~~If~~ For the CIST and a given Port, if the CIST message was received from a Bridge in a different MST Region, i.e. the rcvdInternal flag is clear, the ~~mstiMastered~~ mastered variable for this Port is cleared for all MSTIs.

~~13.27.13 recordMasteredMsti()~~

~~If~~ For a given MSTI and Port, if the MSTI ~~Message~~ message was received on a point to point link and the MSTI Message has the Master flag set, set the ~~mstiMastered~~ mastered variable for ~~the~~ this MSTI. Otherwise reset the ~~mstiMastered~~ mastered variable.

13.27.14 ~~recordProposalCist~~recordProposal()

For the CIST and a given Port, if the received CIST Message conveys a Designated Port Role, and has the Proposal flag set, the CIST proposed flag is set. Otherwise the CIST proposed flag is not changed. Additionally, if the CIST Message was received from a Bridge in a different MST Region, i.e. the rcvdInternal flag is clear, the proposed flags for this Port for all MSTIs are set or cleared to the same value as the CIST proposed flag. If the CIST message was received from a Bridge in the same MST Region, the MSTI proposed flags are not changed.

~~If~~ For a given MSTI and Port, if the ~~CIST~~ received MSTI Message ~~was~~ conveys a ~~Configuration Message received on a point to point link~~ Designated Port Role, and has the Proposal flag set, the ~~CIST~~ MSTI proposed flag is set. Otherwise the ~~CIST~~ MSTI proposed flag is ~~cleared~~not changed.

~~Additionally, if the CIST message was received from a Bridge in a different MST Region i.e. the rcvdInternal flag is clear, the proposed flags for this Port for all MSTIs are set or cleared to the same value as the CIST proposed flag. If the CIST message was received from a Bridge in the same MST Region, the MSTI proposed flags are not changed.~~

13.27.15 ~~recordProposalMsti~~recordTimes()

~~If the MSTI Message was received on a point to point link and has the Proposal flag set, the MSTI proposed flag is set. Otherwise the MSTI proposed flag is cleared.~~

For the CIST and a given Port, sets portTimes' Message Age, Max Age, Forward Delay and remainingHops to the received values held in msgTimes and portTimes' Hello Time to msgTimes' Hello Time if that is greater than the minimum specified in the Compatibility Range column of Table 17-1 of IEEE Std 802.1D, 2004 Edition, and to that minimum otherwise.

For a given MSTI and Port, sets portTime's remainingHops to the received value held in msgTimes.

13.27.16 setRcvdMsgs()

Sets rcvdMsg for the CIST, and makes the received CST or CIST message available to the CIST Port Information state machines.

Additionally and if and only if rcvdInternal is set, sets rcvdMsg for each and every MSTI for which ~~an~~ a MSTI message is conveyed in the BPDU, and makes available each MSTI message and the common parts of the CIST message priority (the CIST Root Identifier, External Root Path Cost, and Regional Root Identifier) to the Port Information state machine for that MSTI.

13.27.17 setReRootTree()

~~This procedure sets~~ Sets reRoot TRUE for this tree (the CIST or a given MSTI) for all Ports of the Bridge.

13.27.18 setSelectedTree()

Sets selected TRUE for this tree (the CIST or a given MSTI) for all Ports of the Bridge.

13.27.19 setSyncTree()

Sets sync TRUE for this tree (the CIST or a given MSTI) for all Ports of the Bridge.

13.27.20 setTcFlags()

~~If the received BPDU is a TCN BPDU, sets rcvdTcn TRUE and sets rcvdTc TRUE for each and every MSTI.~~

~~Otherwise, if the received BPDU is a ConfigBPDU or RST BPDU:~~

For the CIST and a given Port, if the received BPDU is a TCN BPDU, sets rcvdTcn TRUE. Otherwise, if the received BPDU is a ConfigBPDU, a RST BPDU or a MST BPDU:

a) If the Topology Change Acknowledgment flag is set in the CST message, sets rcvdTcAck TRUE.
b) If rcvdInternal is clear and the Topology Change flag is set in the CST message, sets rcvdTc TRUE.
c) If rcvdInternal is set, sets rcvdTc for the CIST if the Topology Change flag is set in the CST message.

For a given MSTI and Port, if the received BPDU is a TCN BPDU, sets rcvdTcn TRUE and sets rcvdTc TRUE. Otherwise, if the received BPDU is a ConfigBPDU, a RST BPDU or a MST BPDU:

a) If the Topology Change Acknowledgment flag is set in the CST message, sets rcvdTcAck TRUE.
b) If rcvdInternal is clear and the Topology Change flag is set in the CST message, sets rcvdTc ~~for the CIST and for each and every MSTI~~TRUE.
c) If rcvdInternal is set, sets rcvdTc for ~~the CIST~~ this MSTI if ~~the Topology Change flag is set in the CST message, and sets rcvdTc for each MSTI for which~~ the Topology Change flag is set in the corresponding MSTI message.

<< Editor's note: what does "CST message" refer to above? >>

13.27.21 setTcPropTree()

~~Sets tcProp TRUE for the given tree (the CIST or an MSTI) for all Ports except the Port that invoked the procedure.~~

If and only if restrictedTcn is FALSE for the Port that invoked the procedure, sets tcProp TRUE for the given tree (the CIST or a given MSTI) for all other Ports.

13.27.22 txConfig()

Transmits a Configuration BPDU. The first four components of the message priority vector (13.24.23) conveyed in the BPDU are set to the value of the CIST Root Identifier, External Root Path Cost, Bridge Identifier, and Port Identifier components of the ~~cistPortPriority~~ CIST's designatedPriority parameter (~~13.24.26~~13.24.4) for this Port. The topology change flag is set if (tcWhile != 0) for the Port. The topology change acknowledgement flag is set to the value of TcAck for the Port. The remaining flags are set to zero. The value of the Message Age, Max Age, Fwd Delay, and Hello Time parameters conveyed in the BPDU are set to the values held in ~~cistPortTimes~~ the CIST's designatedTimes parameter (~~13.24.27~~13.24.5) for the Port.

13.27.23 txMstp()

Transmits ~~an~~ a MST BPDU (14.3.3), encoded according to the specification contained in 14.6. The first six components of the CIST message priority vector (13.24.23) conveyed in the BPDU are set to the value of ~~cistPortPriority~~ the CIST's designatedPriority parameter (~~13.24.26~~13.24.4) for this Port. The Port Role in the BPDU (14.2.1) is set to the current value of the role variable for the transmitting port (13.24.31). The Agreement and Proposal flags in the BPDU are set to the values of the ~~agreed~~ agree (~~13.24.2~~13.24.1) and proposing (13.24 of this standard, 17.~~18~~19.~~20~~ 24 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) variables for the transmitting ~~Port~~ Port, respectively. The CIST topology change flag is set if (tcWhile != 0) for the Port. The topology change acknowledge flag in the BPDU is never used and is set to zero. The learning ~~flag is~~ and forwarding flags in the BPDU are set ~~if~~ to the values of the learning ~~variable for the CIST~~ (13.24 of this standard, 17.~~18~~19.~~9~~ 12 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) ~~is TRUE. The forwarding flag is set if the~~ and forwarding ~~variable for the CIST~~ (13.24 of this standard, 17.~~18~~19.~~5~~ 9 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) ~~is TRUE~~variables for the CIST, respectively. The value of the Message Age, Max Age, Fwd Delay, and Hello Time parameters conveyed in the BPDU are set to the values held in ~~estiPortTimes~~ the CIST's designatedTimes parameter (~~13.24.27~~13.24.5) for the Port.

If the value of the Force Protocol Version parameter is less than 3, no further parameters are encoded in the BPDU and the protocol version parameter is set to 2 (denoting ~~an~~ a RST BPDU). Otherwise, the protocol version parameter is set to 3 and the remaining parameters of the MST BPDU are encoded:

a) The version 3 length.
b) The MST Configuration Identifier parameter of the BPDU is set to the value of the MstConfigId variable for the Bridge (13.23.8).
c) The CIST Internal Root Path Cost (13.24.26).
d) The CIST Bridge Identifier.
e) The CIST Remaining Hops.
f) The parameters of each MSTI message, encoded in MSTID order.

NOTE—No more than 64 MSTIs may be supported. The parameter sets for all of these can be encoded in a standard sized Ethernet frame.

13.27.24 updtBPDUVersion()

Sets rcvdSTP TRUE if the BPDU received is a version 0 or version 1 TCN or a Config BPDU. Sets rcvdRSTP TRUE if the received BPDU is a RST BPDU or a MST BPDU.

13.27.25 ~~updtRevdInfoWhileCist~~updtRcvdInfoWhile()

Calculates the effective age, or remainingHops, to limit the propagation and longevity of received Spanning Tree information for the ~~CIST~~CIST or a given MSTI, setting rcvdInfoWhile to the number of seconds that the information received on a Port will be held before it is either refreshed by receipt of a further configuration message or aged out.

~~If the information was received from a Bridge external to the MST Region (revdInternal FALSE), the effective age of the port information (cistPortPriority and cistPortTimes) is taken as the value of the Message Age parameter carried in a received BPDU (cistMsgTimes), incremented by the greater of (1/16th Max Age) and 1 second, and rounded to the nearest whole second. The value of Message Age and Max Age used in this calculation are taken from the cistPortTimes variable (17.18.18) and, the value assigned to revdInfoWhile is the lower of:~~

~~a) Max Age minus this effective age, and~~
~~b) 3 times the Hello Time,~~

~~or~~

~~c) zero, if the effective age exceeds Max Age.~~

For the CIST and a given Port:

a) If the information was received from a Bridge external to the MST Region (rcvdInternal FALSE):
   1) The value assigned to rcvdInfoWhile is three times the Hello Time, if Message Age, incremented by 1 second and rounded to the nearest whole second, does not exceed Max Age, and is zero otherwise.
   2) ~~and the value of the~~ The remainingHops component of portTimes is set to MaxHops.
b) If the information was received from a Bridge in the same MST Region (rcvdInternal TRUE)~~, the Message Age component of cistPortTimes is set equal to the value received (cistMsgTimes) and the remainingHops component is set equal to the value received decremented by one. The value assigned to rcvdInfoWhile is~~:
~~c) 3 times the Hello Time,~~

~~or~~

~~d) zero, if either cistPortTimes Message Age effective exceeds Max Age or remainingHops is less than or equal to zero.~~

~~The value of Hello Time used by this procedure is taken from cistPortTimes.~~

~~**13.27.26 updtRcvdInfoWhileMsti()**~~

~~Calculates the remainingHops, to limit the propagation and longevity of received Spanning Tree information for an MSTI, setting rcvdInfoWhile to the number of seconds that the information received on a Port will be held before it is either refreshed by receipt of a further configuration message or aged out.~~

   1) The value assigned to rcvdInfoWhile is three times the Hello Time, if Message Age, incremented by 1 second and rounded to the nearest whole second, does not exceed Max Age, and is zero otherwise or if the remainingHops component of portTimes is less than or equal to zero.
   2) ~~mstiPortTimes~~ The remainingHops component of portTimes is set equal to the value received (~~mstiMsgTimes~~msgTimes) decremented by one~~. The value assigned to rcvdInfoWhile is:~~.
~~a) 3 times the Hello Time,~~

~~or~~

~~b) zero, if mstiPortTimes remainingHops is less than or equal to zero.~~

~~The value of Hello Time used by this procedure is taken from cistPortTimes.~~

~~**13.27.27 updtRolesCist()**~~

~~This procedure calculates the following CIST Spanning Tree priority vectors (13.9, 13.10) and timer values:~~

~~a) The *root path priority vector* for each Bridge Port that is not Disabled and has a *port priority vector* (cistPortPriority plus portId—see 13.24.26 and 13.24.25) that has been recorded from a received message and not aged out (infoIs == Received); and~~
~~b) The Bridge's *root priority vector* (cistRootPortId, cistRootPriority—13.23.5, 13.23.6), chosen as the best of the CIST Spanning Tree priority vectors comprising the Bridge's own *bridge priority vector* (CistBridgePriority—13.23.3) plus all the calculated root path priority vectors whose DesignatedBridgeID component is not equal to the DesignatedBridgeID component of the Bridge's own bridge priority vector (13.10); and~~
~~c) The Bridge's *root times*, (cistRootTimes—13.23.7), determined as follows:~~
   ~~1) If the chosen root priority vector is the bridge priority vector, *root times* is equal to CistBridgeTimes (13.23.4).~~

~~2) If the chosen root priority vector is not the bridge priority vector, *root times* is equal to the value of cistPortTimes (13.24.27) for the port associated with the chosen root priority vector.~~

~~d) The *designated priority vector* (cistDesignatedPriority—13.24.4) for each port; and~~

3) The ~~*designated times* for each Port (cistDesignatedTimes—13.24.5)~~ Message Age component of portTimes is set equal to the value ~~of *root times*~~received (msgTimes).

~~The CIST port role for each Port is assigned, and its port priority vector and Spanning Tree timer information are updated as follows:~~

~~e) If the Port is Disabled (infoIs = Disabled), selectedRole is set to DisabledPort. Otherwise:~~
~~f) If the port priority vector information was aged (infoIs = Aged), updtInfo is set and selectedRole is set to DesignatedPort;~~
~~g) If the port priority vector was derived from another port on the Bridge or from the Bridge itself as the Root Bridge (infoIs = Mine), selectedRole is set to DesignatedPort. Additionally, updtInfo is set if the port priority vector differs from the designated priority vector or the Port's associated timer parameters differ from those for the Root Port;~~
~~h) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), and the root priority vector is now derived from it, selectedRole is set to RootPort and updtInfo is reset;~~
~~i) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), the root priority vector is not now derived from it, the designated priority vector is not better than the port priority vector, and the designated bridge and designated port components of the port priority vector do not reflect another port on this bridge, selectedRole is set to AlternatePort and updtInfo is reset;~~
~~j) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), the root priority vector is not now derived from it, the designated priority vector is not better than the port priority vector, and the designated bridge and designated port components of the port priority vector reflect another port on this bridge, selectedRole is set to BackupPort and updtInfo is reset.~~
~~k) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), the root priority vector is not now derived from it, the designated priority vector is better than the port priority vector, selectedRole is set to DesignatedPort and updtInfo is set.~~

~~NOTE—The port role assignment is almost identical to that in 17.19.8 of IEEE Std 802.1D, 1998 Edition, with the addition of the final bullet, erroneously omitted from that specification.~~

For a given MSTI and Port:

1) The value assigned to rcvdInfoWhile is three times the Hello Time or zero if the remainingHops component of portTimes is less than or equal to zero.
2) The remainingHops component of portTimes is set equal to the value received (msgTimes) decremented by one.

The values of Message Age, Max Age, and Hello Time used in these calculations are taken from the CIST's portTimes parameter (13.24.27).

13.27.28 ~~updtRolesMsti~~updtRolesTree()

This procedure calculates the following ~~MST~~ Spanning Tree priority vectors (13.9, 13.10 for the CIST, 13.11 for a MSTI) and timer ~~values~~ values, for ~~an~~ the CIST or a given MSTI:

a) The *root path priority vector* for each Bridge Port that is not Disabled and has a *port priority vector* (~~mstiPortPriority~~ portPriority plus portId—see ~~13.24.19~~13.24.26 and 13.24.25) that has been recorded from a received message and not aged out (infoIs == Received); and

b) The Bridge's *root priority vector* (~~mstiRootPortId~~rootPortId, ~~mstiRootPriority~~rootPriority—~~13.23.11~~13.23.5, ~~13.23.12~~13.23.6), chosen as the best of the ~~MSTI Spanning Tree~~ set of priority vectors ~~for this MSTI~~ comprising the Bridge's own *bridge priority vector* (~~MstiBridgePriority~~BridgePriority—~~13.23.9~~13.23.3) plus all the calculated root path priority vectors ~~whose DesignatedBridgeID component is not equal to the DesignatedBridgeID component of the Bridge's own bridge priority vector (13.11); and~~whose:
   1) DesignatedBridgeID Bridge Address component is not equal to that component of the Bridge's own bridge priority vector (13.10) and,
   2) Port's restrictedRole parameter is FALSE (this second condition applying only when calculating the Bridge's *root priority vector* in the CIST); and

c) The Bridge's *root times*, (~~mstiRootTimes~~rootTimes—~~13.23.13~~13.23.7), ~~determined as follows~~set equal to:
   1) ~~If~~ BridgeTimes (13.23.4), if the chosen root priority vector is the bridge priority vector, ~~*root times* is equal to MstiBridgeTimes (13.23.10).~~otherwise
   ~~2) If the chosen root priority vector is not the bridge priority vector, *root times* is equal to the value of mstiPortTimes (13.24.20) for the port associated with the chosen root priority vector.~~
   3) portTimes (13.24.27) for the port associated with the selected root priority vector, with the Message Age component incremented by 1 second and rounded to the nearest whole second.

d) The *designated priority vector* (~~mstiDesignatedPriority~~designatedPriority—~~13.24.11~~13.24.4) for each port; and

e) The *designated times* (designatedTimes—13.24.5) for each Port ~~(mstiDesignatedTimes—13.24.12)~~ set equal to the value of *root times*.

The CIST or MSTI port role for each Port is assigned, and its port priority vector and Spanning Tree timer information are updated as follows:

f) If the Port is Disabled (infoIs = Disabled), selectedRole is set to DisabledPort~~;~~. Otherwise:

g) If this procedure is invoked for a given MSTI:
   1) If the Port is not Disabled, the selected CIST Port Role (calculated for the CIST prior to invoking this ~~procedure~~procedure for a given MSTI) is RootPort, and the CIST port priority information was received from a Bridge external to the MST Region (infoIs == Received and infoInternal == FALSE), selectedRole is set to MasterPort. Additionally, updtInfo is set if the MSTI port priority vector differs from the designated priority vector or the Port's associated timer parameters differ from those for the Root Port;
   2) If the Port is not Disabled, the selected CIST Port Role (calculated for the CIST prior to invoking this ~~procedure~~procedure for a given MSTI) is AlternatePort, and the CIST port priority information was received from a Bridge external to the MST Region (infoIs == Received and infoInternal == FALSE), selectedRole is set to AlternatePort. Additionally, updtInfo is set if the MSTI port priority vector differs from the designated priority vector or the Port's associated timer parameters differ from those for the Root Port.

Otherwise, if the Port is not Disabled and the CIST port priority information was not received from a Bridge external to the Region (infoIs != Received or infoInternal == TRUE)~~:~~ :

h) If the port priority vector information was aged (infoIs = Aged), updtInfo is set and selectedRole is set to DesignatedPort;

i) If the port priority vector was derived from another port on the Bridge or from the Bridge itself as the Root Bridge (infoIs = Mine), selectedRole is set to DesignatedPort. Additionally, updtInfo is set if the port priority vector differs from the designated priority vector or the Port's associated timer parameters differ from those for the Root Port;

j) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), and the root priority vector is now derived from it, selectedRole is set to RootPort and updtInfo is reset;

k) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), the root priority vector is not now derived from it, the designated priority vector is not better than the port priority vector, and the designated bridge and designated port components of the port priority vector do not reflect another port on this bridge, selectedRole is set to AlternatePort and updtInfo is reset;

l) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), the root priority vector is not now derived from it, the designated priority vector is not better than the port priority vector, and the designated bridge and designated port components of the port priority vector reflect another port on this bridge, selectedRole is set to BackupPort and updtInfo is reset~~.~~;

m) If the port priority vector was received in a Configuration Message and is not aged (infoIs == Received), the root priority vector is not now derived from it, the designated priority vector is better than the port priority vector, selectedRole is set to DesignatedPort and updtInfo is set.

13.27.29 ~~updtRolesDisabledTree~~uptRolesDisabledTree()

This procedure sets selectedRole to DisabledPort for all Ports of the Bridge for a given tree (CIST or MSTI).

<< Editor's note: figures in clauses 13.29, 13.32 and 13.33 below are adapted from IEEE P802.1Q-REV/D0.0, and figures in clauses 13.35 and 13.37 below are adapted from IEEE Std 802.1D, 2004 Edition.

These figures represent:

a) In black, state machine elements as specified in the corresponding draft/standard;

b) In blue, additions to the state machines specified in the corresponding draft/standard;

c) In red, suppressions from the state machines specified in the corresponding draft/standard;

d) In orange, state machine elements subject to change. >>

13.28 The Port Timers state machine

The Port Timers state machine for a given Port is responsible for decrementing the timer variables for the CIST and all MSTIs for that Port each second.

~~The Port Timers state machine shall implement the function specified by the state diagram contained in Figure 13-21.~~

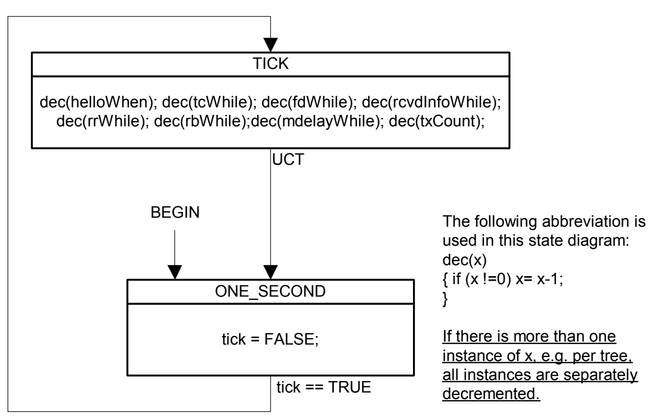

Figure 13-21—Port Timers state machine

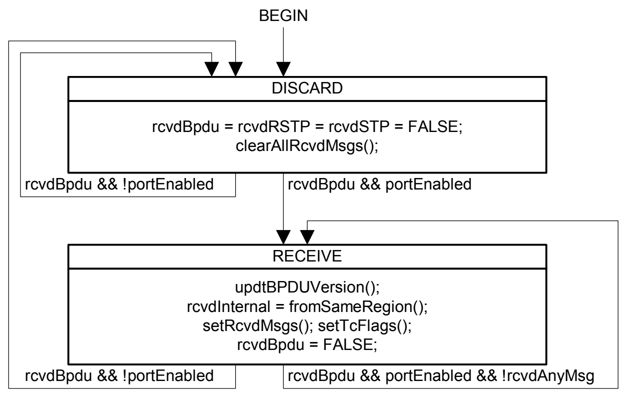

Figure 13-22—Port Receive state machine

The specification of this state machine is identical to that of the Port Timers state machine for RSTP (17.22 of IEEE Std 802.1D, 2004 Edition).

13.29 Port Receive state machine

The Port Receive state machine shall implement the function specified by the state diagram contained in Figure 13-23 and the attendant definitions contained in 13.21 through 13.27.

This state machine is responsible for receiving BPDUs. Its specification is identical to that of the Bridge Detection state machine for RSTP (17.23 of IEEE Std 802.1D, 2004 Edition) but it:

a) ~~This state machine is responsible for receiving BPDUs. Using the updtBPDUversion() procedure it sets either the rcvdRSTP or the rcvdSTP flag to reflect the type of the BPDU for use by the Port Protocol Migration machine. It sets~~ Sets the rcvdInternal flag if the transmitting Bridge and BPDU are internal to the MST Region, i.e. belong to the same MST Region as this Bridge (13.8)~~. It sets rcvdMsg for the CIST, and additionally for each MSTI if the rcvdBPDU is internal. The next BPDU is not processed until all the rcvdMsg flags have been cleared by the per tree state machines~~.

b) Sets rcvdMsg for the CIST, and additionally for each MSTI if the rcvdBPDU is internal.

The next BPDU is not processed until all the rcvdMsg flags have been cleared by the per tree state machines.

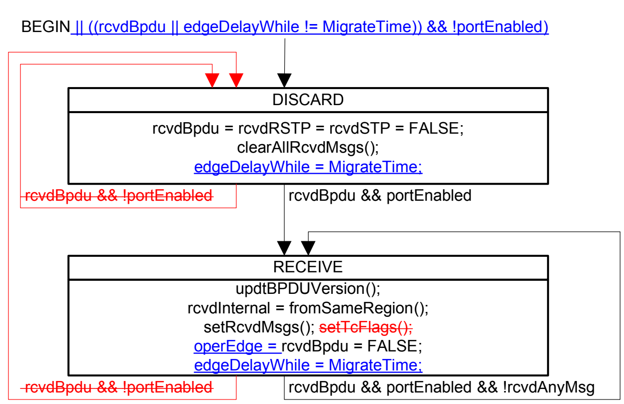

Figure 13-23—Port Receive state machine

13.30 Port Protocol Migration state machine

The specification of this state machine is identical to that of the Port Protocol Migration state machine for RSTP (Clause 17 of IEEE Std 802.1D, 1998 Edition).

13.31 ~~Port Transmit~~ Bridge Detection state machine

The specification of this state machine is identical to that of the Bridge Detection state machine for RSTP (17.25 of IEEE Std 802.1D, 2004 Edition).

<< Editor's note: in the aforementioned 2004 standard, BDM should be corrected as follows:

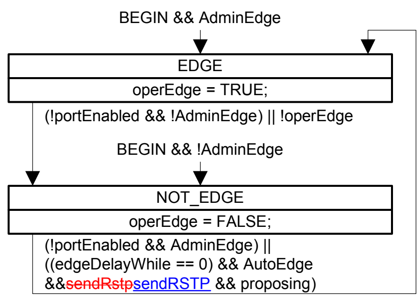

>>

13.32 Port Transmit state machine

The Port Transmit state machine shall implement the function specified by the state diagram contained in Figure 13-25 and the attendant definitions contained in 13.21 through 13.27.

This state machine is responsible for transmitting BPDUs.

NOTE 1—Any single received BPDU that changes the CIST Root Identifier, CIST External Root Path Cost, or CIST Regional Root associated with MSTIs should be processed in their entirety, or not at all, before encoding BPDUs for transmission.This recommendation is made to minimize the number of BPDUs to be transmitted following receipt of a BPDU carrying new information. It is not required for correctness and has not therefore been incorporated into the state machines.

NOTE 2—If a CIST state machine sets ~~newInfoCist~~newInfo, this machine will ensure that a BPDU is transmitted conveying the new CIST information. If MST BPDUs can be transmitted through the port this BPDU will also convey new MSTI information for all MSTIs. If ~~an~~ a MSTI state machine sets newInfoMsti, and MST BPDUs can be transmitted ~~though~~ through the port, this machine will ensure that a BPDU is transmitted conveying information for the CIST and all MSTIs. Separate ~~newInfoCist~~ newInfo and newInfoMsti variables are provided to avoid requiring useless transmission of a BPDU through a port that can only transmit STP BPDUs (as required by the ~~ForceVersion~~ Force Protocol Version parameter or Port Protocol Migration machine) following a change in MSTI information without any change to the CIST.

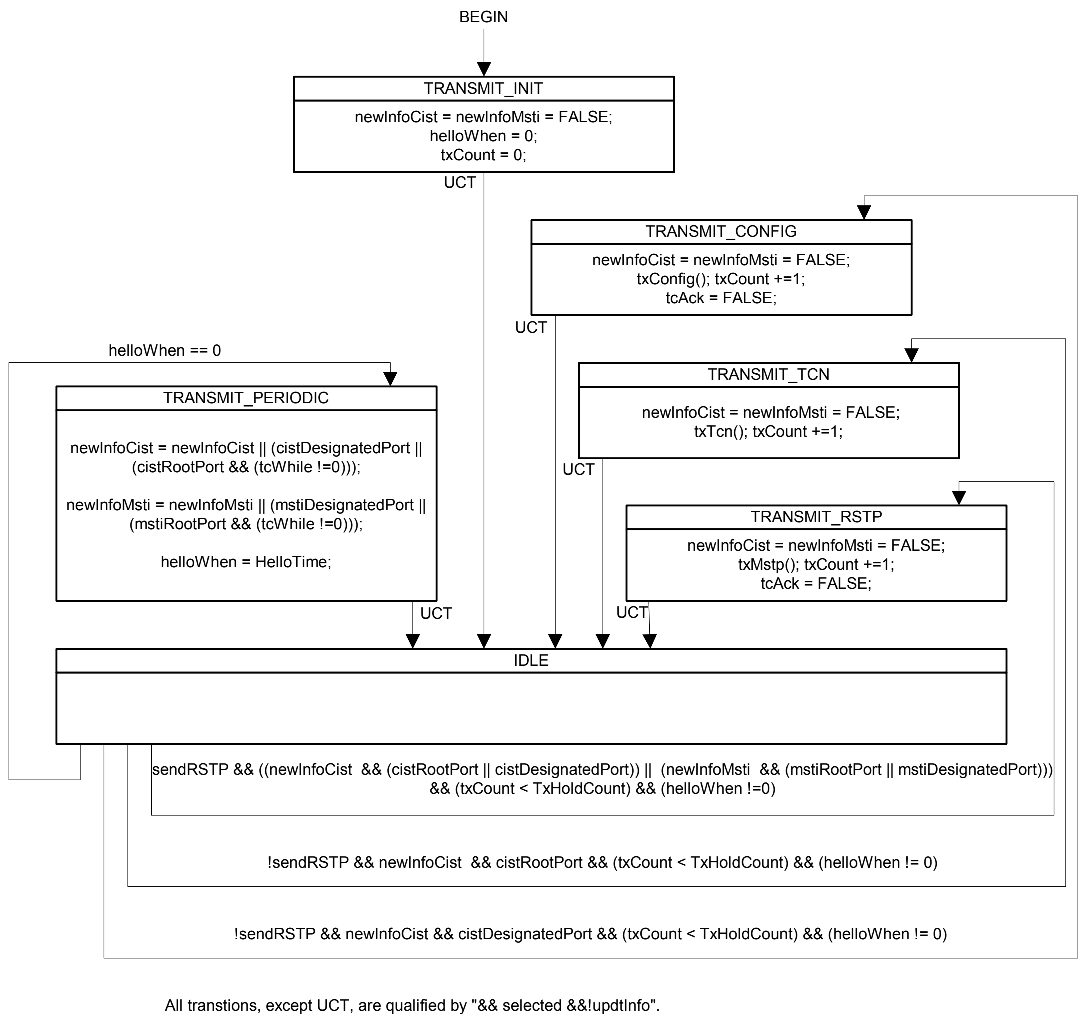

Figure 13-24—Port Transmit state machine

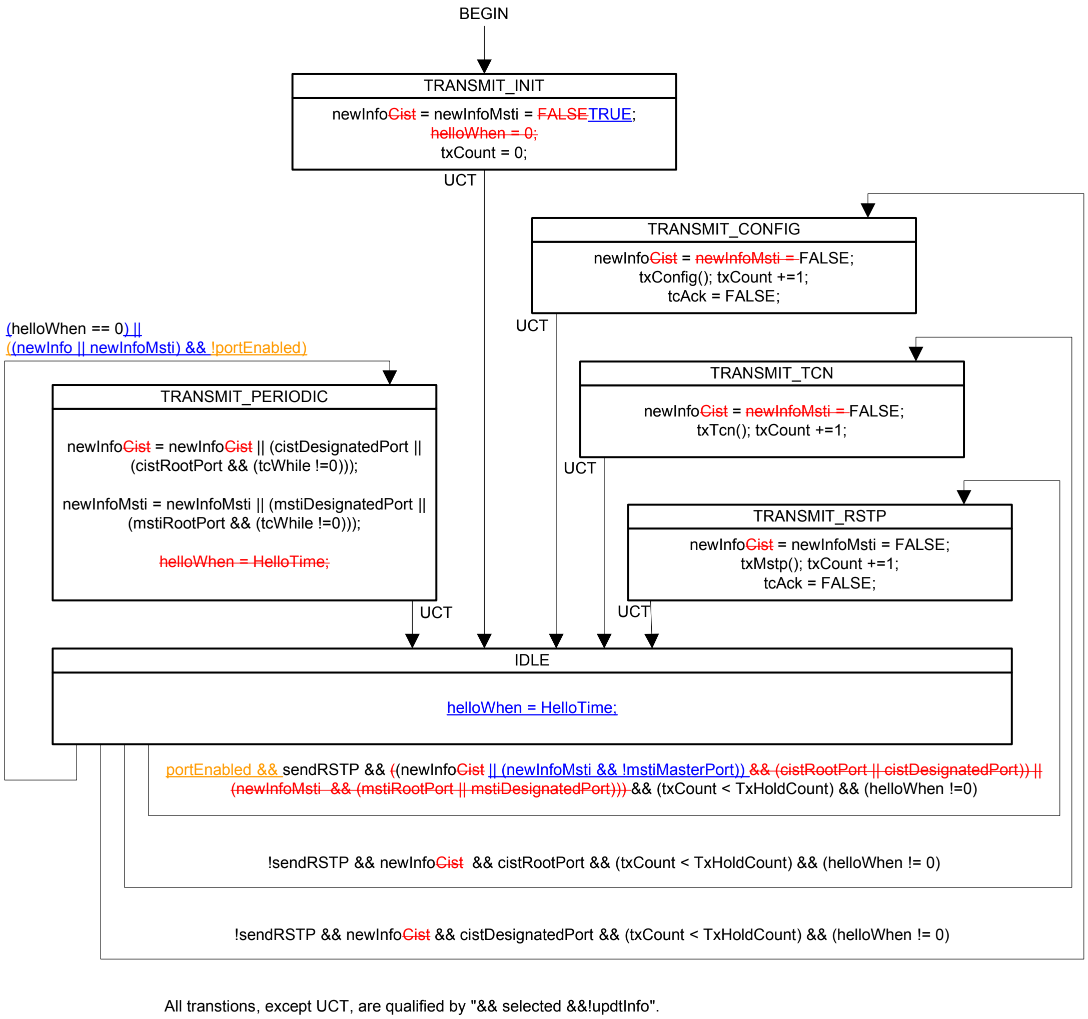

Figure 13-25—Port Transmit state machine

13.33 Port Information state machine

The Port Information state machine for each tree shall implement the function specified by the state diagram contained in Figure 13-27 and the attendant definitions contained in 13.21 through 13.27.

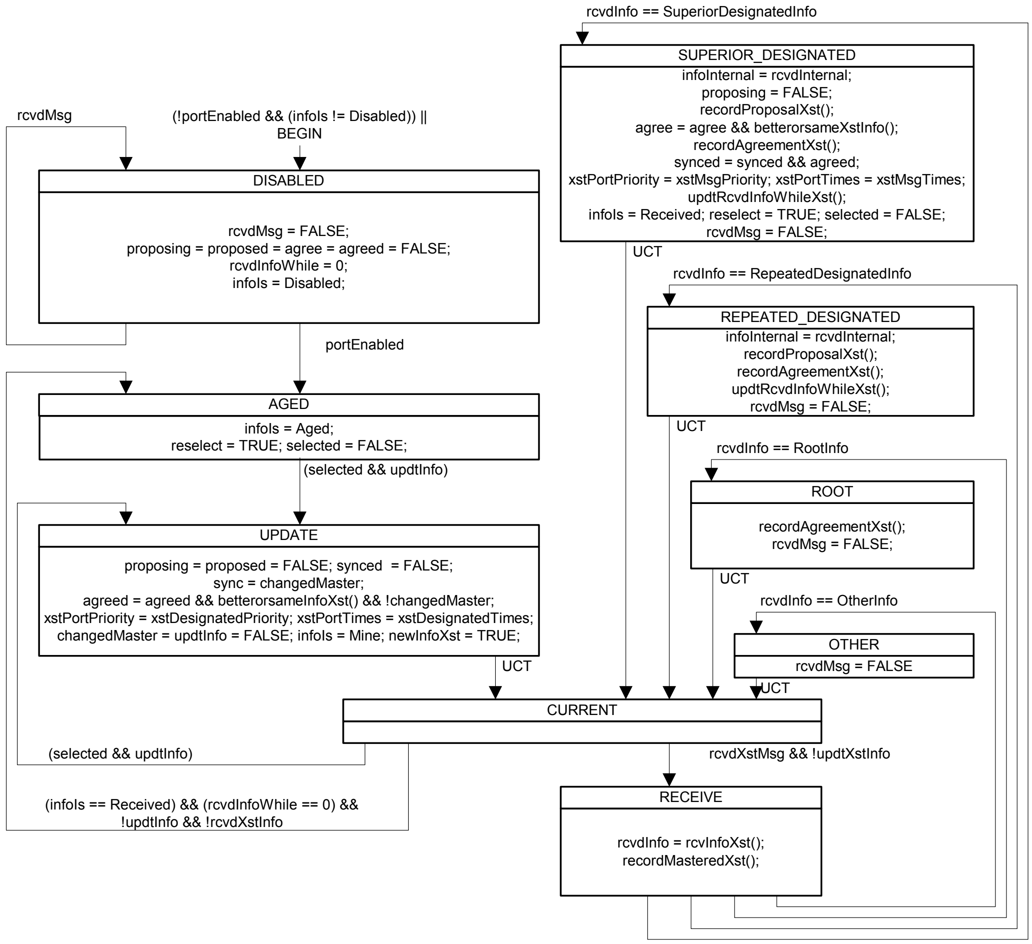

Figure 13-26—Port Information state machine

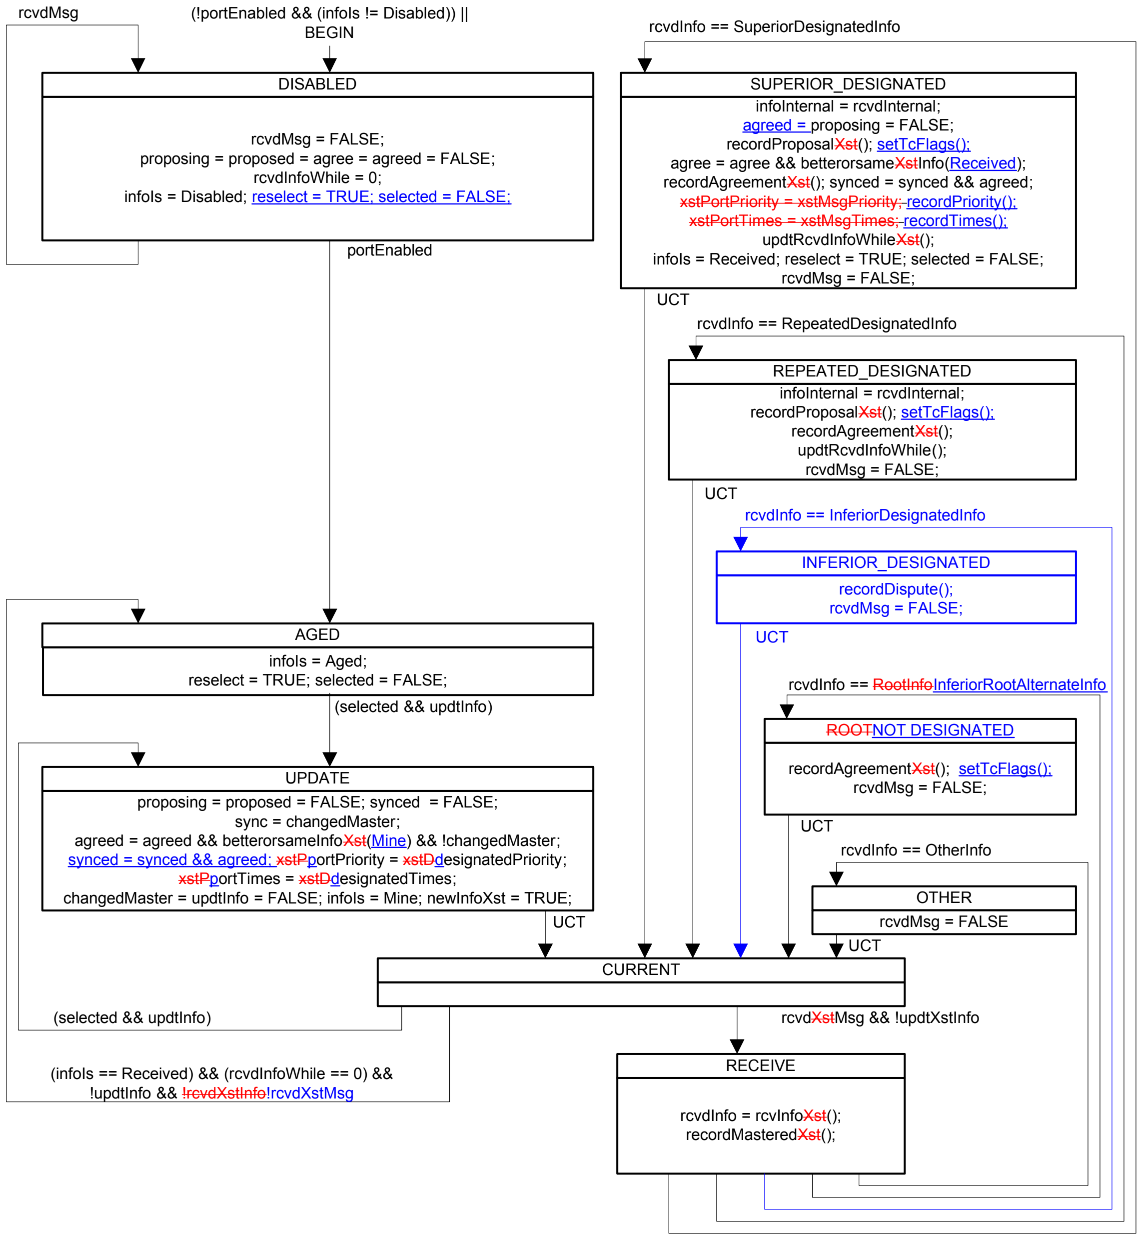

Figure 13-27—Port Information state machine

This state machine is responsible for recording the Spanning Tree information currently in use by the CIST or a given MSTI for a given Port, ageing that information out if it was derived from an incoming BPDU, and recording the origin of the information in the infoIs variable. The selected variable is cleared and reselect set to signal to the Port Role Selection machine that port roles need to be recomputed. The infoIs and portPriority variables from all ports are used in that computation and, together with portTimes, determine new values of designatedPriority and designatedTimes. The selected variable is set by the Port Role Selection machine once the computation is complete.

13.34 Port Role Selection state machine

~~The Port Role Selection machine for each tree shall implement the function specified by the state diagram contained in Figure 13-28 and the attendant definitions contained in 13.21 through 13.27.~~

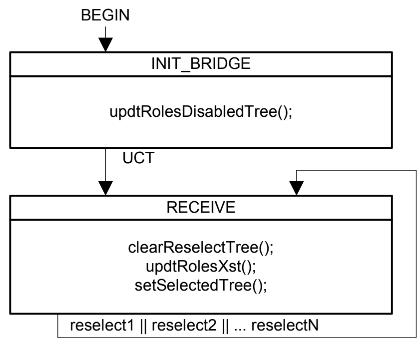

Figure 13-28—Port Role Selection

~~NOTE—~~The specification of this state machine is identical to that of the Port Role Selection state machine for RSTP (~~Clause 17~~ 17.28 of IEEE Std 802.1D, ~~1998~~ 2004 Edition) but makes use of revised procedures as detailed in Clause 13.27 of this specification~~. These procedures have also been renamed relative to IEEE Std 802.1D, 1998 Edition; updtRoleDisabledBridge to updtRoleDisabledTree, clearReselectBridge to clearReselectTree, updtRolesBridge to updtRolesCist and updtRolesMsti, and setSelectedBridge to setSelectedTree~~.

13.35 Port Role Transitions state machine

The Port Role Transitions state machine shall implement the function specified by the state diagram contained in:

a) Part 1: Figure 17-20 of IEEE Std 802.1D, 2004 Edition for both the initialization of this state machine and the states associated with the DisabledPort role; and
b) Part 2: Figure 13-31 for the states associated with the MasterPort role; and
c) Part 3: Figure 13-32 for the states associated with the RootPort role; and
d) Part 4: Figure 13-33 for the states associated with the DesignatedPort role; and
e) Part 5: Figure 13-34 for the states associated with the AlternatePort and BackupPort roles;

~~The Port Role Transitions state machine shall implement the function specified by the state diagram contained in Figure 13-31~~ and ~~Figure 13-32 and~~ the attendant definitions contained in 13.21 through 13.27.

As Figure 17-20 of IEEE Std 802.1D, 2004 Edition, Figure 13-31, Figure 13-32, Figure 13-33, and Figure 13-34 are component parts of the same state machine, the global transitions associated with both diagrams are possible exit transitions from the states shown in either of the diagrams.

~~As Figure 13-31 and Figure 13-32 are component parts~~ Figure 17-20 of ~~the same state machine~~IEEE Std 802.1D, ~~the global transitions associated with both diagrams are possible exit transitions from the states shown in either of the diagrams. Figure 13-31~~2004 Edition and Figure 13-34 ~~shows~~ show the Port Roles for Ports that do not form part of the active topology of the given Tree.

Figure 13-31, Figure 13-32, and Figure 13-33 ~~shows~~ show the Port Roles that form part of the active topology.

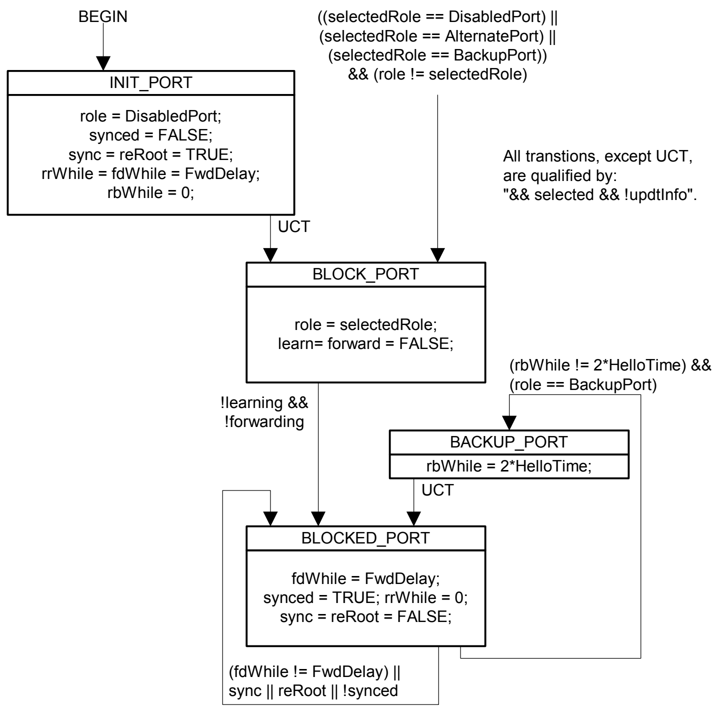

**Figure 13-29—Port Role Transitions state machine—
Part 1: Disabled, Alternate, and Backup Roles**

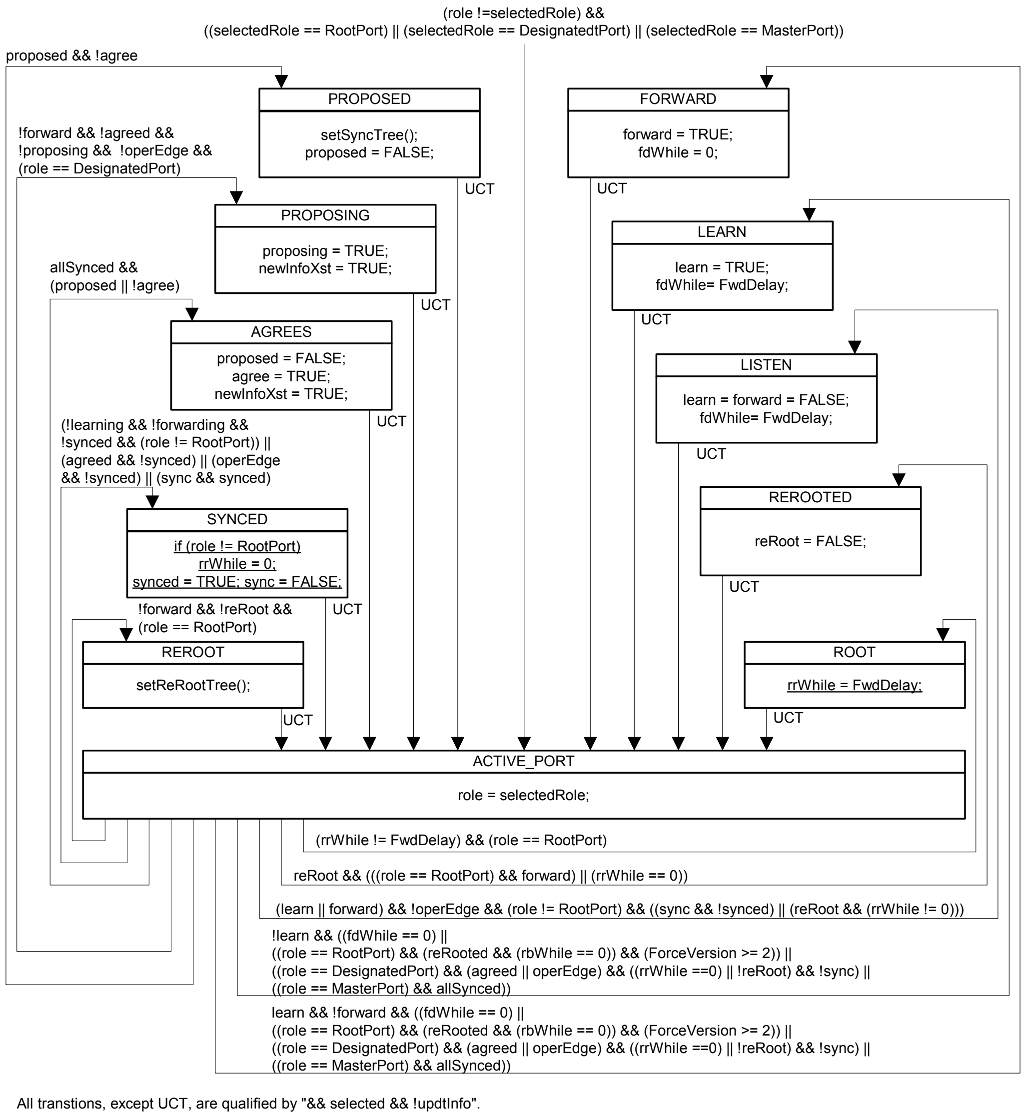

**Figure 13-30—Port Role Transitions state machine—
Part 2: Root, Designated, and Master Roles**

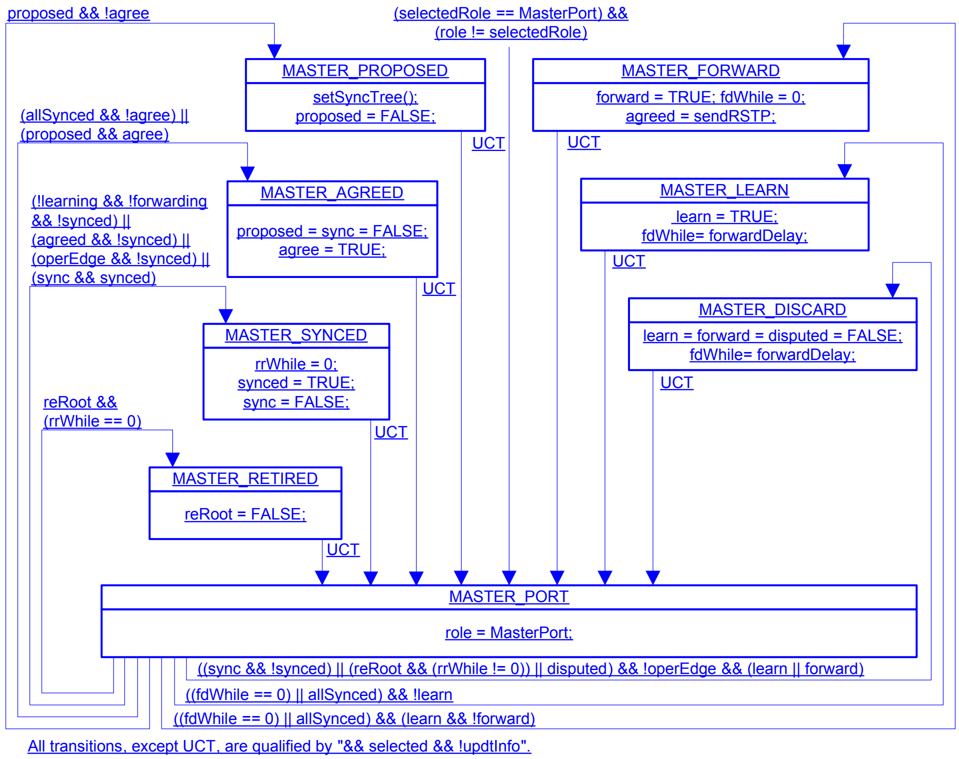

**Figure 13-31—Port Role Transitions state machine—
Part 2: MasterPort role transitions**

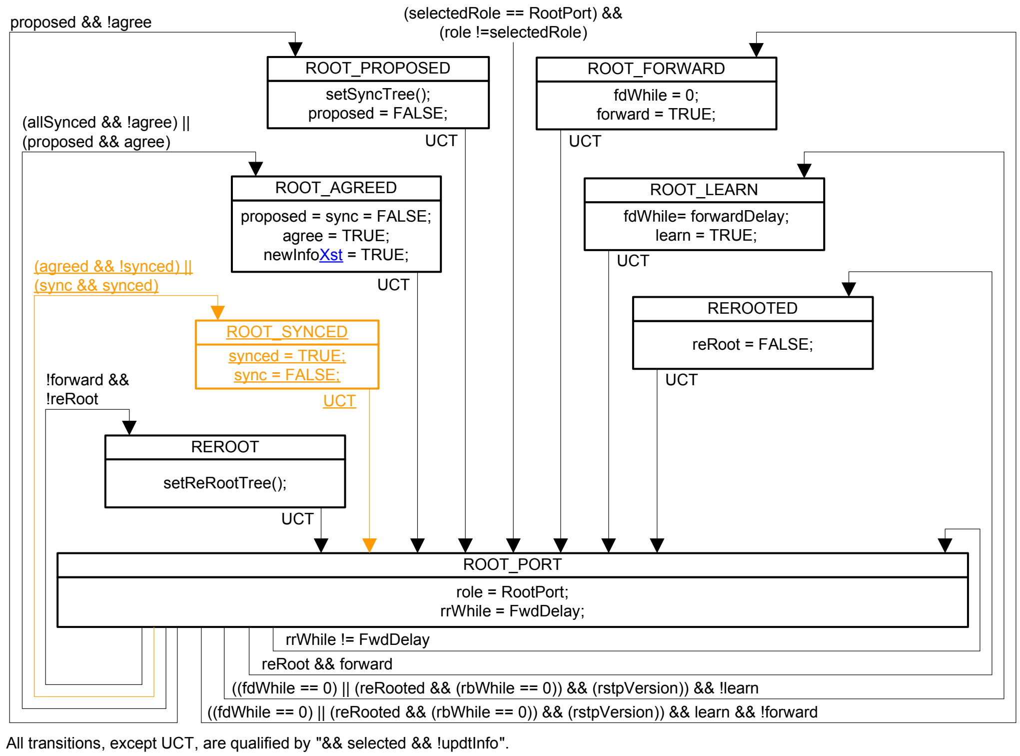

**Figure 13-32—Port Role Transitions state machine—
Part 3: RootPort role transitions**

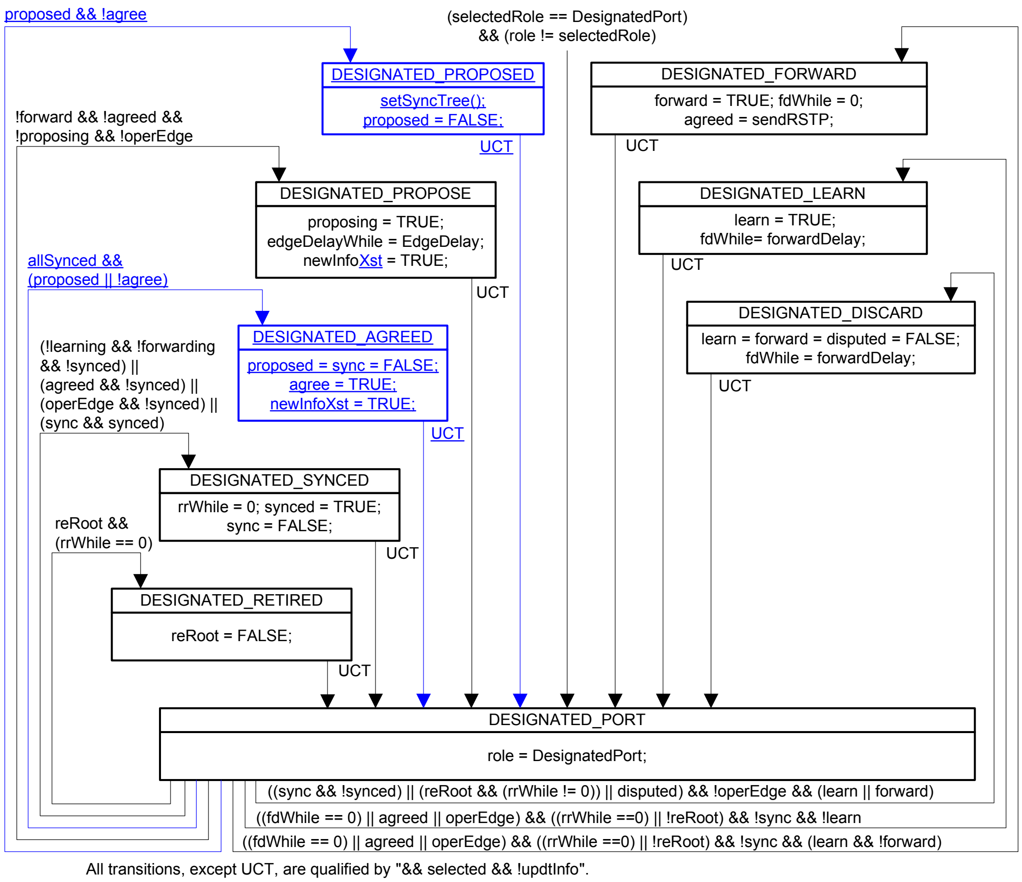

**Figure 13-33—Port Role Transitions state machine—
Part 4: DesignatedPort role transitions**

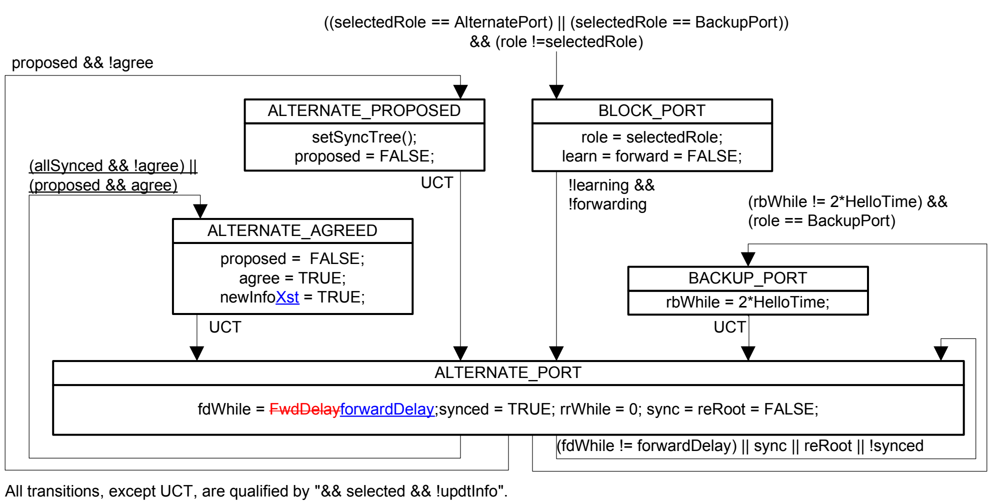

**Figure 13-34—Port Role Transitions state machine—
Part 5: AlternatePort and BackupPort role transitions**

13.36 Port State Transition state machine

~~The Port State Transition state machine for each tree shall implement the function specified by the state diagram contained in Figure 13-35 and the attendant definitions contained in 13.21 through 13.27 above.~~

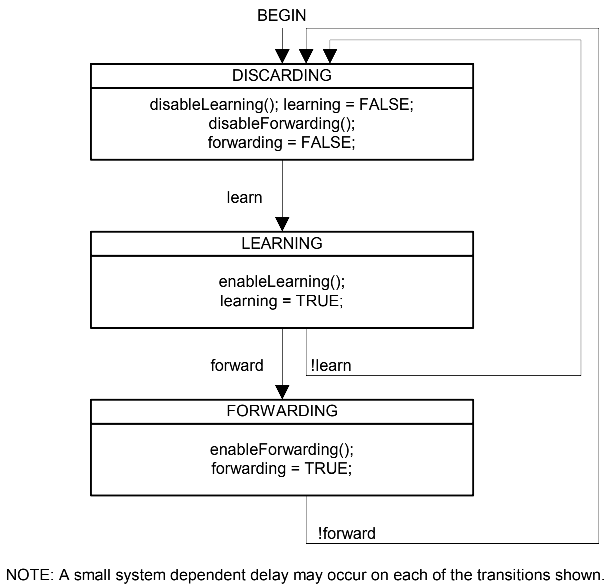

Figure 13-35—Port State Transition state machine

The specification of this state machine is identical to that of the Port State Transition state machine for RSTP (17.30 of IEEE Std 802.1D, 2004 Edition).

NOTE—A small system dependent delay may occur on each of the transitions shown in the referenced state machine.

13.37 Topology Change state machine

The Topology Change state machine for each tree shall implement the function specified by the state diagram contained in Figure 13-36 and the attendant definitions contained in 13.21 through 13.27.

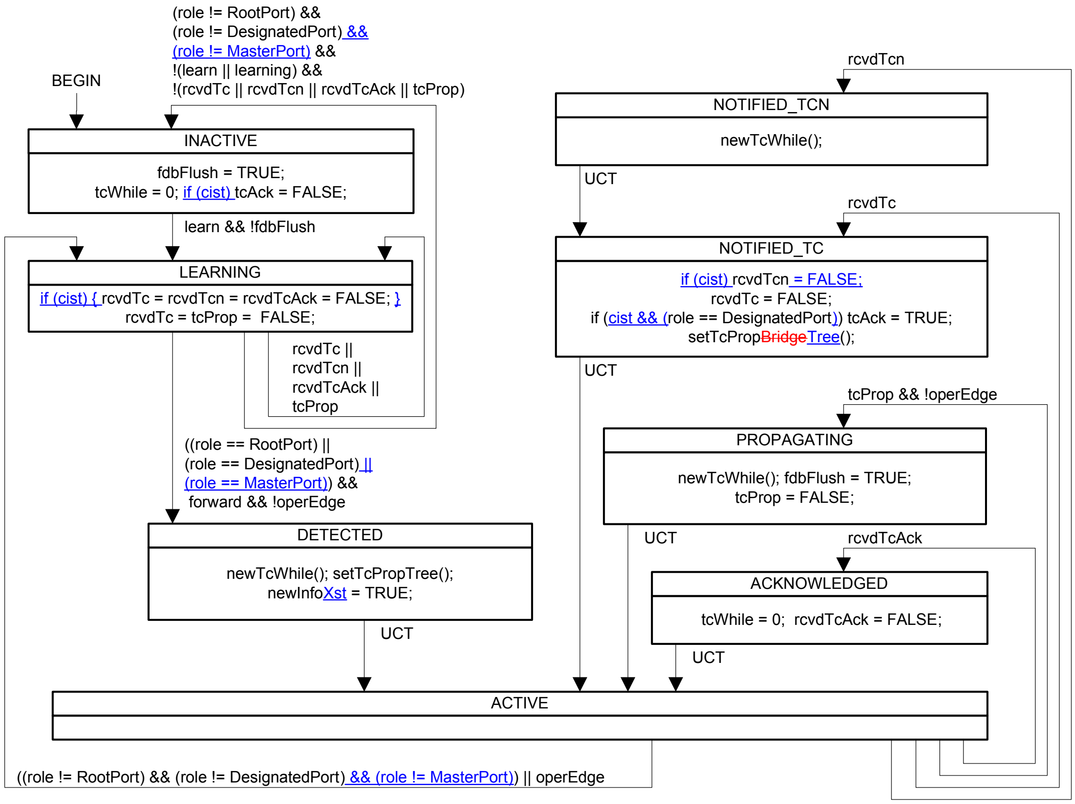

Figure 13-36—Topology Change state machine

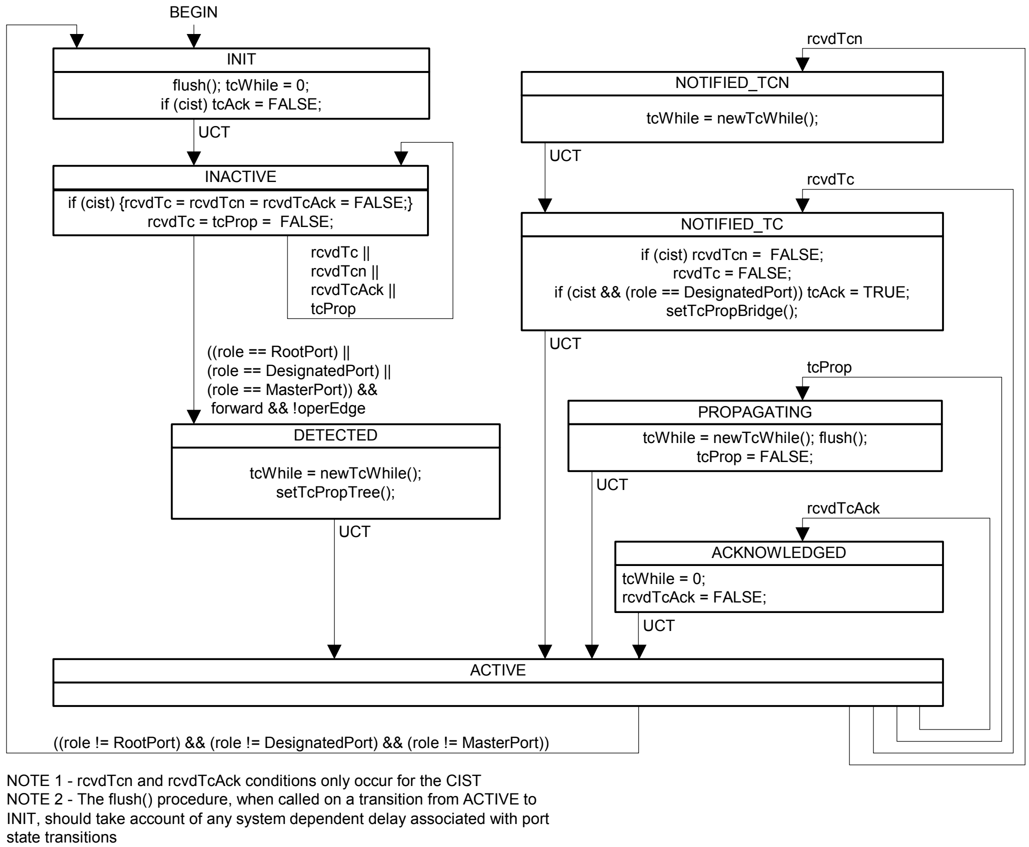

Figure 13-37—Topology Change state machine

13.38 Performance

This subclause places requirements on the setting of the parameters of MSTP. It recommends default operational values for performance parameters. These have been specified in order to avoid the need to set values prior to operation, and have been chosen with a view to maximizing the ease with which Bridged LAN components interoperate.

The constraints on parameter values for correct operation are essentially the same as specified in IEEE Std 802.1D, 2004 Edition for RSTP, and provide for the integration of MST Bridges into LANs that contain Bridges using STP. Bridges using MSTP shall conform to the parameter value requirements of 17.14 and 17.~~39~~ 32 of IEEE Std 802.1D, ~~1998~~ 2004 Edition. Implementations of MSTP and managers of Bridged Local Area Networks should note the recommendations of that Clause. Maximum, minimum, default, and or applicable ranges are specified and recommended for values of the following parameters:

a) Maximum Bridge Diameter
b) Maximum Bridge Transit Delay
c) Maximum BPDU Transmission Delay
d) Maximum Message Age Increment Overestimate
e) Bridge Priority and Port Priority
f) External Port Path Cost (referred to as Path Cost in 17.~~39~~ 14 of IEEE Std 802.1D, ~~1998~~ 2004 Edition)

g) Port Hello Time (HelloTime in 13.22)

NOTE—In MSTP Bridges, Hello Time is manageable on a per-Port basis, rather than per-Bridge in STP and RSTP. Port Hello Time therefore replaces the Bridge Hello Time parameter found in STP and RSTP.

h) Bridge Max Age
i) Bridge Forward Delay

This standard also makes recommendations on the applicable values of Internal Port Path Cost, a parameter specific to MSTP.

13.38.1 Internal Port Path Costs

It is recommended that default values of the Internal Port **Path Cost** parameter for each Bridge Port be based on the values shown in Table 13-3, the values being chosen according to the speed of the LAN segment to which each Port is attached.

Table 13-3—Internal Port Path Costs

Parameter	Link Speed	Recommended value	Recommended range	Range
Internal Port Path Cost	<=100 Kb/s	200 000 000	20 000 000 – 200 000 000	1 – 200 000 000
	1 Mb/s	20 000 000	2 000 000 – 200 000 000	1 – 200 000 000
	10 Mb/s	2 000 000	200 000 – 20 000 000	1 – 200 000 000
	100 Mb/s	200 000	20 000 – 2 000 000	1 – 200 000 000
	1 Gb/s	20 000	2 000 – 200 000	1 – 200 000 000
	10 Gb/s	2 000	200 – 20 000	1 – 200 000 000
	100 Gb/s	200	20 – 2 000	1 – 200 000 000
	1 Tb/s	20	2 – 200	1 – 200 000 000
	10 Tb/s	2	1 – 20	1 – 200 000 000

Where intermediate link speeds are created as a result of the aggregation of 2 or more links of the same speed (see IEEE Std 802.3ad), it may be appropriate to modify the recommended values shown to reflect the change in link speed. However, as the primary purpose of the Path Cost is to establish the active topology of the network, it may be inappropriate for the Path Cost to track the effective speed of such links too closely, as the resultant active topology may differ from that intended by the network administrator. For example, if the network administrator had chosen an active topology that makes use of aggregated links for resilience (rather than for increased data rate), it would be inappropriate to cause a Spanning Tree topology change as a result of one of the physical links in an aggregation failing. Similarly, with links that can autonegotiate their data rate, reflecting such changes of data rate in changes to Path Cost may not be appropriate, depending upon the intent of the network administrator. Hence, as a default behavior, such dynamic changes of data rate should not automatically cause changes in Path Cost for the Port concerned.

NOTE 1—The values shown apply to both full duplex and half duplex operation. The intent of the recommended values and ranges shown is to minimize the number of Bridges in which path costs need to be managed in order to exert control over the topology of the Bridged LAN.

NOTE 2—The values shown are the same as those recommended in IEEE Std 802.~~1t-2001~~1D, 2004 Edition.