Hybrid Buses and the Border PHY
A P1394a Approach
(updated 7/19/1999)

Jerry Hauck

Michael Johas Teener

Colin Whitby-Strevens

ZAYANTE

P1394a Arbitration Tutorial A Refresher

ZAYANTE

P1394a PHY Arbitration Request Rules

A request from a child port is automatically and immediately forwarded towards the root. Other children ports receive DATA_PREFIX.

Requests from the link:

- IsoReq: Immediately forwarded
- FairReq or PriReq: forwarded at
  - Known end of subaction (ack)
  - Gap = Subaction + Arb Delay
  - Gap ≥ Arb Reset + Arb Delay

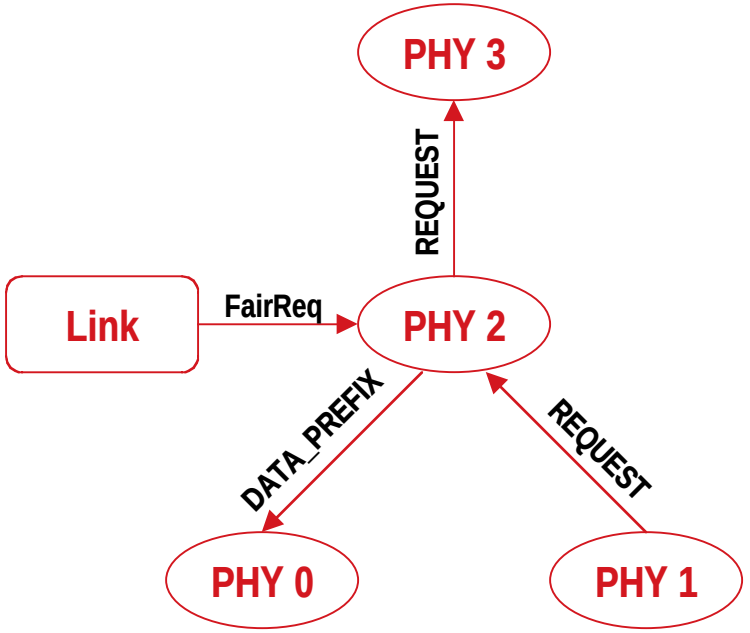

ZAYANTE

P1394a Arbitration Quiet Windows

To ensure consistent, bus-wide detection of gaps, legacy arbitration has quiet intervals during which arbitration must not be initiated:

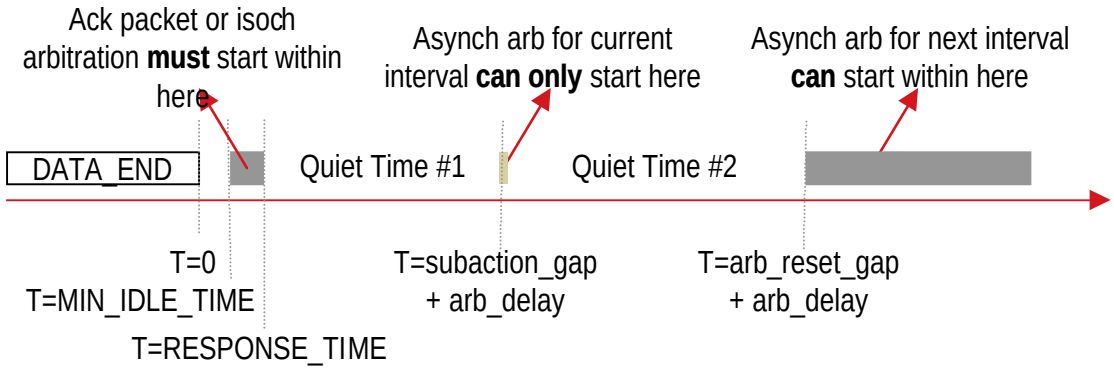

P1394a Grants

Explicit “Loud” Grant

- Grant traverses downward from root along path to requesting node.
- All branches off of “chosen” path are held inactive with leading DATA_PREFIX.

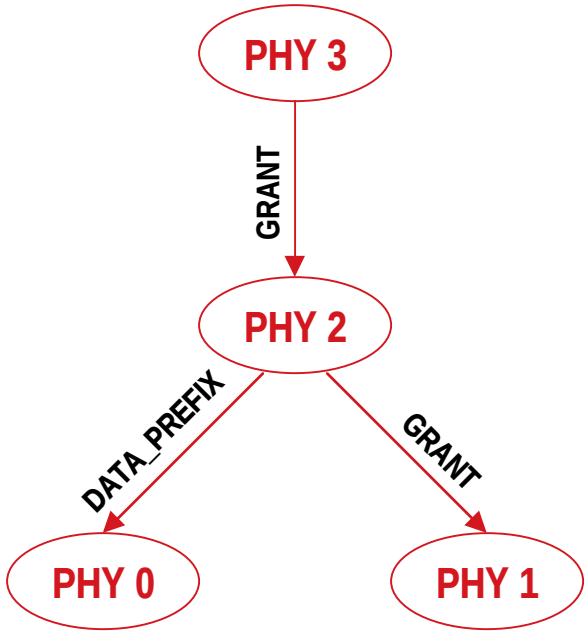

P1394a Grants (cont)

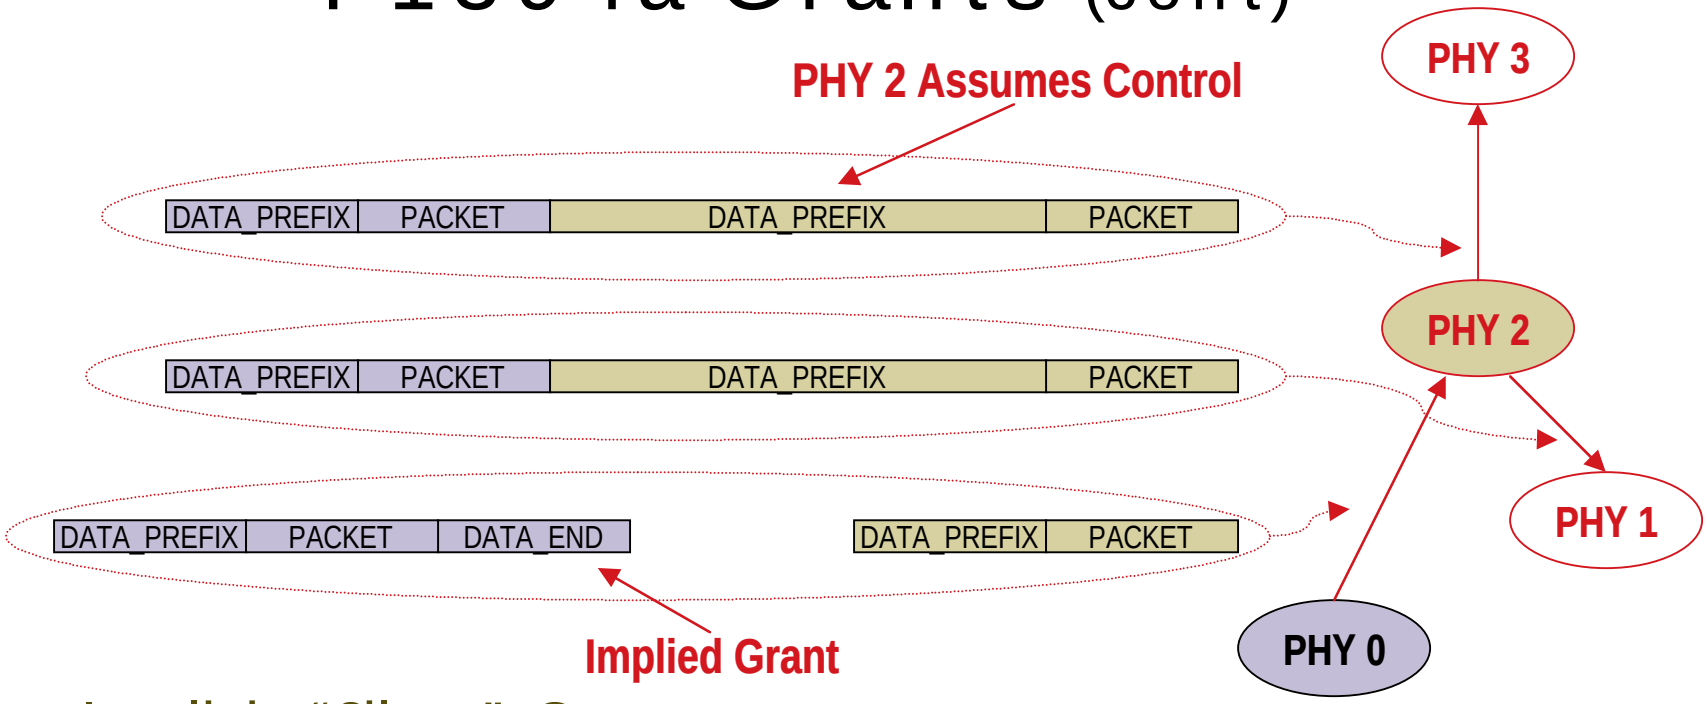

Implicit "Silent" Grant

- After packet transmission, DATA_END traversing upwards along the path to the root is an implied GRANT
- Parents noting DATA_END can perform fly-by concatenation and assume control of the bus.

P1394a Bus Phases

Isochronous Interval

- Begins with a Link-issued Cycle Start Packet
- Concludes when no more isochronous arbitration occurs, causing a subaction gap to elapse on the bus

Fairness Interval (Asynchronous)

- Active whenever not in the isochronous interval
- Fairness interval boundaries marked with arbitration reset gaps which occur when no asynchronous arbitration remains.

BOSS: Reviewed and Slightly Revised

Establishing BOSSship

When does a PHY become BOSS?

- Any beta/border PHY first originating a packet into a given beta cloud is BOSS of that local cloud, immediately.
- A PHY which receives an explicit or implicit GRANT is BOSS.
- The “local root” of a cloud automatically assumes BOSS*ship* after any extended period of inactivity.

When does a PHY surrender BOSS*ship?*

- Whenever a packet is received from a beta-mode port, the receiving PHY ceases to be BOSS.
- When an explicit or implicit GRANT is issued, the issuing PHY ceases to be BOSS.

BOSS Request/Grant Rules

- Current or pipelined requests are issued by P1394b Links
- PHY's constantly forward favorite asynch and favorite isoch requests received from attached link and any cable ports towards the current BOSS.
- At the established end of a subaction, the current BOSS immediately grants any in-phase request.

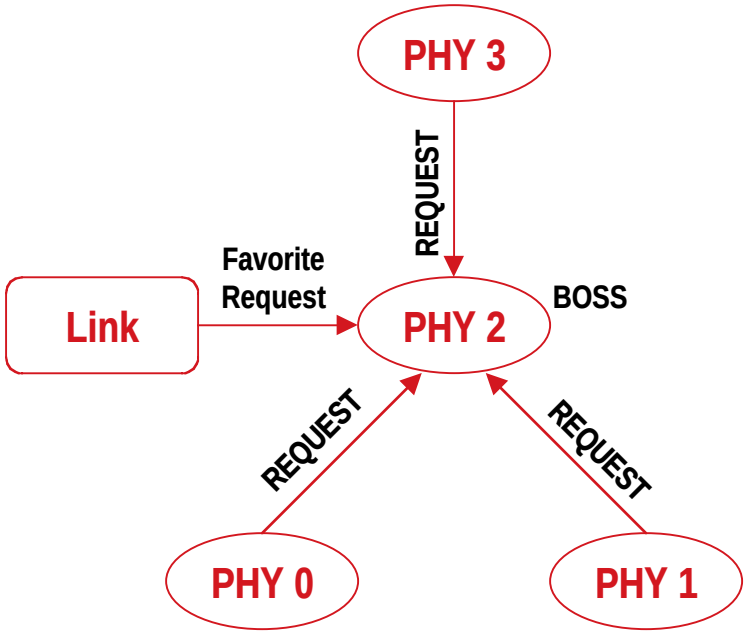

Types of BOSS Grants

Explicit, Subaction Completion Confirmed

- Explicit-Loud
  - Issued after the end of a subaction in response to an active, in-phase asynch, isoch, or legacy request
  - Granted port sees GRANT control symbol

  - Denied ports see DATA_PREFIX control symbols

- Explicit-Quiet
  - Returns control to parent (and ultimately local root) at the end of a subaction when no in-phase asynch, isoch, or legacy requests remain
  - Granted port (parent) sees GRANT control symbol

  - Denied ports (children) see DATA_END

Implicit, Subaction Completion Unconfirmed

- Returns control to parent (and ultimately local root) when subaction has not yet concluded. Local root has responsibility to time-out (i.e., ACK-MISSING)
- All ports see DATA_END, parent port interprets as an implicit grant.

ZAYANTE

Grants Illustrated

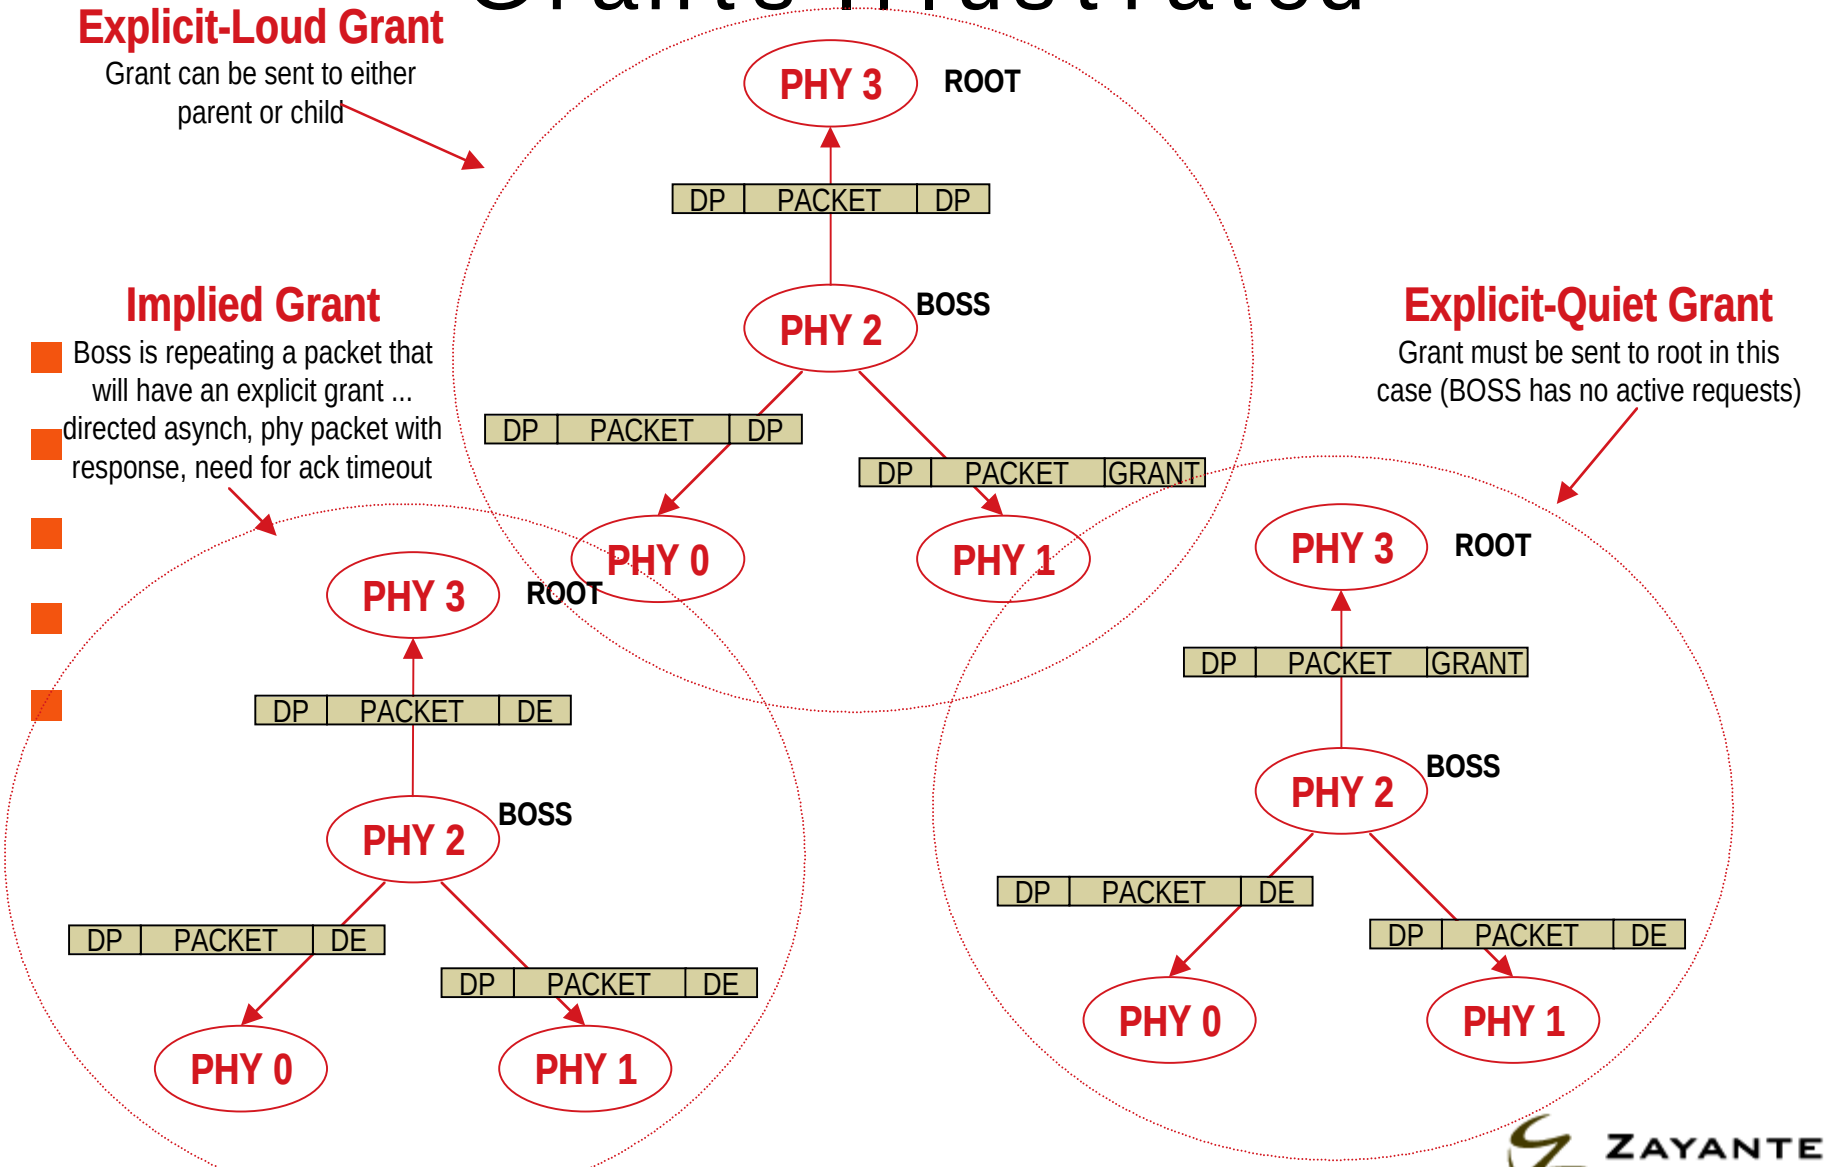

When BOSS is Permitted to Transmit

Explicit Grant

– End of subaction has been explicitly marked, it is safe for BOSS to use received GRANT for any in-phase request.

Implicit Grant

– End of subaction has *not* been communicated. If receiving BOSS can independently determine that the subaction is concluded (e.g., bus is in isoch phase or last packet was an ack in the asynch phase, or request from a legacy Link), GRANT can be used accordingly.

Unarbitrated

– PHY can assume BOSS*ship* immediately and begin transmitting an ACK or a PHY response packet.

P1394b Bus Phases

Isochronous Interval

- Begins with P1394b link issues a Cycle Start packet *and* a CYCLC_START_ODD/EVEN token.
- Concludes when current BOSS has no favorite in-phase isoch request to grant.
- BOSS marks conclusion of interval by issuing ASYNC_START token.

Fairness Interval (Asynchronous)

- Active whenever not in the isochronous interval
- Fairness interval boundaries marked by current BOSS with ARBRST_EVEN/ODD token when no favorite in-phase asynchronous requests remain.

Summary of Proposed Border

Feature	P1934a	Border	Beta-Only
Gap Timer	✓	✓	
Gap Tokens		✓	✓
Quiet Arb Windows	✓	✓	
Retractable Requests	✓	✓	
Pipelined Requests		✓*	✓
Explicit-Loud Grant	✓	✓	✓
Explicit-Quiet Grant		✓*	✓
Implicit Grant to Parent/Root	✓	✓	✓
Accel/Decel Timer	✓	✓*	
DATA_PREFIX Extension		✓	✓

Note: ✓* signifies an optional feature for a simple long-haul repeater, gray rows indicate non-P1394a required features for border nodes

ZAYANTE

Detailed Proposal
Border PHY Theory of Operation

P1394b Scope

- P1394b fundamentally offers an enhanced PHY layer specification for the family of IEEE1394-1995 based protocols.
- Changes to higher layers, services, and programming models are strongly discouraged and avoided (enhanced arbitration services between PHY and Link are exception)
- As a result, certain PHY Layer indications are expected/required for proper LINK and TRANS layer operation whether the underlying bus is composed of P1394a PHYs, P1394b PHYs, or both ...

ZAYANTE

Required PHY Event & Data Indications

Regardless of PHY flavors present in a connected bus, certain PH_EVENT.indication's and PH_DATA.indication's are required. Of interest to the hybrid topology discussions:

- BUS_RESET_COMPLETE, marking the end of the self-ID phase of bus configuration
- SUBACTION_GAP marking the end of an isochronous period or indicating a missing ACK packet
- ARBITRATION_RESET_GAP marking the boundaries of fairness intervals

For proper operation of higher layers and fulfillment of the programming model (isoch period batched before asynch), these indications must be faithfully synchronized across all nodes on a bus.

ZAYANTE

Communicating PHY Indications BOSS vs Legacy

PHY Indication	**BOSS** (Token)	**Legacy** (Gap)
End of Self-ID	ASYNC_START(?)	subaction gap
End of Isoch Interval	ASYNC_START	subaction gap
Missing ACK (only needed at subaction initiator)	ASYNC_START(?)	subaction gap
Boundary of Fairness Interval	ARBRST_ODD/ ARBRST_EVEN	arb reset gap

BOSS and Legacy arbitration algorithms have different ways of communicating/synchronizing PHY Indications across a bus and ultimately to attached nodes ...

ZAYANTE

Hybrid Bus Problem Statement #1

Synchronizing PHY Indications

PHY Indication	BOSS (Token)	Legacy (Gap)
End of Self-ID	ASYNC_START(?	1:1 Ratio subaction gap
End of Isoch Interval	ASYNC_START	1:1 Ratio subaction gap
Missing ACK (only needed at subaction initiator)	ASYNC_START(?	1:1 Ratio subaction gap
Boundary of Fairness Interval	ARBRST_ODD/ ARBRST_EVEN	1:1 Ratio arb reset gap

... Consequently, a hybrid bus requires both BOSS and Legacy style indications to “fire” for each occurrence of a given PHY indication.

ZAYANTE

Hybrid Bus Problem Statement #2

Beta-Only Arbitration Constraints

- Beta-only PHYs, as an assumption, do not have configurable gap timers. If any timeout period is implemented, it is assumed to be fixed and set for a maximum sized topology. As a consequence, attempting to set the gap_count on a beta-only PHY has no effect.
- BOSS currently defines asynchronous and isochronous pipelined requests. The BOSS PHY needs to understand the bus phase (isoch or asynch, even or odd) before it can grant an incoming request.

ZAYANTE

Hybrid Bus Problem Statement #3

Legacy Arbitration Constraints

To ensure consistent, bus-wide detection of gaps, legacy arbitration has quiet intervals during which arbitration must not be initiated:

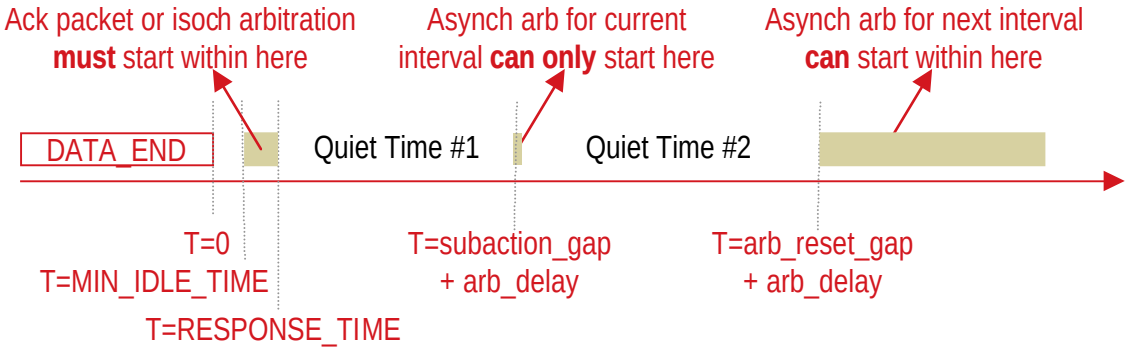

Null packets (consisting of DATA_PREFIX & DATA_END & IDLE) have a minimum specified duration of ~440 ns (140 + 260 + 40)

Proposed Solution - Overview

Synchronizing Legacy Gap Periods with BOSS Tokens

Summary:To synchronize Legacy gap indications with BOSS token indications:

- D/S clouds are prevented from timing gaps during Beta-only transmissions
- Beta-only PHYs are prevented from issuing gap tokens in a hybrid network (since they don't know about legacy gap timings)
- A border PHY will guarantee a gap token is generated whenever a corresponding gap period has expired within the D/S cloud and, by extension, in all D/S clouds connected to the bus

ZAYANTE

Detailed Solution

Preventing Timing of Gaps during Packet Transmission

Particularly during beta-only traffic, we need to make sure that legacy devices don't detect any gaps.

(Example failure: during isoch transmission of beta-only packets, occurrence of a subaction gap in the D/S cloud would cause legacy nodes to fall into the asynch period too quickly and perhaps before their own isoch transmissions occurred.)

- Legacy formatted packets in a beta cloud are repeated by a border directly into any D/S cloud as normal (with DATA_PREFIX replacing the payload if the speed is too great)
- For beta-only packets, a border will begin generating DATA_PREFIX. Since the border must meet minimum DATA_PREFIX/DATA_END timings, and can not predict when the next legacy packet will arrive, it can not safely release DATA_PREFIX at any arbitrary point. Instead, it holds DATA_PREFIX until a legacy packet arrives. The legacy packet is then simply tacked onto the end of the DATA_PREFIX which allows the D/S cloud to return to idle.
- To ensure that the border PHYs don't get stuck in DATA-PREFIX, BOSS PHYs are required to issue a legacy null packet anytime the end of a subaction has been reached AND there are no more in-phase requests to grant AND the last packet sent was not a legacy packet.

ZAYANTE

Detailed Solution

Preventing Beta-Only PHYs from Issuing Gap Tokens

Normally, a beta-only PHY will issue a gap token at the end of a subaction when no eligible in-phase requests remain to be granted. To prevent a beta-only PHY from doing so in a hybrid network,

- Border PHYs with active D/S ports or active attached legacy Link announce their presence within a bus during Self-ID. This is recorded by all beta-capable PHYs.
- If at the end of a subaction the BOSS has no more in-phase requests to grant and it knows that it is in a hybrid network, control is passed towards the local root (after transmitting any necessary null packet to "free" the border nodes stuck in DP). No gap token is issued and the bus will fall IDLE.

Detailed Solution

Guaranteeing Legacy Gaps are Matched with Tokens

Border PHYs are responsible for issuing gap tokens whenever a relevant gap period is detected within the attached D/S clouds

- In a hybrid bus, border nodes reserve the exclusive right to issue gap tokens. If the duration of IDLE at a border reaches a gap timed threshold, that border will issue the corresponding gap token. A bus with multiple borders will end-up with duplicate gap indications as each of the border PHY's reach the gap timeout. This is of no concern given that the gap tokens are required to be idempotent.
- When the local root (which has become the BOSS by default) receives a gap token, it can grant any request that is now in-phase. For example, if the ASYNCH_START token is generated in response to the detection of a subaction gap, the local root can grant any outstanding asynchronous requests immediately upon receipt of ASYNCH_START.

Detailed Solution

BOSS PHY's Unaware of Isochronous Interval

- The BOSS PHY normally uses the current bus phase (asynch/isoch) to determine which of the arbitration requests it is receiving are in phase and are eligible to be granted.
- When a given beta cloud has no P1394b links, the beta PHY's will be unable to detect the start of an isochronous interval because no CYCLE_START_ODD/EVEN tokens are introduced into the beta cloud. Even a border node with a P1394a style link may not be able to infer the start of the isochronous interval if the link has no isochronous data to send.
- Even though the beta PHY's are unable to track the bus phase, there may still be some isochronous packets to send which are being originated from a D/S cloud. To make sure that requests from the D/S cloud can be granted properly absent any bus phase information, a new LEGACY_REQUEST has been created.
- The LEGACY_REQUEST has priority over any normal asynchronous or isochronous request. The LEGACY_REQUEST in itself is neither asynch nor isoch. It communicates to the BOSS PHY that the originator of the LEGACY_REQUEST has enough information to determine that it would be appropriate and valid to immediately grant the request.
- The BOSS must either grant the LEGACY_REQUEST or grant a higher priority request within ARB_RESPONSE_DELAY of receiving the LEGACY_REQUEST to prevent gaps from occurring.
- Unlike other BOSS request types, the LEGACY_REQUEST operates like it does in P1394a and is expected to be withdrawn if denied.

ZAYANTE

Detailed Solution

BOSS PHY's Unaware of Cycle Start Priority

In a beta cloud that does not contain the cycle master, there is no priority request mechanism that will prevent beta nodes from concatenating packets (ping-pong BOSS-ship). To prevent this, the border node that is the local root in a beta cloud sends a high priority BORDER request whenever a cycle start is expected.

- Done by "cycle start expected" signal from an attached P1394b or P1394a Link
- If no active "a" or "b" link is attached, the BORDER request is ALWAYS sent by the local root, effectively preventing all BOSS accelerations in the cloud

This is equivalent to the P1394a accelerate / decelerate control

Detailed Solution
Guaranteeing P1394a "Quiet Times"

Quiet Time #1 (Making sure a subaction gap is consistently heard)

- Quiet time #1 only has to be respected when detecting ACK_Missing, the end of the isoch period, or the end of self-ID. The quiet period is at risk if ack packets arrive late, if isochronous arbitration is granted late, or primary asynchronous packet arbitration starts too quickly.
- All nodes (including Beta-only) are required to meet RESPONSE_TIME and ARB_RESPONSE_DELAY, meaning data prefix or arbitration must be initiated by a responding PHY and repeated by intermediate PHYs within P1394a defined limits. The P1394a gap count analysis then applies and guarantees that ack packets and isochronous arbitration will be seen by all PHYs before the beginning of Quiet Time #1.
- Specifically, after transmitting a packet that marks the end of a subaction, the current BOSS must grant the next PHY immediately. If there are no valid requests to grant, control is passed towards the local root. When the local root becomes BOSS, it must grant a request within RESPONSE_TIME or refrain from generating any subsequent grant until an ASYNCH_START token is received/generated by a border or until a LEGACY_REQUEST is received. This is identical to the role a P1394a arbitration state machine plays in determining when it is okay to initiate arbitration.
- Border PHYs forward D/S style requests as high priority LEGACY requests. Given that all D/S PHYs respect Quiet Time #1 when generating a LEGACY_REQUEST, and given the beta-PHYs repeat legacy requests within ARB_RESPONSE_DELAY, LEGACY requests can only be present outside of the quiet period.

ZAYANTE

Detailed Solution

Guaranteeing P1394a "Quiet Times" (cont)

Quiet Time #2 (Making sure an arbitration reset gap is consistently heard)

- Quiet time #2 is at risk if a late arriving asynchronous request is granted, or if arbitration for the next fairness interval is granted too soon.
- If the current BOSS (which should be the local root) has no requests to grant at the time a subaction gap token (ASYNC_START) arrives, it enters the second quiet period and will wait until a border signals the ARBRST_* event or until a BORDER request is received.
- Border PHYs forward D/S style requests as high priority LEGACY requests. Given that all D/S PHYs respect Quiet Time #2 when generating a LEGACY_REQUEST, and given that beta-PHYs repeat legacy requests within ARB_RESPONSE_DELAY, LEGACY requests will only be present outside of the quiet period.

Border Request Mapping

D/S to BOSS

- RX_REQUEST from D/S child ports (only heard between active packet transfers) or requests from legacy Link are mapped to new LEGACY_REQUEST immediately. Note that the border PHY may not know the phase of the bus (isoch or asynch), so it can not generally try to map RX_REQUEST to a BOSS asynch or isoch request.

BOSS to D/S

- The border PHY always respects the P1394a quiet times when forwarding eligible BOSS requests into the D/S cloud.
- The border PHY determines which requests are eligible based on the phase of the bus, asynch or isoch, even or odd. If the border PHY is unaware of the start of the isochronous interval, then it can't possibly be receiving any isochronous requests. (The presence of isochronous requests implies the presence of P1394b links which are required to send CYCLE_START_EVEN/ODD tokens.) If the border PHY thinks the bus is in the asynch phase when it isn't, the border will still refrain from forwarding asynch requests too soon since it is observing the quiet times and knows it can't arbitrate until after a subaction gap elapses.
- LEGACY BOSS requests are always eligible.

Detailed Revisions to BOSS Operation in Support of Borders

ZAYANTE

BOSS Request Types and Priorities

Asynch Request Name	Priority Level	Comment
BORDER	7	Used to force control back to senior border (local root) when cycle start expected from parent D/S cloud
CYCLE_START_REQ	6	
NEXT_ODD/EVEN	5	
CURRENT	4	
NONE_EVEN/ODD	3	
NEXT_EVEN/ODD	2	
NONE_ODD/EVEN	1	

- add LEGACY_REQUESTs
  - Has priority over all asynchronous and isochronous requests, regardless of current bus phase
  - Coded as configuration request type

When BOSS is Permitted to Transmit

Explicit Grant

– End of subaction has been explicitly marked, it is safe for BOSS to use received GRANT immediately (within RESPONSE_TIME for hybrid bus) for in-phase beta request or a LEGACY_REQUEST.

Implicit Grant

– End of subaction has *not* been communicated. If receiving BOSS can independently determine that the subaction is concluded (e.g., bus is in isoch phase or last packet was an ack in the asynch phase, or request from a legacy Link), GRANT can be used immediately (within RESPONSE_TIME for hybrid bus)

Unarbitrated

– PHY can assume BOSS*ship* immediately and begin transmitting (within RESPONSE_TIME) to send an ACK or a PHY response packet.

When BOSS can/must Grant

Hybrid Bus

- At conclusion of each packet transmission, a beta-aware PHY must immediately grant an active request or it's parent as a default.
- A beta-aware PHY which receives a GRANT (thereby becoming BOSS) must either use the GRANT within RESPONSE_TIME or pass it along within ARB_RESPONSE_DELAY.
- For the special case of the local root, a grant can additionally be issued:
  - Within RESPONSE_TIME after receiving an ASYNC_START token.
  - Anytime after receiving an ARBRST_EVEN/ODD token.
  - Within RESPONSE_TIME after receiving a BORDER request.

Border-Imposed BOSS Requirements (1)

- Beta-only PHYs must meet RESPONSE_TIME when generating ACK's or queuing in-phase isoch arbitration.
- Beta-only PHYs must repeat arbitration within ARB_RESPONSE_DELAY, particularly propagation of GRANT and DATA_PREFIX.
- When explicitly granting a particular requesting port, other ports must start sending DATA_PREFIX to stop the timing of legacy gaps.
- When no borders are present, current BOSS is expected to issue gap tokens immediately. Control of the bus is retained by ending previous packet with DP, generating token, and then ending with DE/GRANT as appropriate.
- Packet ending symbols can now include DP, DE, GRANT, CYCLE_START_*, ARB_RESET_*, and ASYNCH_START. Only DP needs +/- disparity versions. As a byproduct, concatenated packets are now defined for beta-only.

ZAYANTE

Border-Imposed BOSS Requirements (2)

- New LEGACY_REQUEST type which is granted with a higher priority than asynch or isoch requests. It is granted or denied immediately without regard to bus phase. Furthermore, when it is denied, it will be withdrawn/cancelled as in P1394a. Due to the longer cables of P1394b and shorter packet sizes, some filtering of LEGACY_REQUESTs may be required.
- New BORDER request type (equivalent priority to CYCLE_START_REQ) to force the return of BOSS*ship* to the local root.
  - If local root has an attached Link, sent when cycle start is expected,
  - Used all of the time for local roots without an attached link.

ZAYANTE

Observations/Axioms

- For a Beta-Only device, Link, if present, is Beta-style link and able to mark each and every end of subaction. If link is not beta-style, then the attached PHY is a border PHY by definition.
- Explicit grant can only be issued after isoch packet if a P1394b link is present. However, it can be issued after an asynch packet by a Beta PHY (PHY response packet) or by a border node after an ack even when a P1394b link is not present.

Issues for Later Discussion

- Loss of Synch
- Timeout of Beta-only root must be > than worst case subaction gap detection time and token delivery time.

ZAYANTE

Interactive Examples