Updated Proposal for RPR Rate Synchronization

Sept 2002, New Orleans, LA

Rhett Brikovskis, Lantern Communications
David James, JG&G
John Lemon, Luminous Networks
Mike Takefman, Cisco Systems

Purpose

- Review the rate synchronization problems that have been presented in mt_over_01.pdf (May 02) and rb_sync_01.pdf (July 02).
- Review the rate synchronization function that was presented in rb_sync_01.pdf (July 02).
- Propose that the rate synchronization function be adopted for P802.17:
  - Implementation of the algorithm and transmit functions to be optional
  - Receive behavior must be implemented to preserve interoperability

Review of Rate Synchronization Issues

- The draft standard allows RPR implementations with either synchronous or asynchronous PHYs.
- In a synchronous ring, the transmit data rate at each station is exactly identical to the received data rate. A station cannot receive more traffic than it is able to transmit.
- In an asynchronous ring, the transmit data rate at each station is determined by a local clock source, and varies slightly from the nominal link rate. If a station transmits at a lower data rate than the preceding station, the transit queue may fill.
- This rate synchronization problem has been presented in mt_over_01.pdf (May 02) and rb_sync_01.pdf (July 02).

An Example

- A simplified example is shown below.
  - Station A is transmitting at the nominal link rate plus 100 PPM.
  - Station B is receiving only transit traffic from Station A, but is transmitting only at the nominal link rate minus 100 PPM.
  - If the link is fully utilized, PTQ in B will fill at (link rate)*200 PPM.
  - If this condition is sustained for a sufficient period of time, the PTQ in B will overflow and lose packets.

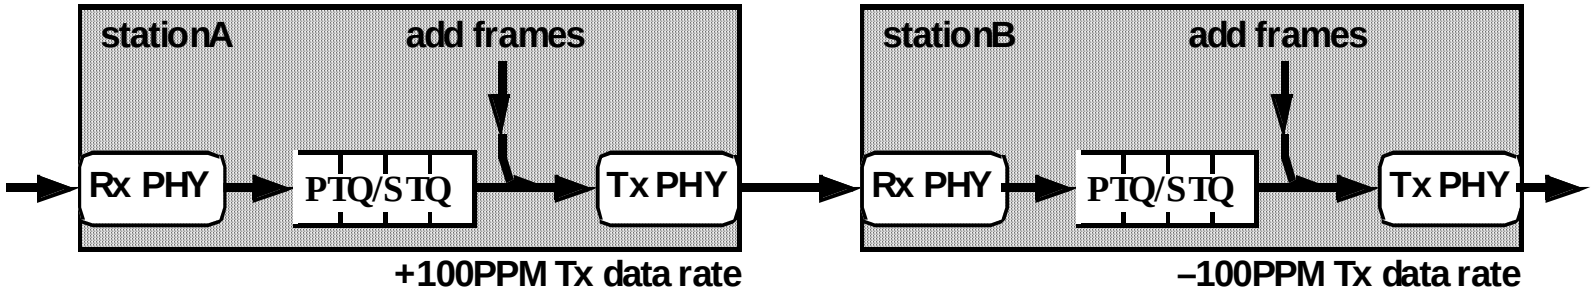

How to solve the rate sync issue?

- Why not use RPR fairness to provision the link for <100%?
  - Fairness operates on a time-averaged basis, and has a non-zero response time. Fairness cannot ensure that link utilization won't burst to 100% for long enough to overflow the PTQ. Link also may be over-provisioned and reach full utilization. Provisioning guidelines may resolve this issue in specific applications.
- Why not force the link utilization to <100% by inserting small idle periods (similar to Ethernet IPG)?
  - When link reaches maximum utilization, idle periods will be inserted to limit data rate to <100%. However, the following station must still re-transmit the same minimum idle periods it receives, and the receive/transmit rate difference is unchanged. The PTQ will fill.
- Need to insert small, variable idle periods in the transmitted datastream, and vary their size based on the TX/RX rate difference.

A Simple Rate Synchronization Function

- Insert small idle periods into the transmitted datastream of each RPR station, and vary the time interval between the idle periods in accordance with the depth of the PTQ:
  - Any time the PTQ level is below a first threshold (PTQ occupancy is low), insert fixed idle periods at a regular interval corresponding to 500 PPM of the nominal link data rate.
  - Above the first threshold, inhibit the addition of any local traffic (add traffic or control packets). This ensures that the PTQ can only continue to fill if the RX rate exceeds the TX rate and link utilization is high.
  - Above a second PTQ threshold, increase the interval between idle periods to a rate corresponding to 250 PPM of the nominal link data rate. This increases the TX data rate to reduce the PTQ level.

A Graphical View

- Operation of the algorithm vs PTQ depth

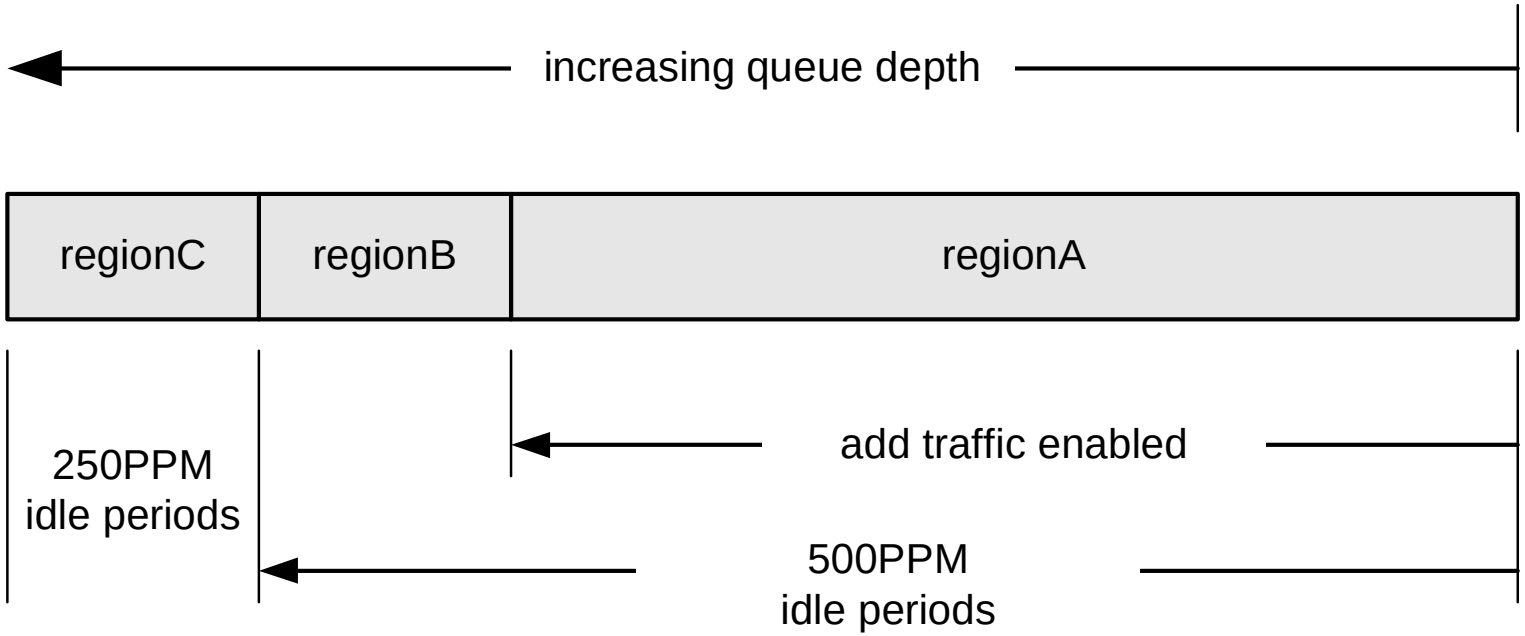

Idle Packets

- To generate idle periods, the MAC could transmit nothing, but these idle periods may not be preserved in all implementations.
- Use an “Idle Packet” instead:
  - Define an Idle Packet (a reserved “Type 0” Fairness Message variant) compatible with all PHYs.
  - Idle Packet is a point-to-point message only, and is discarded by recognition or by TTL.
  - Idle Packet is a fixed 16-byte length with a fixed payload value.

Implementation requirements

- Generation of the rate synchronization Idle Packets (implementing the algorithm) is optional
  - MACs intended for use only with synchronous PHYs do not need this function, other implementations may also choose to omit this function
- Reception of Idle Packets is mandatory for all implementations to ensure interoperability
  - Must be able to recognize and discard an Idle Packet in any implementation

Summary of the rate synchronization proposal

- Adopt a rate synchronization function based on this proposal for P802.17
  - Preliminary text proposal is included in DvjComments.pdf (Sept 02) for consideration.
- Eliminate Annex B, and move the rate synchronization function to Clause 6, since it is part of the MAC datapath.
- Implementation is optional except that a MAC must recognize and discard Idle Packets.