



- 
- 
- 
- 
- 

**HYBRID BUSES AND THE BORDER PHY**  
**A P1394A APPROACH**  
(UPDATED 12 OCT 1999)

---

Jerry Hauck

Michael Johas Teener

Colin Whitby-Strevens



**P 1 3 9 4 A A R B I T R A T I O N  
T U T O R I A L  
A R E F R E S H E R**



# 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.

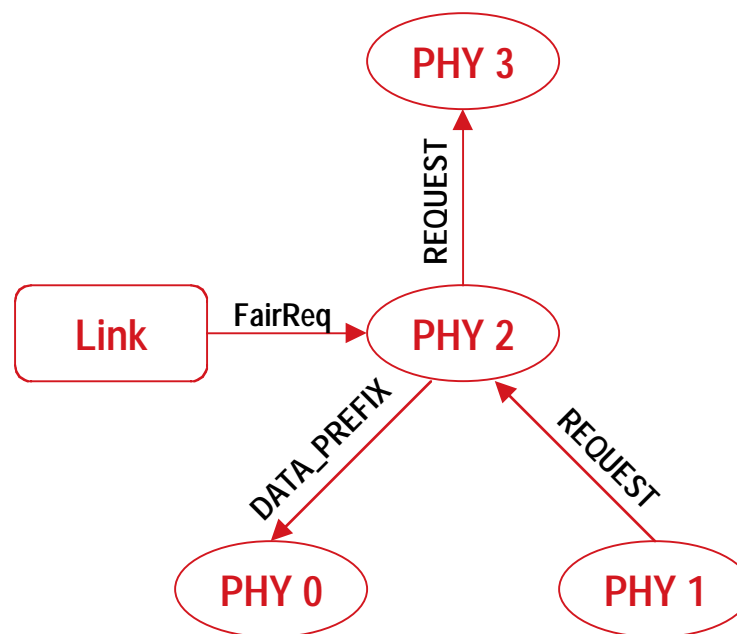
Requests from the link:

- IsoReq: Immediately forwarded
- FairReq or PriReq: forwarded at

Known end of subaction (ack)

Gap = Subaction + Arb Delay

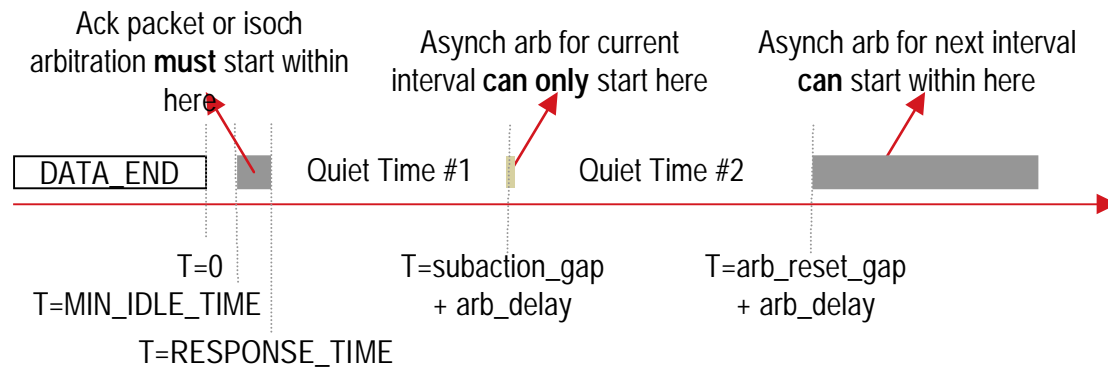
Gap  $\geq$  Arb Reset + Arb Delay



# P1394A ARBITRATION QUIET WINDOWS

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

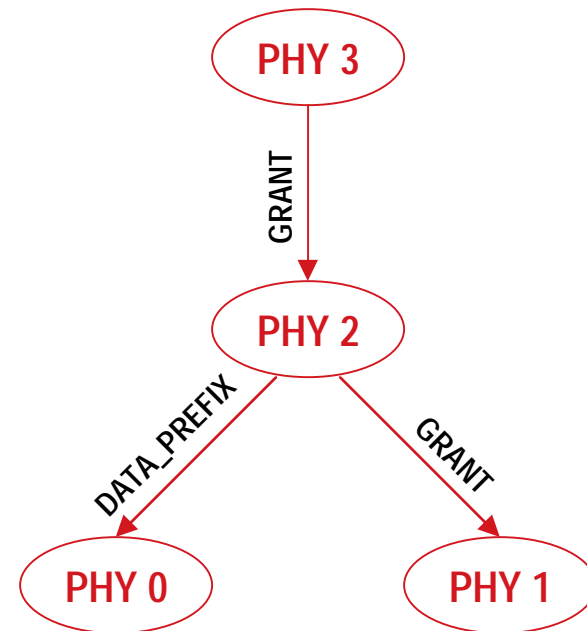
- 
- 
- 
- 
- 



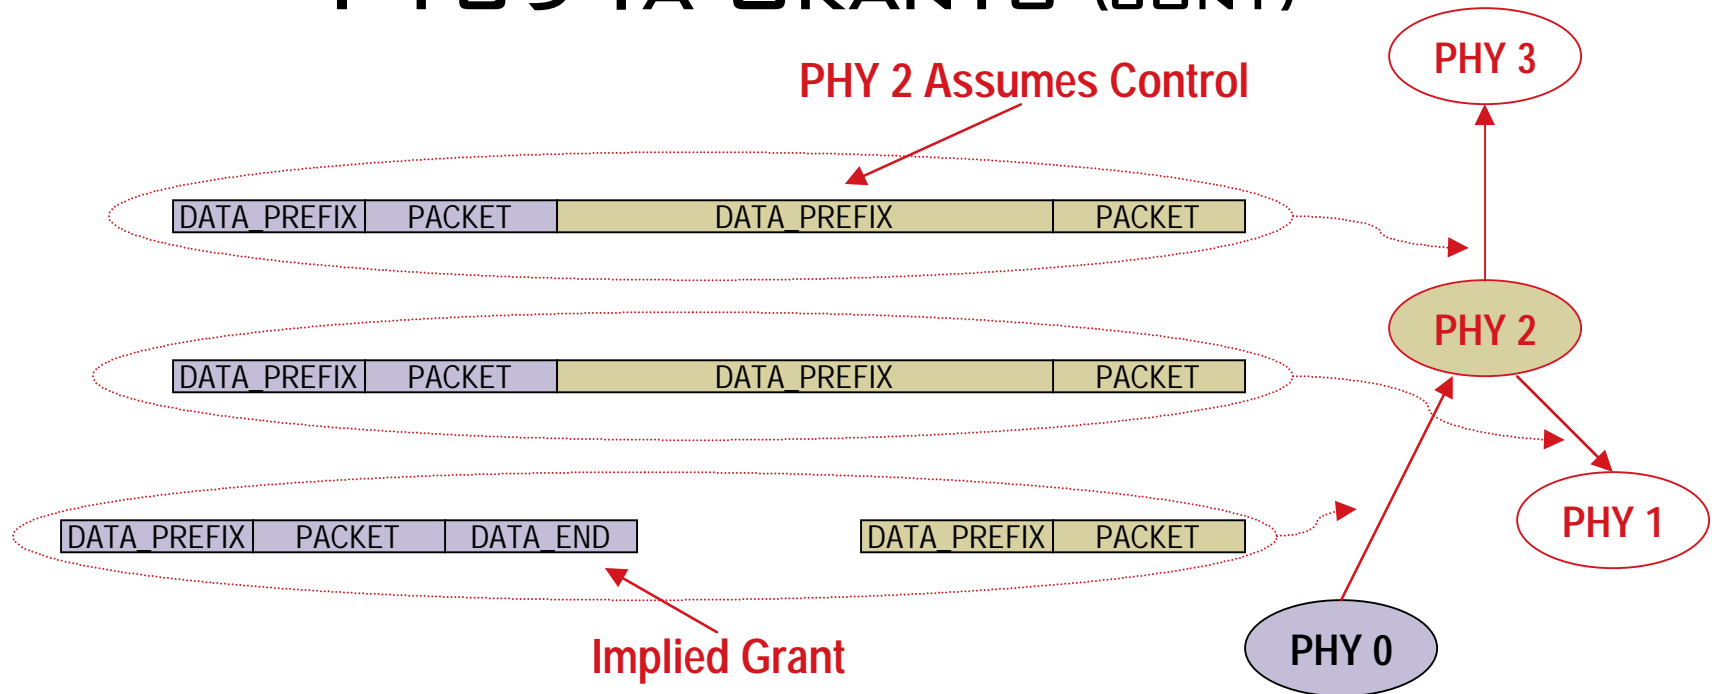
# P 1 3 9 4 A GRANTS

## 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

# 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 “preferred border” of a cloud automatically assumes BOSS*ship* after any extended period of inactivity.

## When does a PHY surrender BOSS*ship*?

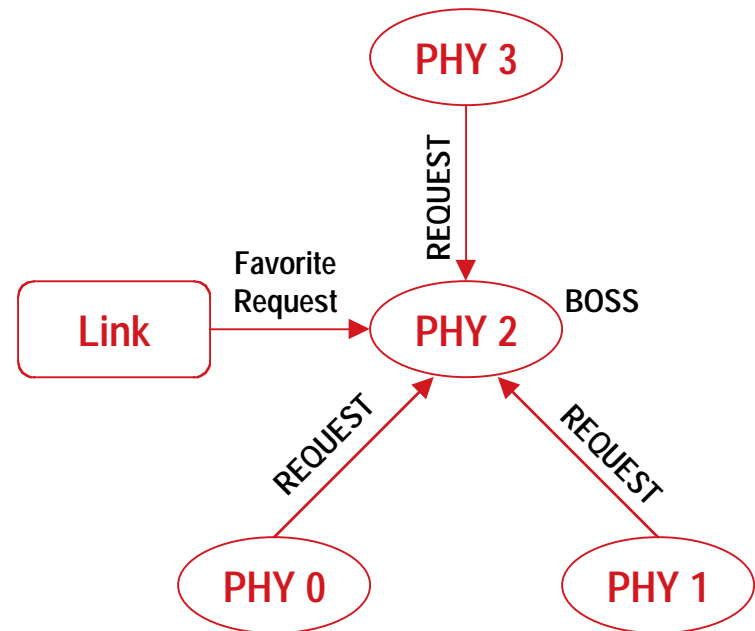
- Whenever a packet is received from a beta-mode port, the receiving PHY ceases to be BOSS.
- When an explicit or implicit GRANT is issued, the issuing PHY ceases to be BOSS.

# 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\_NULL 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)

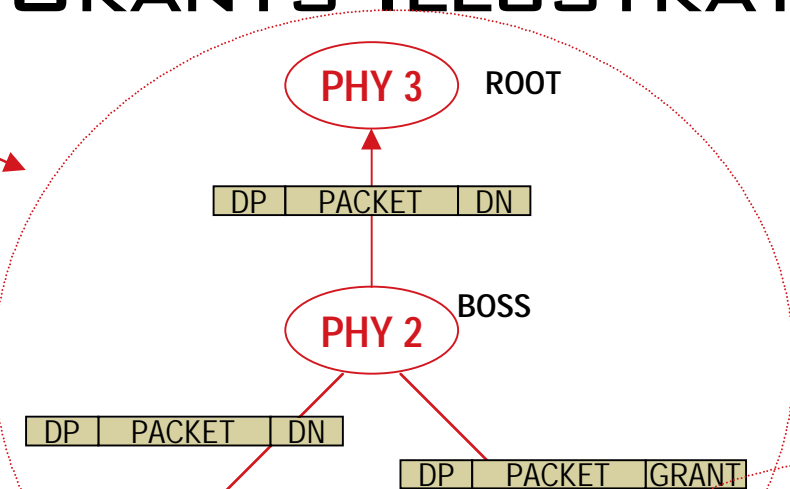
- All ports see DATA\_END, parent port interprets as an implicit grant.



# GRANTS ILLUSTRATED

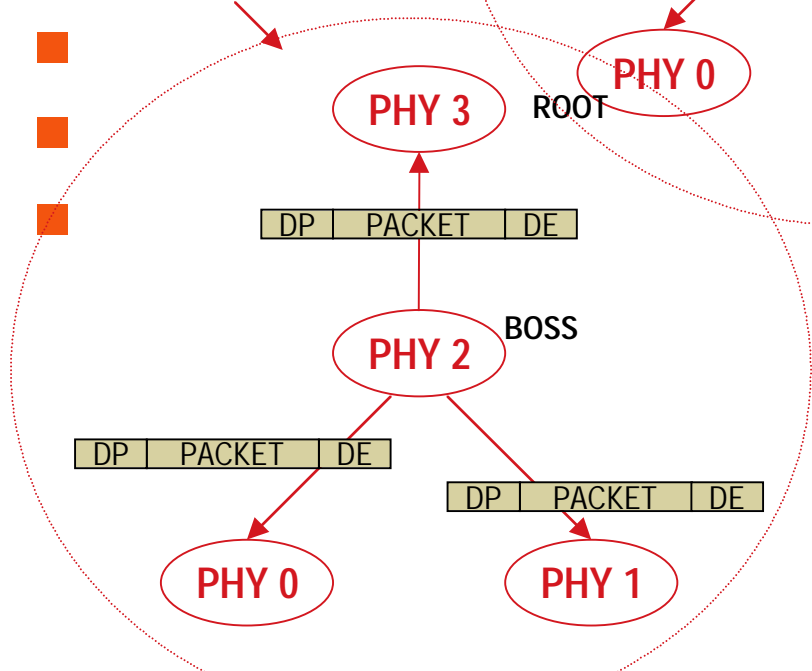
## Explicit-Loud Grant

Grant can be sent to either parent or child



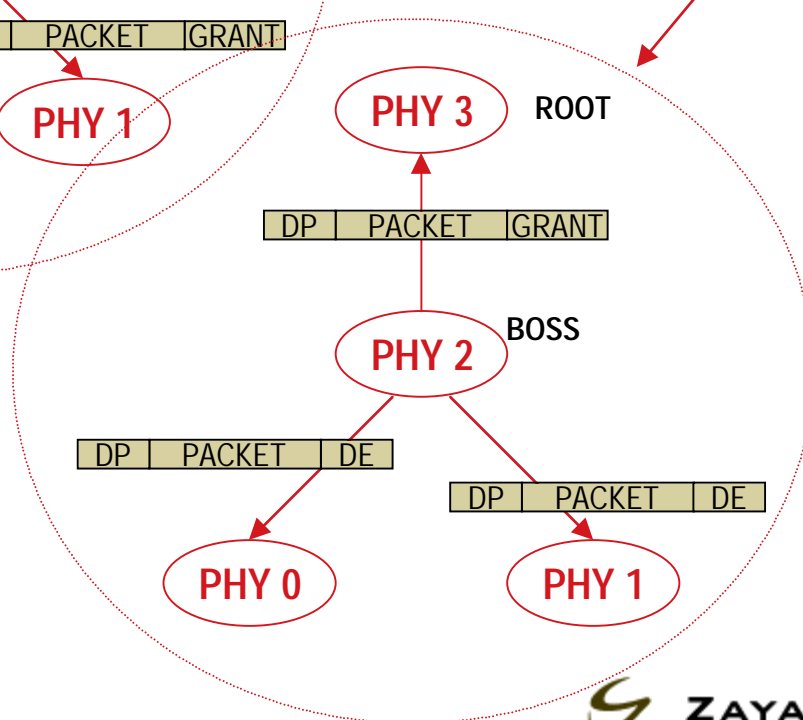
## Implied Grant

- Boss is repeating a packet that will have an explicit grant ...
- directed asynch, phy packet with response, need for ack timeout



## Explicit-Quiet Grant

Grant must be sent to root in this case (BOSS has no active requests)



# 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<sub>ship</sub> immediately and begin transmitting an ACK or a PHY response packet.

# WHEN BOSS CAN/MUST GRANT

## Hybrid Bus

- At conclusion of each packet transmission, a beta-mode PHY must immediately grant an active request or it's parent as a default.
- A beta-mode 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 preferred\_border, a grant can additionally be issued (if the root is within the cloud) or arbitration for a parent D/S cloud can begin:
  - Within ARB\_RESPONSE\_DELAY after the ARB\_DELAY following the issuance of an ASYNC\_START token.
  - Anytime after the ARB\_DELAY following the issuance of an ARBRST\_EVEN/ODD token.
  - Within ARB\_RESPONSE\_DELAY after receiving a LEGACY\_REQUEST.

# P1394B BUS PHASES

## Isochronous Interval

- Begins with P1394b link issues a CYCLE\_START\_ODD/EVEN token *and* a Cycle Start packet.
- 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.



- 
- 
- 
- 
- 

# 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 ...

# 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.

# 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 ...

# HYBRID BUS PROBLEM STATEMENT # 1

## SYNCHRONIZING PHY INDICATIONS

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

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

# 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.

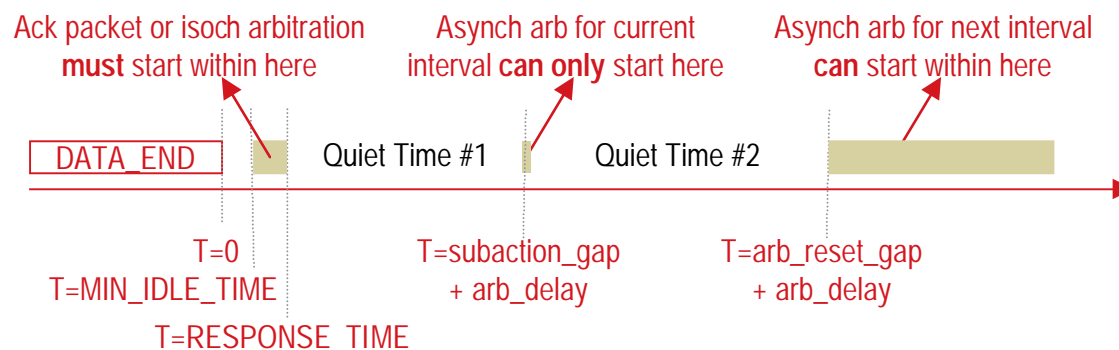
- BOSS currently defines asynchronous and isochronous pipelined requests. The BOSS PHY needs to understand the bus phase (isoch or asynch, even or odd) before it can grant an incoming request.

# 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)

# BORDER 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 single border PHY within each beta cloud will guarantee a gap token is generated whenever a corresponding gap period has expired within an attached D/S cloud and, by extension, in all D/S clouds connected to the bus

# 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.

# 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 preferred border (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

The preferred border PHY is responsible for issuing gap tokens whenever a relevant gap period is detected within the attached D/S clouds

- In a hybrid bus, the preferred border node reserves the exclusive right to issue gap tokens. If the duration of IDLE at the preferred border reaches a gap timed threshold, that border will issue the corresponding gap token.
- After the preferred border (which has become the BOSS by default) issues a gap token, it can attempt to 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 preferred border can arbitrate for any parent D/S cloud on behalf of any outstanding asynchronous requests received from within the beta cloud. (Note that the preferred border obeys P1394a timings and will wait an additional arb\_delay after issuing a token before initiating arbitration.)

# DETAILED SOLUTION

## BOSS PHY'S UNAWARE OF ISOCHRONOUS INTERVAL

The BOSS PHY normally uses the current bus phase (asynch/iso) 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 LEGACY\_REQUEST is defined.

- The LEGACY\_REQUEST has priority over any normal asynchronous or isochronous request. The LEGACY\_REQUEST in itself is neither asynch nor iso. 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.

# 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, beta PHYs refrain from passing a grant received from a junior port to either the link or another junior port when:

- – The bus is in the asynch phase
- – The cycle master is outside of the beta cloud (determined by observing the self-id issued by the root)
- – AND, a cycle start is expected

As in P1394a, links inform PHYs when a cycle start is expected by issuing `CYCLE_START_DUE` and `SEND_CYCLE_START_TOKEN` indications. If a PHY has no active link attached, it assumes that a cycle start is always imminent.

# DETAILED SOLUTION

## INTRODUCING CYCLE START TOKENS

Nominally, a P1394b cycle master marks the beginning of a cycle start interval with both a cycle start token AND a cycle start packet. However, when the cycle master is located outside of a given beta cloud, the cycle start packet is stripped of its corresponding token.

- 
- 
- 
- 
- 

To introduce a cycle start token within a given beta cloud, beta links instruct beta PHYs whenever a cycle start packet has been received via the SEND\_CYCLE\_START\_TOKEN service. In turn, the PHY will issue the appropriate CYCLE\_START\_EVEN/ODD token on the cable if one has not already been received.

# DETAILED SOLUTION

## GUARANTEERING 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 preferred border. When the preferred border becomes BOSS, it must grant a request within RESPONSE\_TIME or refrain from generating any subsequent grant until a subaction gap is timed and ASYNCH\_START token is generated, 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.

# 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 is the preferred border in this case) has no requests to grant at the time a subaction gap token (ASYNC\_START) is issued, it enters the second quiet period and will wait until an arbitration reset gap is timed and an ARBRST\_\* token is issued, or until a LEGACY\_REQUEST is received before granting.
- – 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 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\_REQUESTs are always eligible to be forwarded.

# 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\_NULL 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 DN, generating token, and then ending with DE/GRANT as appropriate.

Packet ending symbols can now include DP, DE, GRANT and DN. Only DP needs +/- disparity versions.

# BORDER-IMPOSED BOSS REQUIREMENTS (2)

- LEGACY\_REQUEST type 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 is required.
- A grant (explicit or implicit) received on a junior port can not be passed as an explicit grant to another junior port or the link port when:
  - The bus is in the asynch phase,
  - The cycle master is outside of the local beta cloud,
  - AND, a cycle start is expected.

# INTERACTIVE EXAMPLES



- 
- 
- 
- 
-