EFM Copper

Exploring MAC-PHY rate matching

EFM November, 2001

Hugh Barrass (Cisco Systems)

MAC-PHY rate matching problem

- **It is safe to assume that the EFM copper data rates will not match exactly existing 10Mbps or 100Mbps Ethernet**
  - **Odd numbers, asymmetry, non-deterministic…**
- **Safe to assume that line rate will be <100Mbps**
  - **What about loop aggregation at short reach?**
  - **MAC-PHY interface based on MII**
- **Need to determine goals / objectives for interface**
  - **… then examine solutions**
- **Note: this presentation does not propose a solution**
  - **The exploration should invite proposals**

Goals and Objectives

- **Optimal use of PHY buffering**
  - Whole frame or multi-frame buffer will limit PHY integration
- **Re-use of existing MAC silicon**
  - Including highly integrated multi-MAC, bridge devices
- **Not mess with other protocols**
  - Including 802.3x?
- **Be mindful of RMII, SMII**
  - Although out of scope…

PHY → MAC

- **This direction should be easy!**
  - Assume MAC is always capable of 100Mbps (& full duplex)
  - Support for 10Mbps or half duplex MAC unnecessary?
- **Open loop solution OK**
  - No need for MAC to hold off PHY
  - 802.3x allows pause of other station (nb latency!)
- **Simple solution – stretch IPG**
  - PHY assembles whole frame & sends when ready
  - Needs 1 whole frame buffer
  - Adds more jitter (1 max frame time)

Can we eliminate frame buffer?

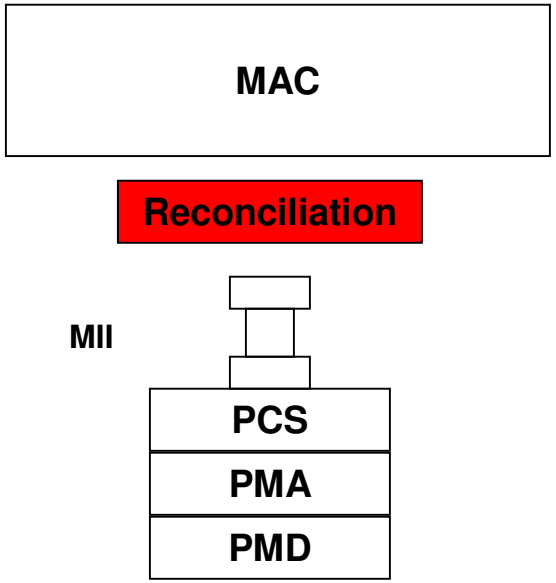

- **Remember the reconcillation sublayer?**
  - **Allows some adjustment of MII interface**
  - **Could be used for finer grain rate matching (to reduce buffer)**

Two more methods

- **Adjust the Rx clock**
  - Stretch the clock cycle to match the line rate
  - May cause problems with some installed silicon (also may not work with RMII / SMII)
- **Use RxDv to "skip" cycles**
  - MAC (via reconciliation) ignores clock cycles with RxDv de-asserted
  - Almost certainly won't work with currect MAC silicon – does work with RMII / SMII

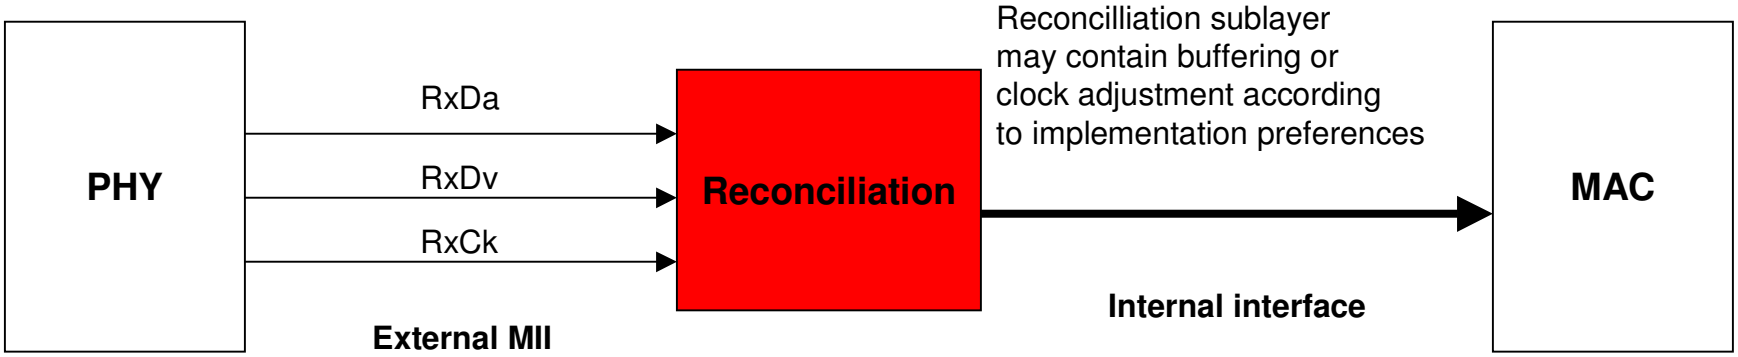

MAC → PHY

- **Open loop only works with fixed line rate**
  - IPG stretch (as per 10G), MAC must know real line rate
  - Does not work with HDLC
  - Would need pre-emptive PHY-MAC notification for rate adaption (or automatic loop aggregation)
- **Closed loop solution**
  - PHY must tell MAC when buffer is filling
  - Varying levels of granularity → various PHY buffer sizes

Closed loop – pause frame

- **Send 802.3x pause to stop transmission**
  - Frame insertion in PHY not advisable
  - Control of outbound data effects inbound service
  - Multi-frame resolution – requires large buffer
  - Causes very large jitter (max frame size)
- **Could interfere with station to station .3x**
  - Difficult to predict the interactions

Closed loop – other methods

- **Fake collision to apply back-pressure**
  - Well known technique
  - Single frame resolution
  - Forces MAC to be in half-duplex
  - Confuses SNMP stats
- **TxClk stretch or skip**
  - Same issues as for Rx
  - Maybe re-interpret CRS for skip (instead of new pin)
- **Reconciliation sublayer could rescue us**
  - Hide the rate matching implementation from the MAC

Reconciliation sublayer

- **Either skip or collision**
  - **MAC doesn't see the events – no SNMP problem**
  - **Buffering or clock adjustment in RS - implementation specific**
  - **Could work with RMII / SMII**

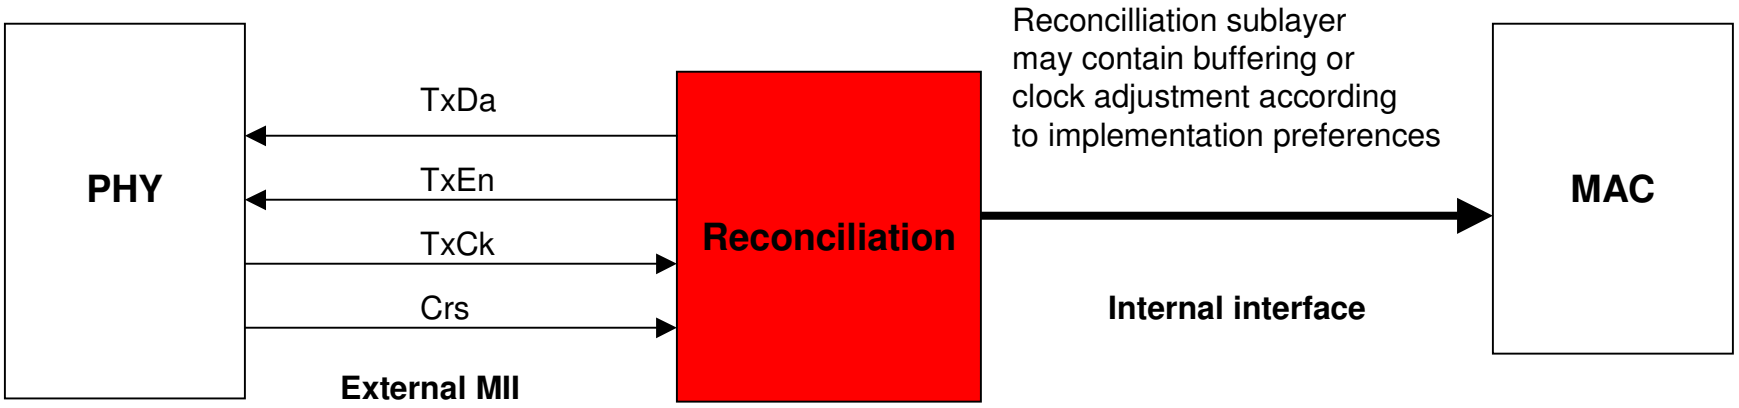

In conclusion

- **Multiple possible solutions**
  - Need to define priorities / objectives
- **PHY buffering vs MAC silicon changes**
  - Chip vendors vs systems vendors?
- **Proposals invited**
  - … consensus?