INTERNET TEACHING LAB: SPANNING TREE PROTOCOL

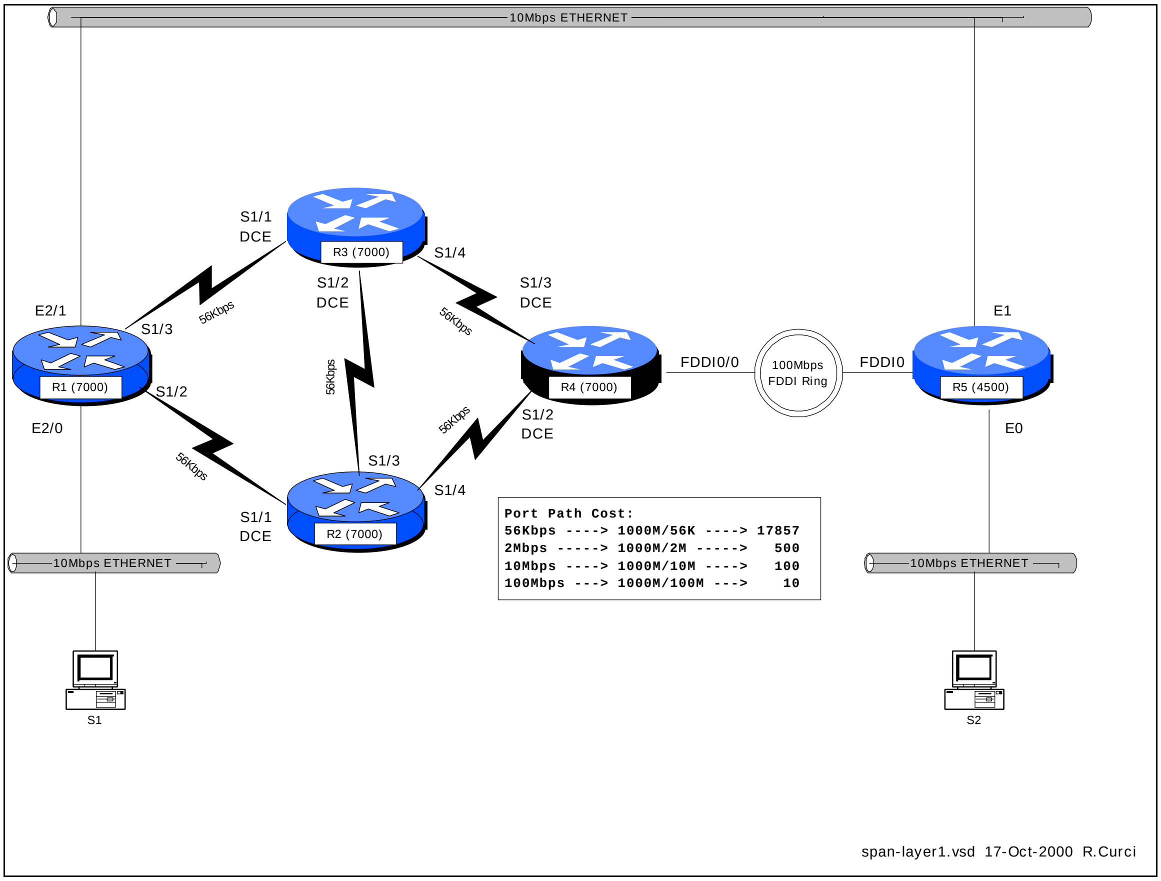

```
Port Path Cost:
56Kbps ----> 1000M/56K ----> 17857
2Mbps -----> 1000M/2M ----->   500
10Mbps ----> 1000M/10M ---->   100
100Mbps ---> 1000M/100M --->    10
```

span-layer1.vsd 17-Oct-2000 R.Curci

Overview

The Spanning Tree Protocol, also known as the Djistrja's Algorithm, is documented in the IEEE 802.1D standard. It is implemented in many current routers, bridges, and switches to provide a loop-free network topology. It is popular to build layer2 networks with redundant network connections to improve reliability, but the redundancy can lead to broadcast storms. Spanning Tree Protocol provides a mechanism for network devices to learn the network topology, elect a root bridge, and selectively block ports to form a loop-free spanning tree. We will explore some of the capabilities of this protocol, advantages, and limitations. The IEEE spanning tree protocol was first implemented in the DEC LAN Bridge 100 in the mid 1980s by Dr. Radia Perlman whose text book, *Interconnections*, now in the second edition, is the definitive reference.

Configuration

We will explore the Cisco Router implementation of 802.1D. Set up the physical cabling as specified in diagram above. The initial configuration for all five routers is listed at the end of this document also also on text file *span-config.txt*. Log into each of the five routers R1, R2, R3, R4, and R5, go into router configuration mode, and paste the appropriate configuration commands. Verify that all appropriate interfaces are up and that everything is cabled to the correct routers and ports. Use the commands “show ip interface”, “show ip interface brief”, and “show cdp neighbors” for verification.

Setup PCs

Configure PCs S1 and S2 with IP addresses in the same IP network. Verify that you can PING between the two PCs. (Hint: If this does not work you can test the PCs by temporarily connecting them to the same physical Ethernet segment or by using a 10baseT Ethernet crossover cable. You may have difficulty if your router interface accidently has an IP address on one of the bridge interface in which case it may be routing IP protocol and bridging non-IP traffic. You can verify that the router is bridging IP traffic on the appropriate interfaces with the command “show interface crb”)

Try sending a series of PINGs from S1 → S2 using both small 64-byte packets and large 1500-byte packets and note the average round-trip time. Repeat this test while S1 and S2 are temporarily directly connected. Compare the numbers and if substancially different, explain why.

There are redundant connections in your network and we want to determine the physical path between S1 and S2 used by the PING packets. First, determine the Ethernet MAC addresses for the NIC cards in S1 and S2. (Hint: If two devices on the same IP network have recently communicated, you will find each other’s Ethernet MAC address inside their respective ARP caches which can be displayed with the command “arp –a”)
Use the command “show bridge 1” on each router to display the bridge forwarding table and find the S1 and S2 entries. Record the forwarding path on your network diagram.

Bridge IDs and Port Path Cost

Using the command “show span 1”, determine which router is the root bridge and indicate it on your network diagram. This implementation of 802.1D computes the port path cost by dividing 1,000,000,000 by the bandwidth of the port in bits/second. This gives us the following port costs for the connections in your network:

INTERFACE TYPE	BANDWIDTH	PORT PATH COST
56K SERIAL	56,000 bits/sec	17857
10M ETHERNET	10,000,000 bits/sec	100
FDDI	100,000,000 bits/sec	10

Given your diagram, knowledge of the root bridge, and above table, manually compute the spanning tree algorithm. For each bridge port, indicate the port state (F=forwarding, B=blocking) as well as the port type (RP=root port, DP=designated port, NDP=non-designated port).

Verify your calculations by comparing them with the output of the command "show spanning-tree 1" on each router.

Bridge Protocol Data Units

On one of your routers with a blocked bridge port, issue the command "show interface xxx" where xxx is the name of the blocked interface/port. Note the input and output packet counters. Are they incrementing? If so, why are they incrementing? Instead of doing the arithmetic, you may find it easier to "clear counters" to zero the counters before you start.

The Cisco router has a number of debug modes used to diagnose network problems. Although sometimes dangerous to use on a production network, they are very good tools in a lab environment. The command "term monitor" will enable debug messages to be displayed on your router session and disabled with "term no monitor". Try turning on the spanning tree topology change debug with "debug spanning tree" until you collect a few messages, then turn it off with "undebug all". You should see some bridge protocol data unit packets represented in hexadecimal. You should be able to spot the MAC address of your root bridge embedded in the packet. Using the following table, decode the root bridge ID (priority and MAC address), sending bridge ID (priority and MAC address), root path cost, and timers.

FIELD	**OCTETS**	**FUNCTION**
Protocol ID	2	future (always zero)
Version	1	future (always zero)
Type	1	BPDU Type (0=config BPDU)
Flags	1	LSB (topolgy chg flash), MSB (Topology chg ACK)
Root BID	8	Bridge ID of root (16bit priority + 48bit MAC)
Root Path Cost	4	Cumulative cost to root bridge
Sending BID	8	Bridge ID of sender (16bit priority + 48bit MAC)
Port ID	2	Port ID that sent this BPDU
Message Age	2	Age of root BPDU
Max Age	2	Max age to save BPDU info (default = 20s)
Hello Time	2	Time between sending consecutive BPDUs (default = 2s)
Forward Delay	2	Time spent in listening and learning states in FSM (default = 15s)

Finite State Machine

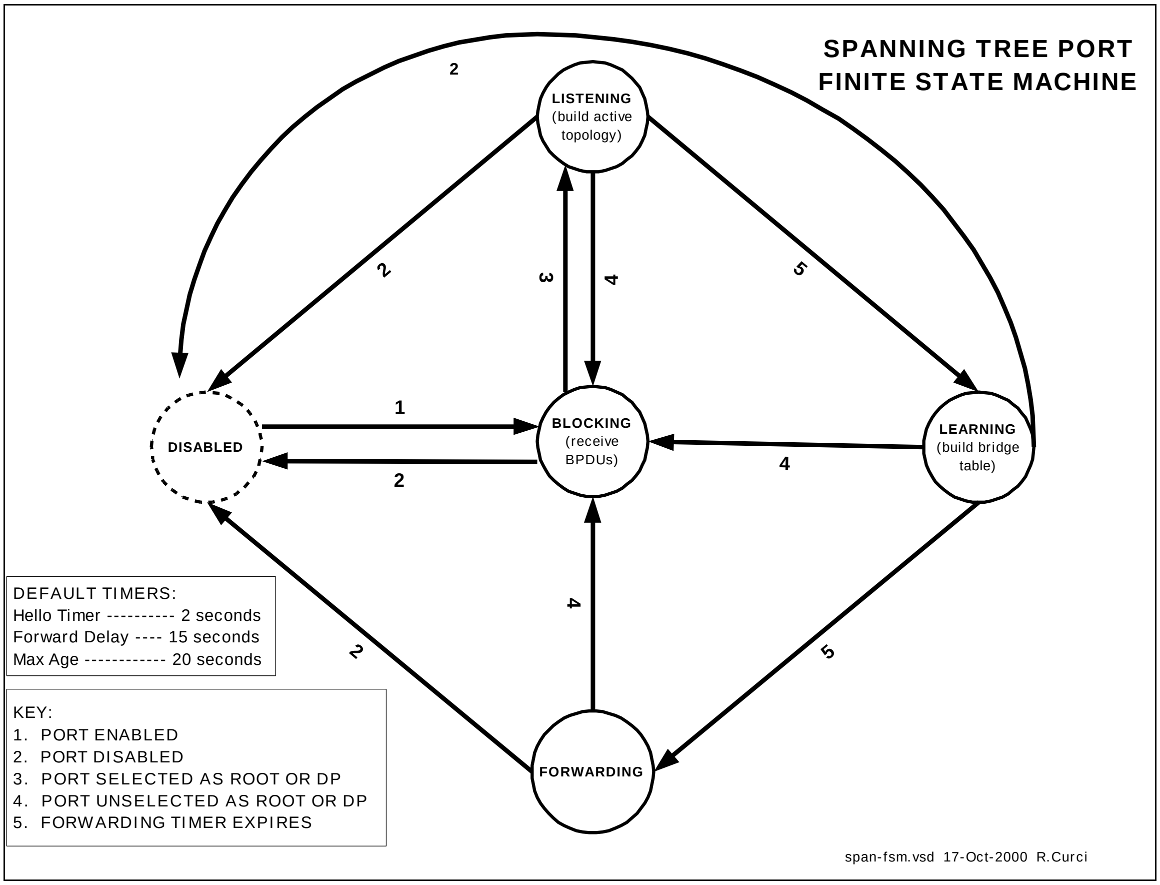
SPANNING TREE PORT
FINITE STATE MACHINE
LISTENING
(build active topology)
BLOCKING
(receive BPDUs)
LEARNING
(build bridge table)
DISABLED
FORWARDING
1
2
3
4
5
DEFAULT TIMERS:
Hello Timer ---------- 2 seconds
Forward Delay ---- 15 seconds
Max Age ------------ 20 seconds
KEY:
1. PORT ENABLED
2. PORT DISABLED
3. PORT SELECTED AS ROOT OR DP
4. PORT UNSELECTED AS ROOT OR DP
5. FORWARDING TIMER EXPIRES
span-fsm.vsd 17-Oct-2000 R.Curci

Bridge ports can be in one of five states: disabled, blocking, listening, learning, and forwarding. See the diagram *span-fsm.pdf* to see what events cause transitions between different states. Log into one of your routers and identify a bridge interface in the forwarding state. Turn on spanning tree topology events debugging with “debug spanning events” and shut down the interface with “interface xyz” and “shutdown”. Wait a minute, then turn it back on with “no shutdown”. Note the state changes as it transitions from the disabled to the forwarding state including intermediate states. Record how much time was spent in each state. Turn off debugging with “undebug all”.

TEST TCP

Locate the program TTCP by searching the Internet. At the time of this writing, it was available for anonymous/ftp download at ftp://FTP.ARL.MIL/pub/ttcp. It is a TCP/IP benchmarking program. There are both C-language versions, usually named ttcp.c, and java implementations that work on Windows systems. You basically start this program on one system in receive mode, then start the other copy in transmit mode and supply the IP address of the receiver. The utility sends several blocks of data (you specify how many blocks and how many bytes per block) then displays statistics in Bytes/Second and Bits/Second on speed of the transfer. Use this tool to measure the network performance from S1 → S2 traversing your network. How many bits per second did you achieve? Study your network diagram paying particular attention to your router link speeds and which interfaces are blocked. As packets traverse your network, your throughput is affected factors such as the speed of the links traversed, congestion, router CPU load and switching method, errors, etc. If you focus on the link speeds, is there a better (faster) path through your network that is not used? Determine which bridge should be made the root bridge in order to maximize the S1 → S2 throughput and change your configuration to make it so. Is there an optimal solution or more than one equally good solution? Repeat your S1 → S2 test and compare results with the first time. (Hint: The bridge with lowest bridge ID is elected the root. BIDs are 64-bit numbers by concatenating the bridge priority with the bridge MAC address. Although you normally cannot change the MAC address, you can change the bridge priority.) What is the slowest link traversed in the new network configuration? Was your throughput significantly less than your slowest link speed? Why? (Hint: read up on CSMA/CD)

INITIAL ROUTER CONFIGURATION:

COMMON:

```
service timestamps debug uptime
enable password cisco
no ip domain-lookup
ip classless
line con 0
 exec-timeout 0 0
line vty 0 4
 password cisco
 login
```

R1:

```
hostname r1
interface Serial1/2
 description Link to R2 S1/1
 no ip address
 bandwidth 56
 bridge-group 1
 no shutdown
interface Serial1/3
 description Link to R3 S1/1
 no ip address
 bandwidth 56
 bridge-group 1
 no shutdown
interface Ethernet2/0
 description Link to S1
 ip address 192.168.10.1
255.255.255.0
 bridge-group 1
 no shutdown
interface Ethernet2/1
 description Link to R5 E1
 no ip address
 bridge-group 1
 no shutdown
bridge crb
bridge 1 protocol ieee
 bridge 1 route ip
```

R2:

```
hostname r2
interface Serial1/1
 description Link to R1 S1/2
 no ip address
 bandwidth 56
 clockrate 56000
 bridge-group 1
 no shutdown
interface Serial1/3
 description Link to R3 S1/2
 no ip address
 bandwidth 56
 clockrate 56000
 bridge-group 1
 no shutdown
interface Serial1/4
 description Link to R4 S1/2
 no ip address
 bandwidth 56
 bridge-group 1
 no shutdown
bridge crb
bridge 1 protocol ieee
bridge 1 priority 100
```

R3:

```
hostname r3
interface Serial1/1
 description Link to R1 S1/3
 no ip address
 bandwidth 56
 clockrate 56000
 bridge-group 1
 no shutdown
interface Serial1/2
 description Link to R2 S1/3
 no ip address
 bandwidth 56
 clockrate 56000
 bridge-group 1
 no shutdown
interface Serial1/4
 description Link to R4 S1/3
 no ip address
 bandwidth 56
 bridge-group 1
 no shutdown
bridge crb
bridge 1 protocol ieee
```

R4:

```
hostname r4
interface Fddi0/0
 description Link to R5 FDDI0
 no ip address
 bridge-group 1
 no shutdown
interface Serial1/2
 description LINK to R2 S1/0
 no ip address
 bandwidth 56
 clockrate 56000
 bridge-group 1
 no shutdown
interface Serial1/3
 description LINK to R3 S1/0
 no ip address
 bandwidth 56
 clockrate 56000
 bridge-group 1
 no shutdown
bridge crb
bridge 1 protocol ieee
 bridge 1 route ip
```

R5:

```
hostname r5
interface Ethernet0
 description Link to S2
 no ip address
 bridge-group 1
 no shutdown
interface Ethernet1
 description Link to R1 E2/0
 no ip address
 media-type 10BaseT
 bridge-group 1
 no shutdown
interface Fddi0
 no ip address
 bridge-group 1
 no shutdown
bridge crb
bridge 1 protocol ieee
 bridge 1 route ip
```