

Comparison of Rate Control Methods

IEEE P802.3ae
10 Gigabit Ethernet Task Force
24-May-2000
Ottawa, ON
Howard Frazier - Cisco Systems

Goal For This Presentation

- Present options for rate control, in support of our objective to:
 - Define a mechanism to adapt the MAC/PLS data rate to the data rate of the WAN PHY
- Compare the options
- Stimulate discussion

Why Do We Need Rate Control


- The MAC/PLS service interface operates at 10.000 Gb/s
- The WAN PHY is constrained by the available payload rate of OC-192c/SDH VC-4-64c
- For the UniPHY the available payload rate is 9.29419 Gb/s

Candidate Rate Control Methods

- Clock stretching
- Word hold
- Busy idle
- Open loop
- Frame based

Clock Stretching


- Stretch or “gap” the transmit and receive clock references at the XGMII



Word Hold

- Add signals to the XGMII to hold off the MAC transmitter and receiver

– http://www.ieee802.org/3/10G_study/public/jan00/figueira_1_0100.pdf



Busy Idle

- PHY sends normal Idle to MAC when it can accept data
- PHY sends Busy Idle to MAC when it wants MAC to slow down
 - MAC inhibits further transmission at packet boundary

– http://www.ieee802.org/3/10G_study/public/jan00/frazier_1_0100.pdf

Open Loop

- The MAC stretches its IPG in proportion to the length of the packet proceeding the IPG
- The MAC knows the desired rate, and stretches its IPG appropriately

– http://www.ieee802.org/3/10G_study/public/mar00/muller_1_0300.pdf

Frame Based

- Similar to 802.3x Pause Flow Control
- PHY generates Pause frames with a timer value of either FF (XOFF) or 00 (XON)

Clock Stretching

- Pros
 - Precise
 - No extra pins
 - Minimal FIFO required in PHY ~ 16 bytes
 - PHY controls the rate
 - Adapts to a wide range of rates and varying rates
- Cons
 - Difficult timing
 - Interrupts flow of data through pipeline stages
 - Makes buffer pre-fetching difficult
 - Doesn't work with XAUI

Word Hold

- Pros
 - Precise
 - Minimal FIFO required in PHY ~ 16 bytes
 - PHY controls the rate
 - Adapts to a wide range of rates and varying rates
- Cons
 - Extra pins
 - Difficult timing
 - Interrupts flow of data through pipeline stages
 - Makes buffer pre-fetching difficult
 - Doesn't work with XAUI

Busy Idle

- Pros
 - Precise
 - No extra pins
 - PHY controls the rate
 - Works with XGMII and XAUI
 - Adapts to a wide range of rates and varying rates
- Cons
 - Small FIFO required in PHY ~700 bytes
 - Constrains physical distance between MAC and PHY

Open Loop

- Pros
 - No extra pins
 - Works with XGMII and XAUI
 - MAC controls the rate
 - Not sensitive to physical distance between MAC and PHY
- Cons
 - Imprecise
 - Small FIFO required in PHY ~600 bytes
 - Does not adapt to varying rates
 - MAC must know the rate a priori

Frame Based

- Pros
 - No extra pins
 - Familiar mechanism
 - Already implemented in most Gig MACs
 - PHY controls the rate
 - Works with XGMII and XAUI
- Cons
 - Confusion between PHY to MAC rate control and end to end flow control frames
 - Small FIFO required in PHY ~800 bytes
 - Relatively large amount of logic required in PHY

Summary

- Several candidate rate control methods exist
- All of them can meet the objective
- The choice comes down to how much weight you give each of the pros and cons