

Data-Link Bridge Spanning Tree Algorithm

Assumptions:

1. Each bridge has a globally unique identifier. Each port on a bridge has an identifier unique for the bridge. Identifiers can be sorted in a strict order.
2. There is a special MAC address for all bridges on a LAN.

All bridge ports will be one of the following types:

root port - The port that is the first hop on the minimum cost path to the root bridge. If the cost is equal for multiple ports, the port with the lowest ID is selected. Exactly one root port is defined for each non-root bridge. The root bridge does not have a root port.

designated port - The next hop port on the minimum cost path from the bridge away from the root. Frames are only forwarded out designated ports. There is only **one** designated port per LAN.

blocked port - A port not used for forwarding frames.

The basic idea is to build a spanning tree of bridges and LANs, eliminating the problem of network cycles. To do this, we determine a root bridge and then construct a shortest-path tree to all LANs.

Root Bridge/Port Selection

We arbitrarily designate the bridge with the lowest ID to be the **root bridge**. To determine which bridge has the lowest ID, initially all bridges believe that they are the root. A root bridge periodically asserts its "rootness" to the world by broadcasting a BPDU (Bridge Protocol Data Unit) to the "all bridges" address on every port. A BPDU contains the following information:

- BPDU Transmitting Bridge ID
- Bridge ID of root for transmitting bridge
- Cost of reaching the root via the transmitting bridge (e.g. # LANs to cross getting to root). A root bridge announces the cost as 0. Before forwarding the BPDU, each bridge increments this cost by the cost of crossing the LAN to get to the bridge. (ex, If all LANs cost 1 to traverse, then each bridge increments this cost by 1.)

Initially, every bridge believes it is the root and sends a BPDU. A bridge drops its claim to being root when it receives a BPDU with a root bridge ID lower than its own. A **non-root** bridge designates its root port as the one on which it is receiving BPDU's with the lowest cost to root, breaking a tie in favor of the lowest port ID. A **non-root** bridge forwards the BPDUs received on its root port out its **designated** ports (defined later) after incrementing the BPDU cost. A bridge will forward the BPDU that caused it to become a non-root bridge on all ports (except the incoming port).

Designated/Blocked Port Selection

A bridge believing itself to be the root, considers all of its ports to be designated. A non-root bridge classifies a port as designated if it is 1) not receiving any BPDUs on the port, 2) receiving BPDUs on the port with a root ID that is larger than the ID of its currently believed root, or 3) receiving BPDUs on the port with a root ID equal to its own believed root ID but with a cost higher than the bridge's root cost plus the cost of the LAN connected to the port. (In the case of equal costs, the port is classified as designated if all BPDUs have a *higher* transmitting bridge ID than the receiving bridge's ID.) If a port is not classified as either designated or a root port, it is blocked. Frames are not forwarded on blocked ports.

A bridge receives BPDUs on its blocked ports. If it does not hear a BPDU before a timeout expiration, it assumes a failure has occurred and changes the port classification to designated.