SOLARFLARE Communications

10GBASE-T EEE Synchronization

11/11/2008 Gavin Parnaby

Energy Efficient Ethernet

Supporters

- George Zimmerman, Solarflare
- Brett McClellan, Solarflare
- Dimitry Taich, Teranetics
- Mike Grimwood, Broadcom

Synchronization objectives

- Maximize the quiet time for best power savings
- Maximize predictability
  - **architecture optimization**
  - **eliminate interoperability risk**
  - **simplify validation**
  - **simplify testing**
- Ideally require no new signaling
- Stable; no link degradation due to corner cases of synchronization mechanism

Proposed solution

- Master and slave establish synchronization by using the transition to PAM-16 as a common reference
  - **PAM-2 infofield during initial training contains a counter that indicates the number of frames to the PAM-16 transition (state PCS_Test, clause 55.4.2.5), for master and for slave**
  - **Master sends infofields to slave indicating exactly when the transition to PAM-16 will happen on the master to slave side**
  - **Slave sends infofields to master indicating exactly when the transition to PAM-16 will happen on the slave to master side**

Proposed solution

- Generate the refresh timing for both sides from fixed timing references
  - **Start tx LDPC frame counter when local phy transitions from training to data mode during initial training**
  - **Start rx LDPC frame counter when LP transitions from training to data mode during initial training**
- Active pair on both sides derived using simple modulo arithmetic from the appropriate counter
- Refresh_active for both sides generated by simple modulo arithmetic from the appropriate counter
- Both PHYs know exactly when they send and receive refresh, and the active pair
  - **Following the PAM-16 transition event both sides know when refreshes are due, and which pair is active, with no ambiguity**

Precise synchronization

- Existing text requires only that master and slave countdowns are synchronized to within 1 PMA training frame [= 64 LDPC frames] [clause 55.4.5.14]
  - **This is not precise enough to ensure refreshes do not overlap**
- New requirement that for EEE compliant PHYs, the slave synchronize its transition to PAM-16 to within 1 LDPC frame of the master, at the MDI
  - **Ensures that refresh offset is < 1 LDPC frame from ideal at the slave MDI**
  - **Ensures that refresh offset is < 3 LDPC frames from ideal at the master MDI**
- Ample opportunity to do this during the countdown to PCS_Test

Definitions

- Tc = Tq + Tr
  - **The time for a complete quiet refresh cycle on a single pair**
- To
  - **The desired offset time between the master and slave refresh signals**
- U
  - **Master transmitted LDPC frame count, reset to 0 at master's Tx transition to PAM-16 at PCS_Test**
- V
  - **Slave transmitted LDPC frame count , reset to 0 at Slave's Tx transition to PAM-16 at PCS_Test**

Assumptions

- Use Tc=128, Tq+Tr=128

Tr	Tq
4	124
8	120
16	112
32	96

Generating logic

	Master Tx / Slave Rx	Master Rx / Slave Tx
Refresh active	Tq ≤ Mod (U,Tc)	To ≤ Mod (V,Tc) < To+Tr
Pair A	0 ≤ Mod(U,4Tc) < Tc	To+Tr ≤ Mod(V,4Tc) < To +Tr +Tc
Pair B	Tc ≤ Mod(U,4Tc) < 2Tc	To +Tr +Tc ≤ Mod(V,4Tc) < To +Tr + 2Tc
Pair C	2Tc ≤ Mod(U,4Tc) < 3Tc	To +Tr + 2Tc ≤ Mod(V,4Tc) < To +Tr + 3Tc
Pair D	3Tc ≤ Mod(U,4Tc) < 4Tc	To +Tr + 3Tc ≤ Mod(V,4Tc) < 4Tc OR 0 ≤ Mod(V,4Tc) < To

U = master transmit LDPC frame counter

V = master receive LDPC frame counter

Tc = Tq + Tr = 128; To = Tc/2 - Tr

Tx_mode control

- Lpi_tx_mode =
  - **NORMAL when !tx_lpi_active**
  - **REFRESH_A when tx_lpi_active * (tx_active_pair==PAIR_A * tx_refresh_active)**
  - **REFRESH_B when tx_lpi_active * (tx_active_pair==PAIR_B * tx_refresh_active)**
  - **REFRESH_C when tx_lpi_active * (tx_active_pair==PAIR_C * tx_refresh_active)**
  - **REFRESH_D when tx_lpi_active * (tx_active_pair==PAIR_D * tx_refresh_active)**
  - **QUIET when tx_lpi_active* !tx_refresh_active**
  - **ALERT when tx_lpi_active * tx_alert_active**

Synchronization at slave

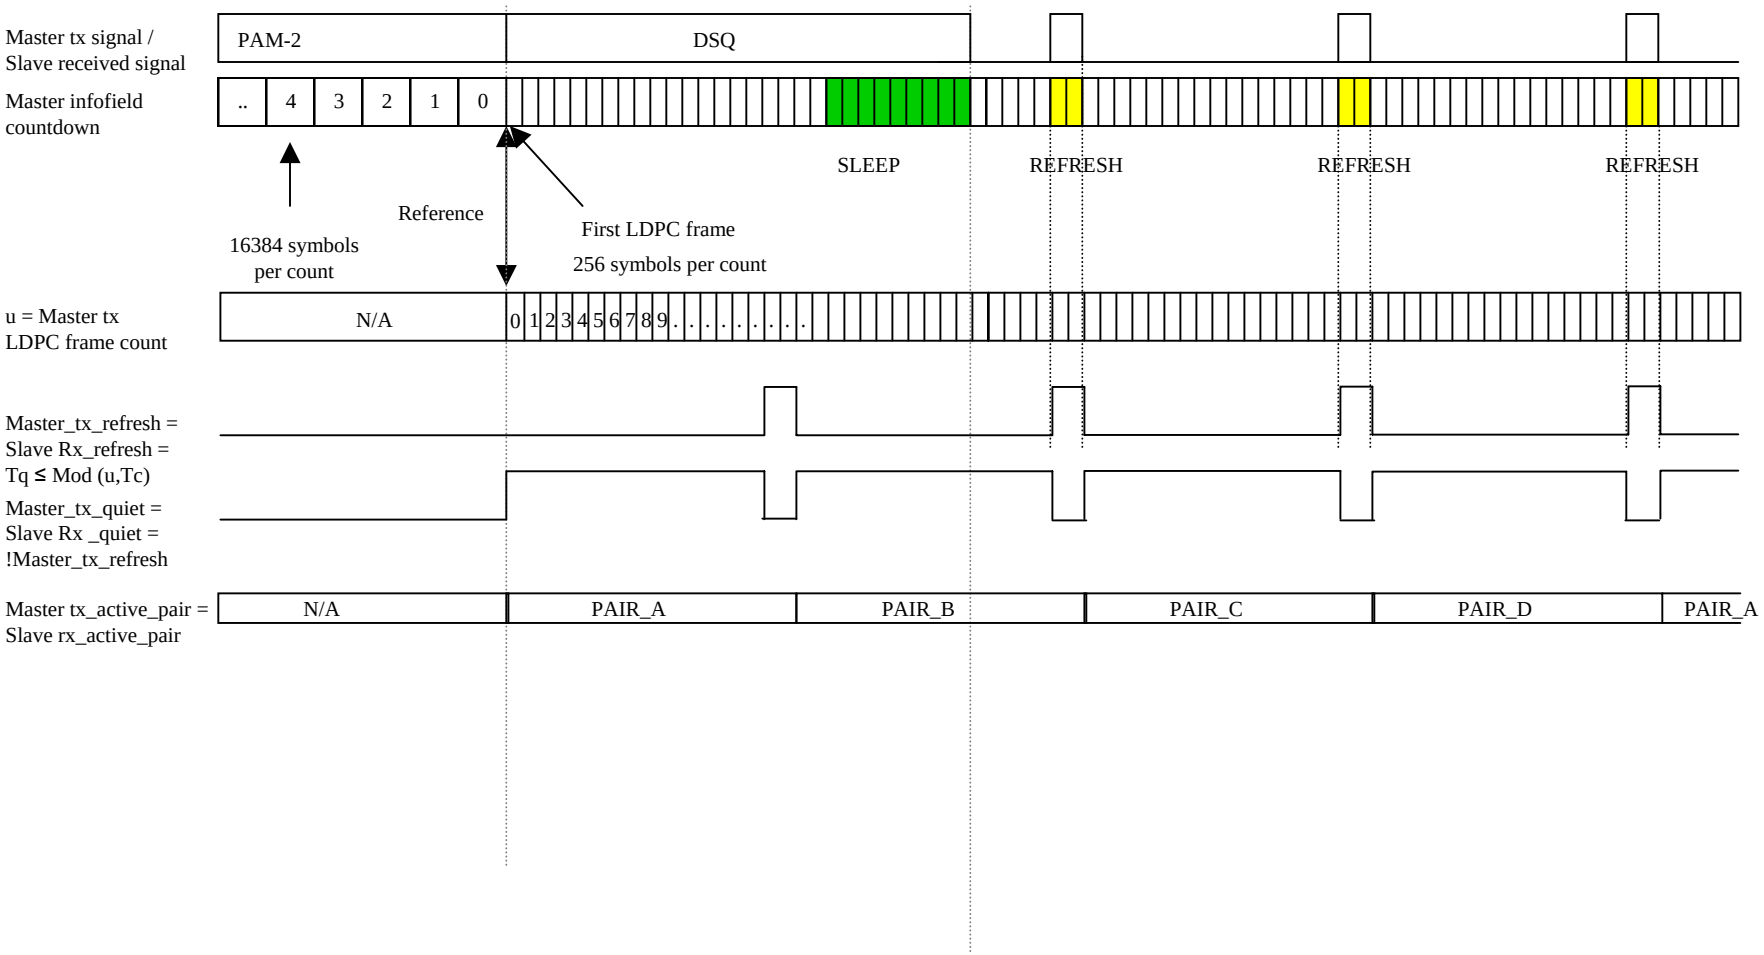

SOLARFLARE Communications

Synchronization at master

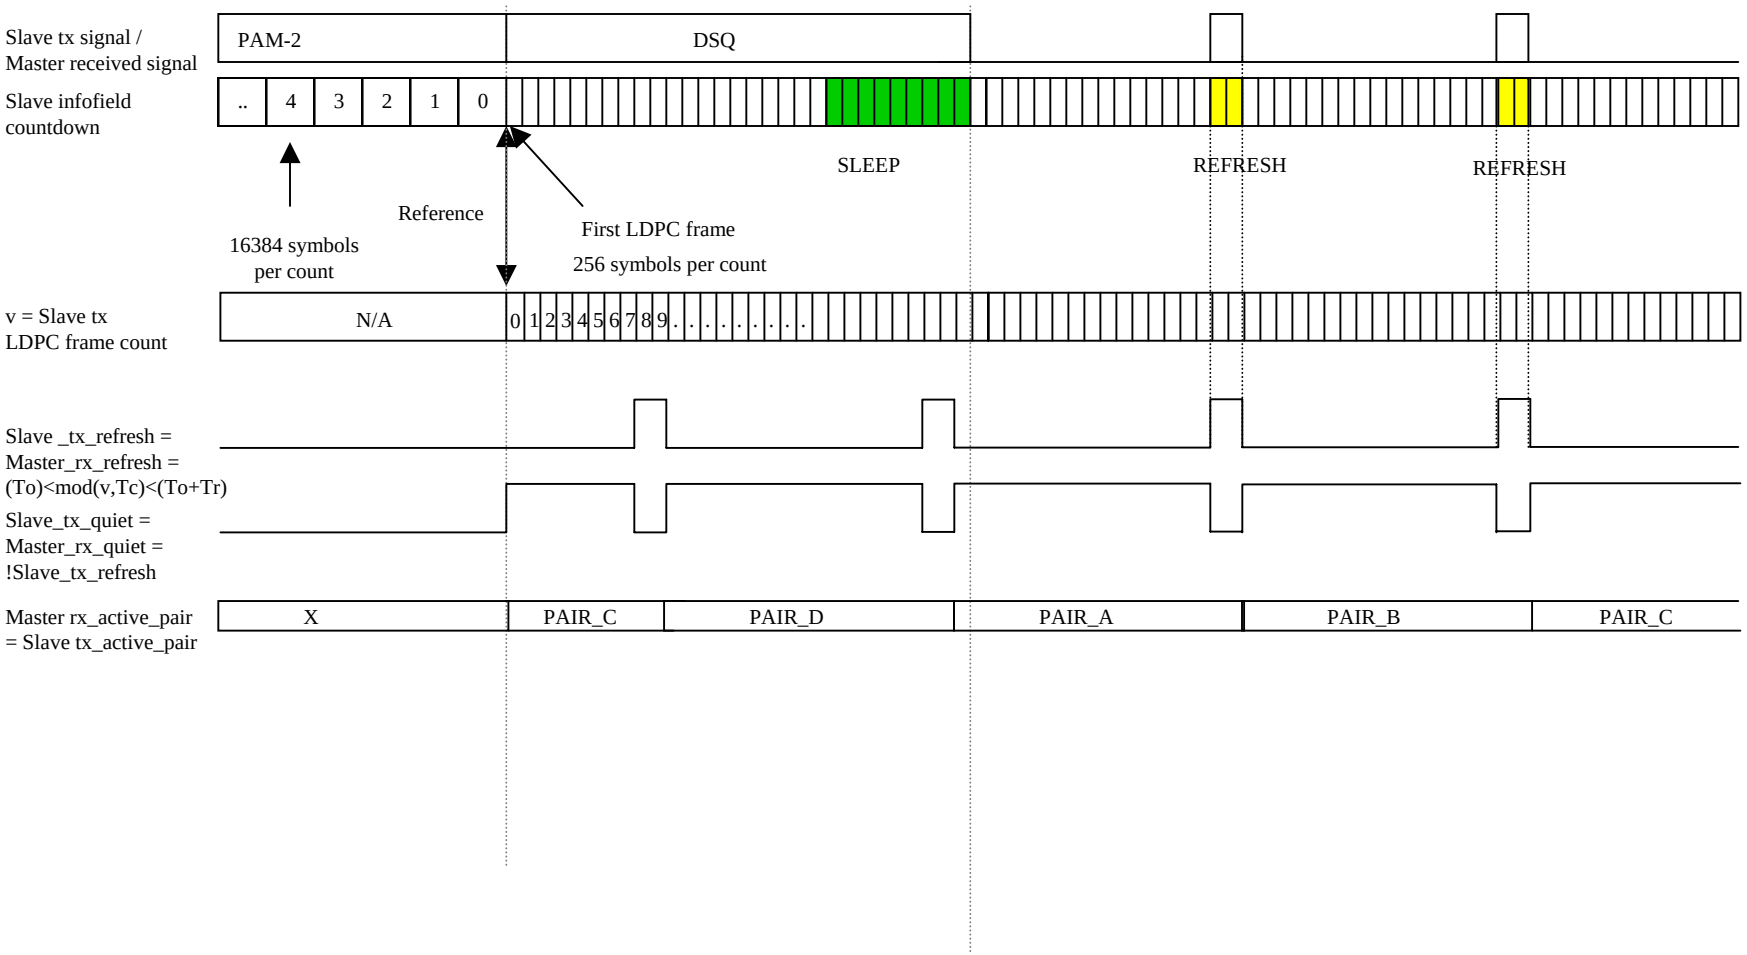

SOLARFLARE Communications

Synchronization between master and slave

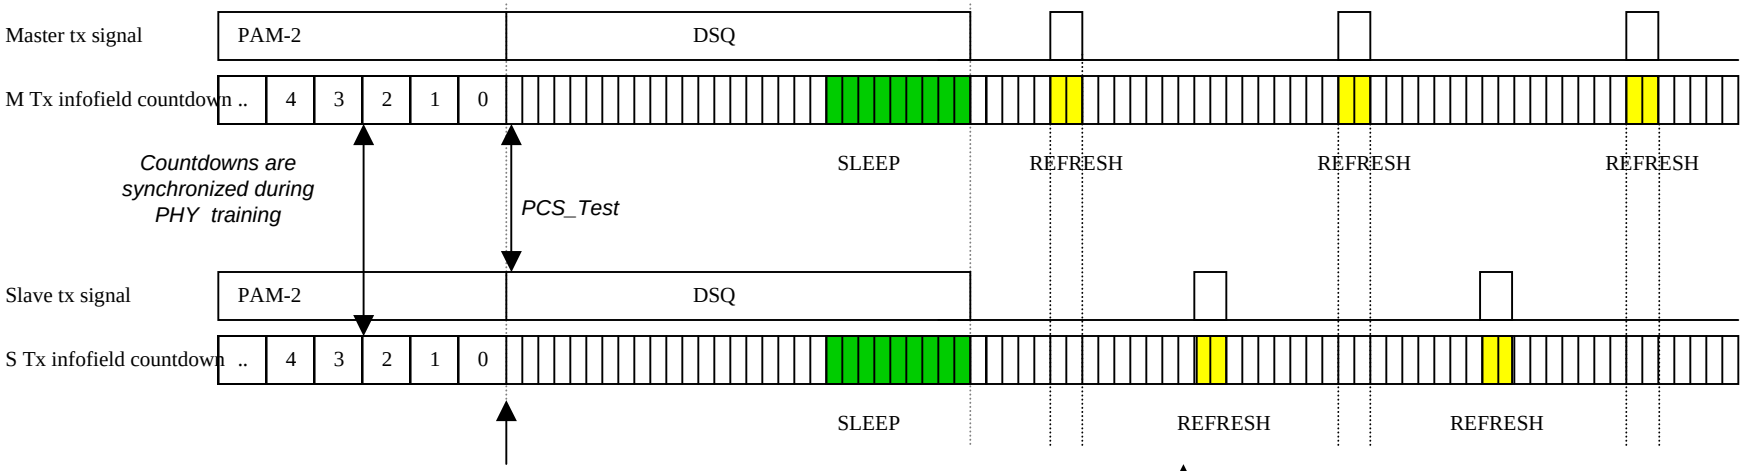

New requirement: Slave is required to synchronize transition to data-mode / PAM-16 with the master's transition
[The 802.3an requirement is within 1ms]
This guarantees that the transmit and receive LDPC frame counters are synchronized to within 1 LDPC frame + d on the master side

Quiet/refresh control generated as previously described

- *Each PHY has a Tx LDPC frame counter, and an Rx LDPC frame counter*
- *The Tx LDPC frame counter is reset when the transmit signal transitions to data mode*
- *The Rx LDPC frame counter is reset when the receive signal transitions to data mode*
- *Transmit and receive refresh periods are deterministic thereafter, on both sides*

Further details

- Refreshes that begin during SLEEP are not transmitted
  - **To eliminate partial refreshes that overlap with SLEEP**
  - **Means that LPI state always begins with quiet or full refresh**

Advantages

- Quiet / refresh is completely predictable on both sides
- Can bound refresh offset to a very small number of frames on both sides
- Easily testable
  - **Only requirement is slave countdown synchronization**
- Potential for higher power savings due to predictable power downs

Conclusion

- Method achieves robust synchronization without introducing new signaling
  - **Key variables are deterministic on both sides throughout**