Assumptions: All connections use same connection event time (time slot). High bandwidth connections use common base connection interval. Low bandwidth devices use connection interval of 4x base interval. Network connections are generally static, although some may be bonded but transient. Once network is commissioned, addition of new connections will be rare.

Ideal scheduling example, 4 high bandwidth/low latency connections, 6 lower bandwidth/higher latency connections (4x base interval)

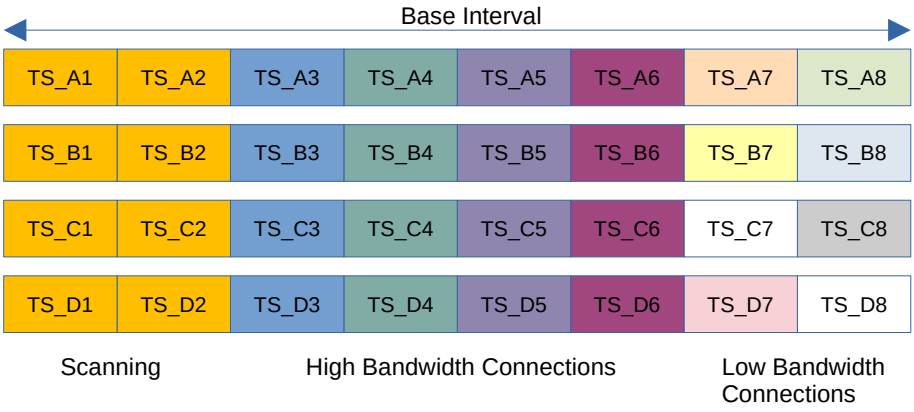

Connection Scheduling Process

Step 1: Scanning for Connection Requests to detect all pending connections

Step 2: Open low bandwidth slots and start accepting connections only for low bandwidth devices

Step 3: After all low bandwidth connections set up, open up 1 high bandwidth slot and accept the first high bandwidth connection. Non transient connections should be initiated before transient connections.

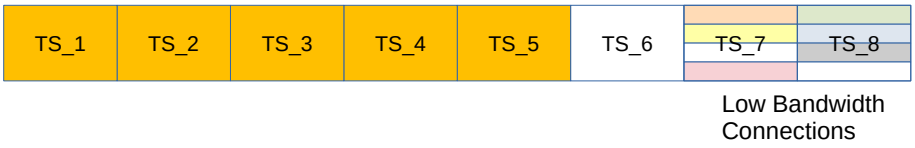

Step 4: After first high bandwidth connection is established, open next time slot and start accepting next high bandwidth connection. Repeat for remaining high bandwidth connections.

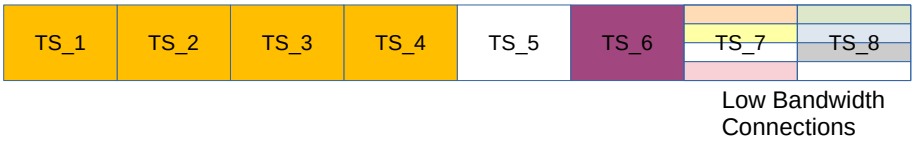

Step 5: If there are any transient high bandwidth connections that are bonded but not available for connection, leave its slot allocated to scanning until the device makes a connection request.

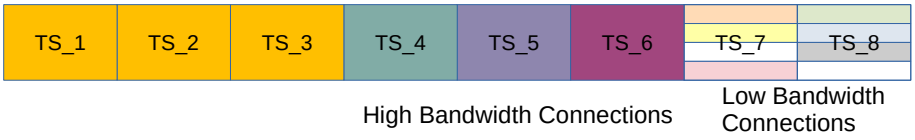

If one of the high bandwidth connections gets dropped, wait for re-connection request and accept. Assuming schedule picks first available slot, it should connect in slot TS_5 again rather than any empty low bandwidth slots.

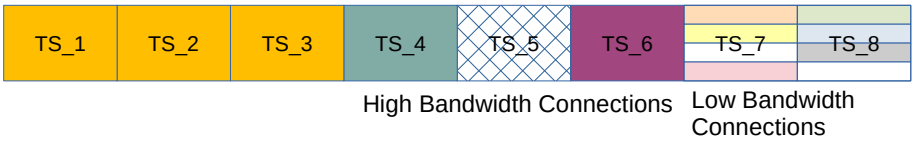

If one of the low bandwidth connections gets dropped, wait for re-connection request and accept. All high bandwidth slots are filled, so ideally it should be scheduled to an empty low bandwidth slot.

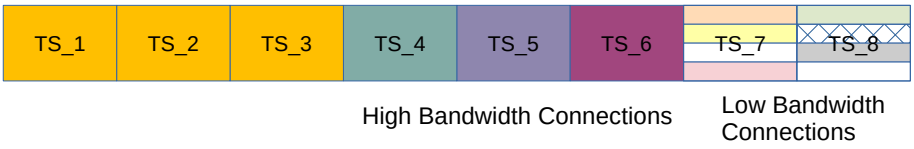

If multiple connections are dropped, should always reconnect the high bandwidth connections first, then the low bandwidth connections. If a significant number of connections drop, it could be that it is best to drop all connections and restart the connection process.

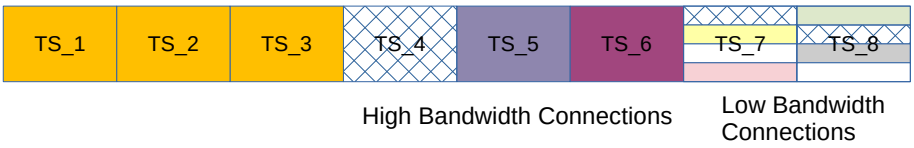

One question is what happens if a low bandwidth connection request is received and connection initiation started during base interval with no empty slots but there are available slots in subsequent base intervals. Does the scheduler only look at the current base interval, or because the low bandwidth connection interval is 4x base, will it look further out and find the next available slot (TS_B7 in this case)?