



-
-
-
-
-

IEEE P1394b

Report to the 16th DVB-IHDN meeting (18th May 1999)

Colin Whitby-Strevens

Zayante Inc

Agenda

Introduction

P1394b summary

■ P1394b status

■ – stable items

■ – work to do

■ Actions to address DVB concerns

■ Summary

Introduction

P1394b has been in development for 2.5 years

DVB-IHDN needs a stable specification

- for long-haul, rather than high speed

■ Much of P1394b is stable

■ Some critical items need accelerating to completion

■ Philips made a presentation to April 29th meeting

- simplification decisions taken

■ Review P1394b

- main features
- describe the actions to accelerate stability
- schedule

Firstly, a word on 1394a

1394a: improvements in efficiency

- Arbitration acceleration
- Reset improvements
- Suspend/resume
- Miscellaneous

P1394a nearly finished

- Ballot in Q2, 98
- Other issues arising from implementations and liaison
 - IP/1394
 - P1394.1
- Recirculation ballot in July 99
 - BRC meets on June 24/25 (Chicago)

1394b: go faster, go further

Full backwards compatibility with 1394-1995 and 1394a

Extend speed ranges to 800 and 1600 Mbit/sec

- 3200 Mbit/sec and higher architectural support

Support long distance systems

- 50m, 100m

New connection model

New encoding and signaling scheme

Arbitration improvements

Miscellaneous

1394b backwards compatibility - some definitions

Refer to the new signaling and arbitration scheme as **Beta mode**

Refer to the 1394-1995/1394a signaling and arbitration as **DS mode**

- A **bilingual port** is a port which is capable of operating either in Beta mode or in DS mode

- connection management performs the necessary autonegotiation

- A **border PHY** has some ports which are operating in Beta mode and some ports which are operating in DS mode

- individual ports may be dedicated or bilingual
- example - a simple long-haul repeater - has one Beta and one DS (or bilingual)
- also need a border PHY if there's a mixture of Beta and DS when taking the link layer into account

1394b speeds and feeds

1394a cable and connectors designed for cluster interconnect

- optimized to 4.5m

1394b aims to go faster, further

- PC's need S800 and above
- Home network requires 100m
 - (source VESA Home Network Committee)

Range of trade-offs

- STP, UTP5, POF, HPCF, Glass fibre

1394b media - common properties

All media use same encoding scheme and binary transmission

Continuous full-duplex operation

■ Scrambling for emissions control

- scramble before encoding

■ 8B10B (modified) for DC balance and limited run-length

- permits AC coupling and excellent control of signal integrity

■ BER of 10^{-12} for all specifications

■ Same connection management

- allows same PHY to be used for different media
- external optical transceiver
- electrical isolation at the connector

STP (short haul)


New Beta-mode connector required

- need connector/cable to provide road-map to S3200

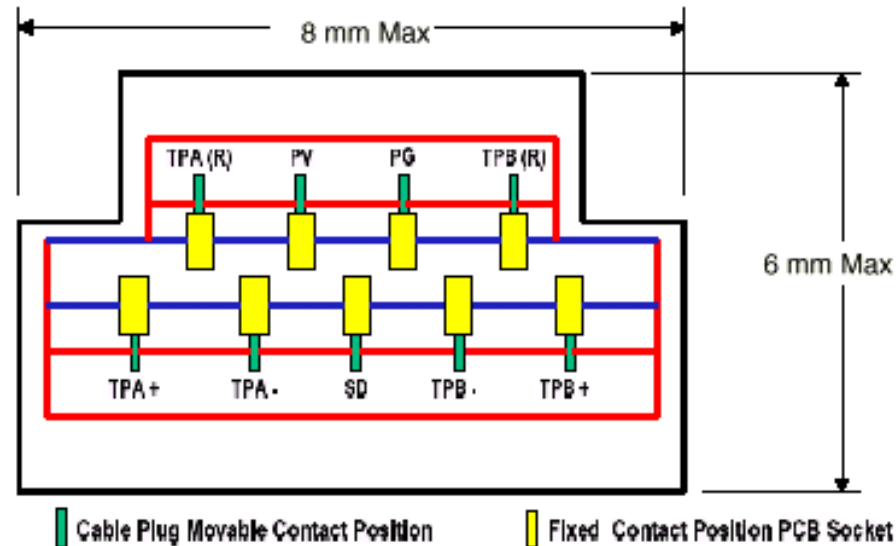
Bi-lingual connector will be a variant of the new Beta-mode connector

- supports 1394-1995, P1394a and P1394b
- note: existing 4-ckt and 6-ckt connector not used for Beta mode

Legacy connection



Beta connector proposal (preliminary)



Very nearly as small as the current 4-ckt connector

– much smaller than current 6-ckt

Cheap, very low cost, inexpensive and affordable

Includes power

Keyed for bi-lingual (bi-lingual plug will NOT go into a beta-only socket)

UTP-5 (ISO/IEC 11801 ch7)

Cat-5 Unshielded Twisted Pair and RJ-45 connector

- very commonly installed in commercial buildings
- some installation in new-build houses
- low cost, easy installation

Well-known for ethernet

TX on pins 1/2, RX on pins 7/8

- transformer isolation

Similar electrical specification as other standards (1V pk-pk binary)

Adaptive equalization in the receiver

- allows 100m operation


Plastic Optic Fiber/ Hard Polymer Clad Fiber

POF: 1000 μm step index multimode plastic optic fiber

- suitable for up to 50m

HPCF: 225 μm graded index multimode hard polymer clad fiber

- - suitable for up to 100m
- Same transceiver spec for both
- Low cost, easy to install
- Fiber alleviates emissions and interference problems
- PN connector



Glass fiber (MMF)

Leverage VCSEL (Vertical Cavity Surface Emitting Laser) technology

Leverage Fibre Channel and Gigabit Ethernet specifications

- 50 micron multimode fiber (MMF)
 - – Media gives road-map to S3200
 - – Transceiver spec development required
- LC connector
- Leverage Fibre Channel/ Gigabit Ethernet specifications

1394b media summary

	Reach	S100	S200	S400	S800	S1600	S3200
UTP5	100m	•					
POF	50m	•	•				
HPCF	100m	•	•				
MMF	100m			•	•	•	(•)
STP	4.5m			•	•	•	(•)

1394b arbitration

Main issue is inefficiencies introduced by arbitration “gaps”

1394a starts to address this issue


- 1394b exploits new media properties to overlap arbitration and data transmission (“Boss mode”)
- Boss arbitration is simple for Beta-only systems
- Boss arbitration gives rise to subtle issues in border nodes
-

Background: 1394 operation


Wait for idle bus

Arbitrate

- Send packet
- Responder sends ack



Bus timing for asynchronous arbitration

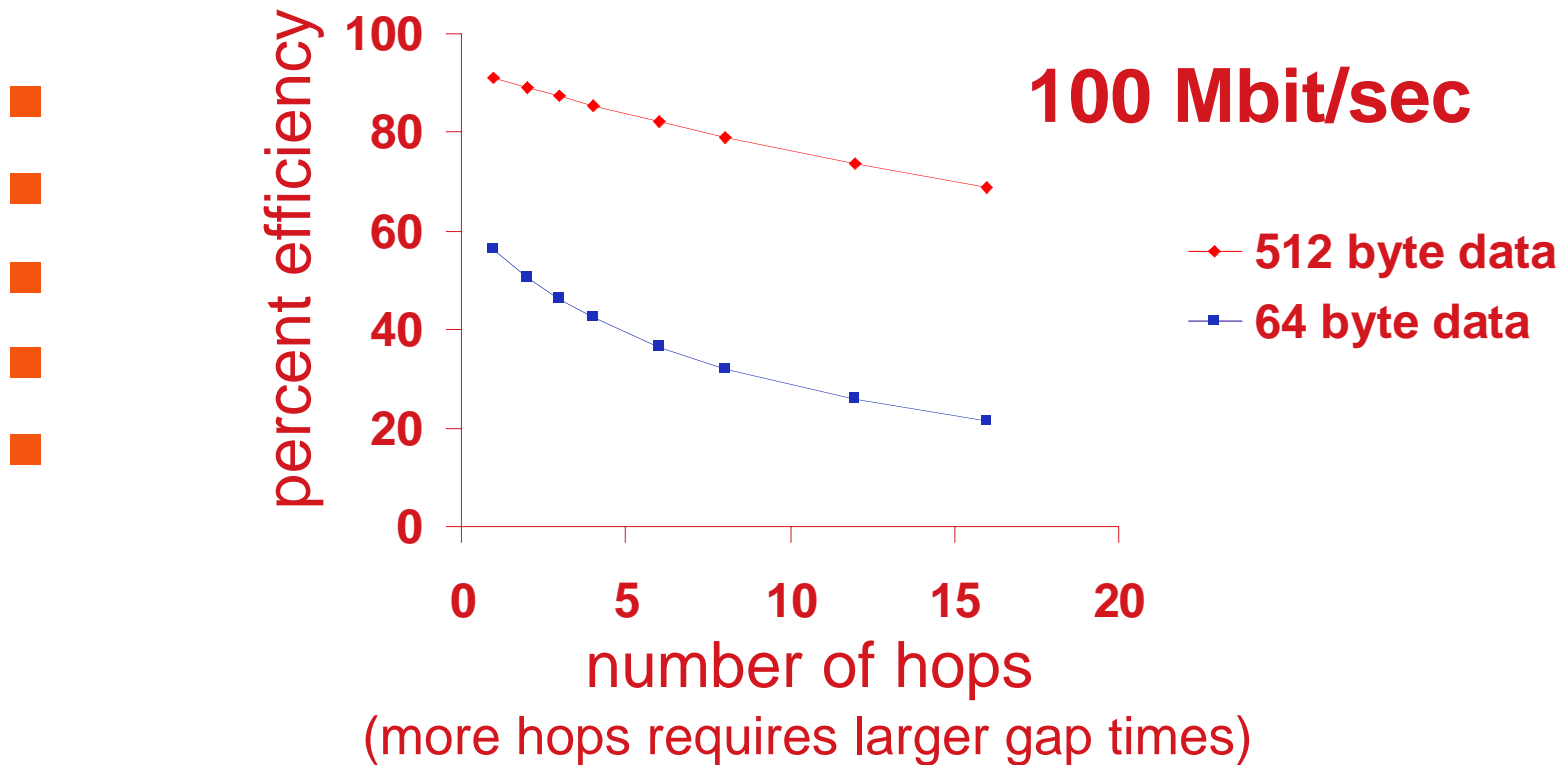


- Subaction gap time is greater than round trip delay time for most distant pair of nodes on a 1394 bus

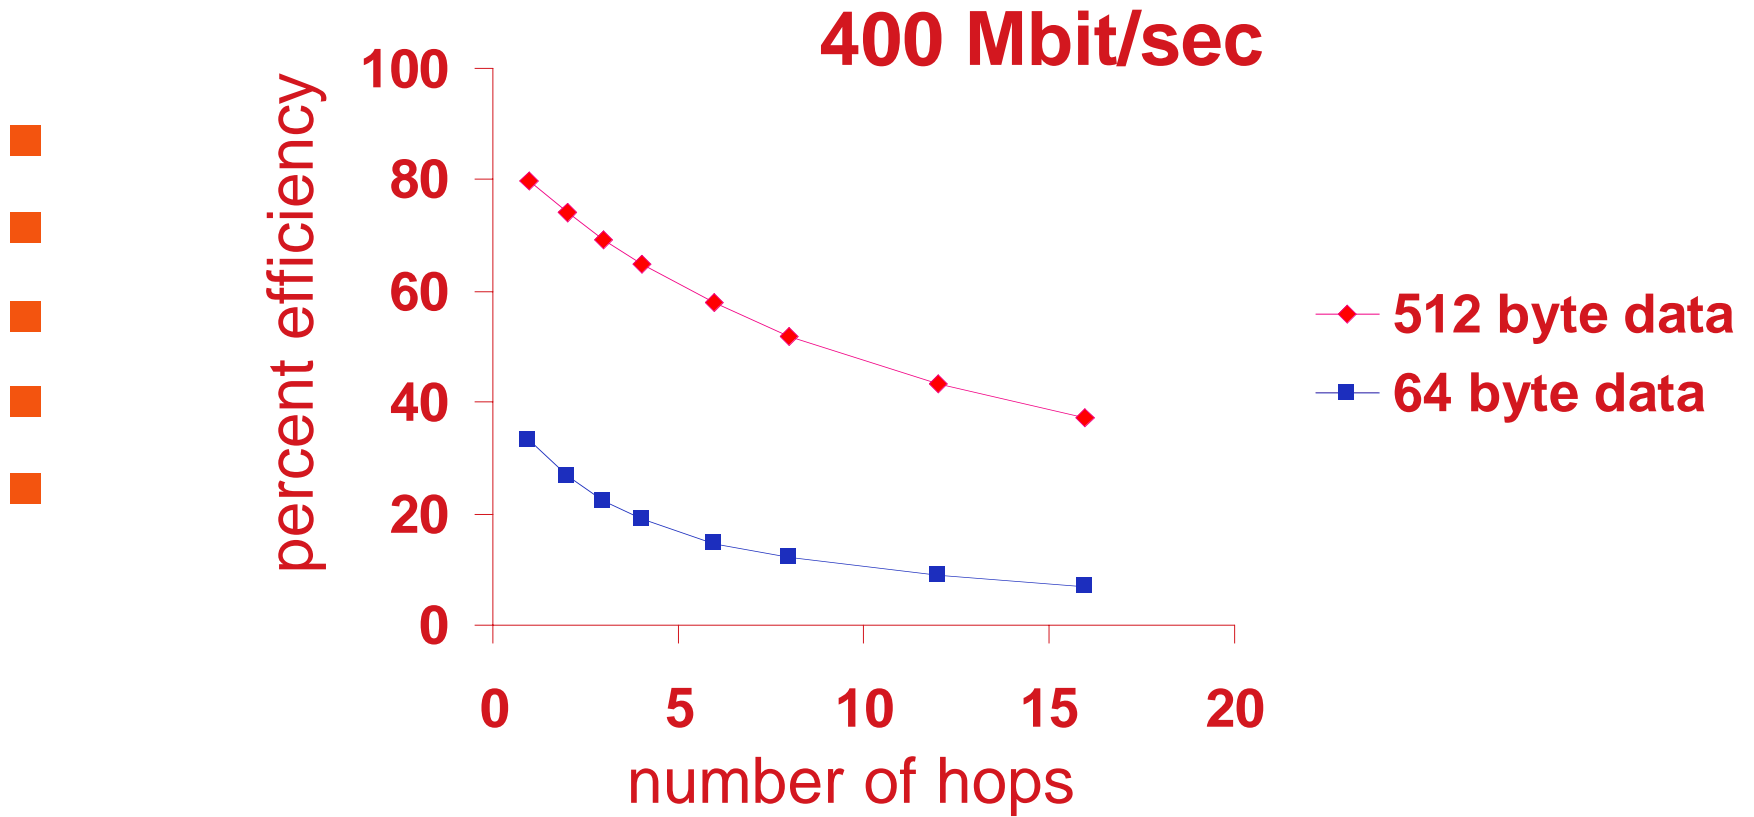
For asynchronous arbitration, waiting a Subaction gap time after passage of last packet before start of new arbitration insures that new arbitration will not interfere with an ack

Efficiency effects of gap time

OK for large packets, terrible for small ones



Gets worse for high speeds



1394a starts to address the issue

ACK accelerated arbitration

Fly-by arbitration

- Multi-speed concatenated packets

- – helps isochronous transfers

- Priority arbitration

- Transaction fairness

- – no limitation on responses (they don't count against the fairness interval)

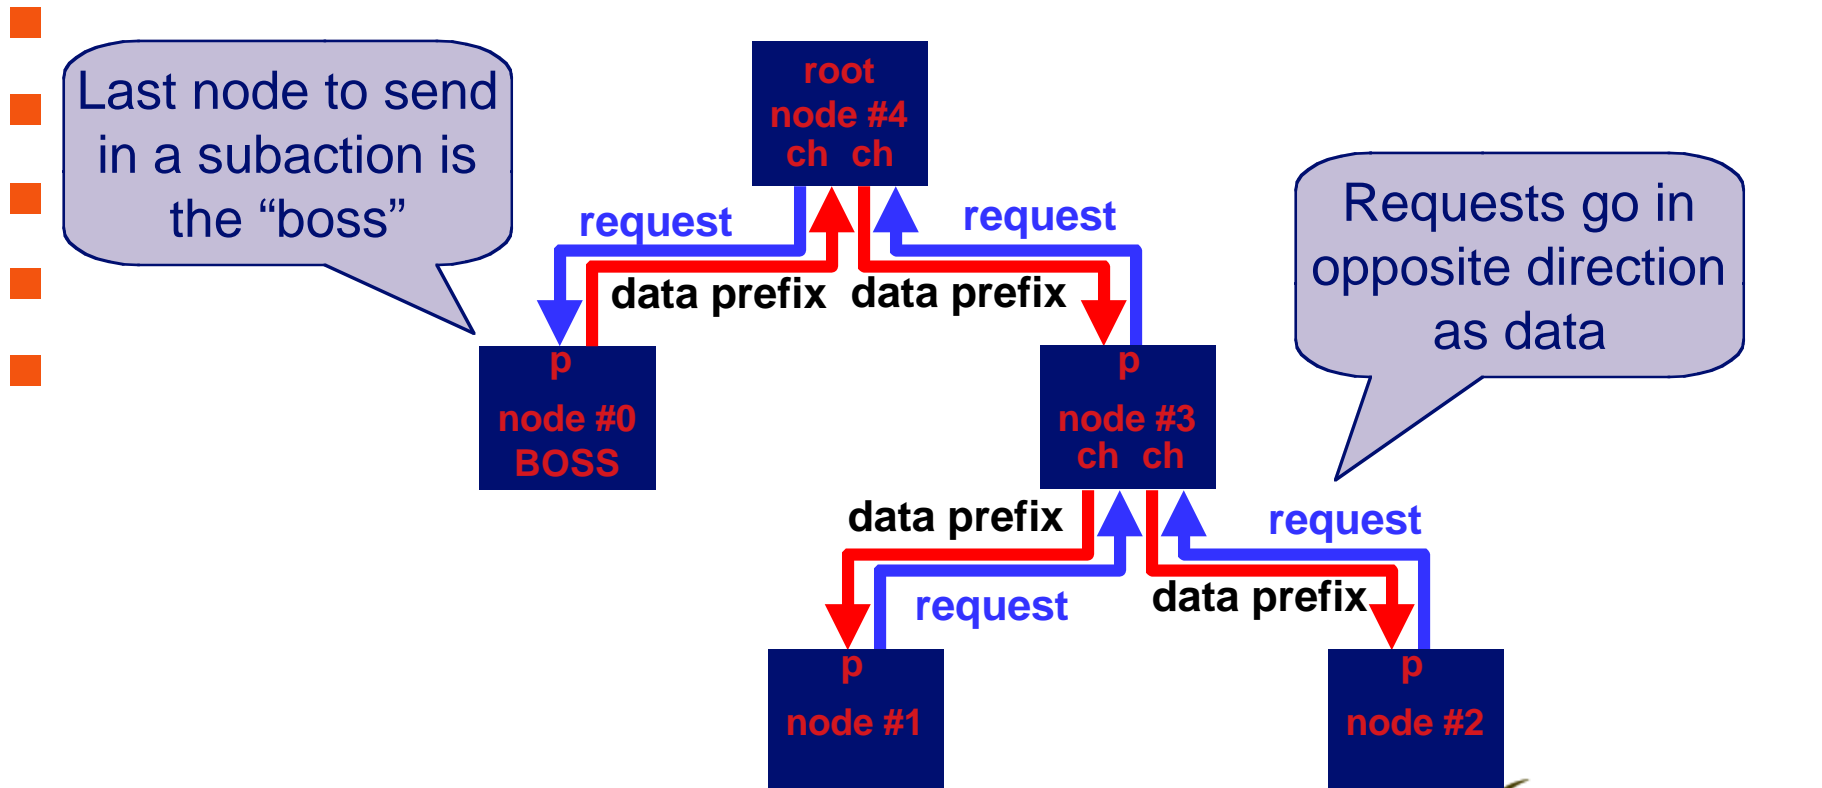
Fairness optimization

- “unfair fairness”

- 64 64 transactions per fairness interval, regardless of the no of nodes on the bus

1394b arbitration improvements (1)

Full duplex connections means that arbitration can be pipelined!



1394b arbitration improvements (2)

General rule: send requests continually, and in every direction that is not carrying packet information

Sender is the “Boss” if there is no other implicit winner of arbitration

- if sending a directed asynchronous packet, then the sender of the ACK is the implicit winner ... will become the Boss.

Eliminates gaps

- use tokens to indicate sub-action and arb-reset “gaps”
- More cool technology, David LaFollette of Intel, Jerry Hauck and Michael Johas Teener of Zayante

1394b miscellaneous

Startup methods (connection management) use autonegotiation

- May be connected via optical (no DC bias), so no direct connection signal
 - use “tones”
- May connect using copper
- Always use “beta” mode if possible

Loop breaking

- Loops are automatically removed during initialization

P1394b status - history

Started in 1997

- currently hold 10 meetings per year

Decided on speeds and media in 1997

- POF and Glass media and connectors
- UTP added in early 1998

Decided on a single signaling scheme (scrambled 8B10B) in 1997

First version of optical specifications in 1997

- Glass extended to include S400 in mid 1998

Boss mode proposal in September 1997

Connection management techniques agreed in early 1998

First version of the draft to combine Boss mode, port mechanisms and connection management in June 98

Early '99 - reorganized to have Steering Committee of technical leaders

P1394b status - today

Chapter	Topic	Status	Comments
4	High speed connectors and cables	WIP	New Beta-only connector
5	Electrical specification	Stable	Revised jitter budget
6	PHY test modes	Open	Improved jitter methodology
7	Optical specification (Glass)	Stable	Revised jitter budget
8	POF	Stable	
9	UTP	Editorial	Handover to the editor for June draft
10	Beta mode port	Stable	Tracking changing Boss/border requirements
11	Connection management	Stable	Recent updates to P1394a BRC
12	PHY-Link	WIP	Recent strategic decisions
13	Register map	Stable	
14	Arbitration and control	WIP	Boss - editorial New simplified border - stable in July Open issue- Loop healing
15	C-code collation		
16 Annex K	Bulk cable specification	Stable	
17	Miscellenary		

P1394b status - arbitration

Work on Boss mode slowed in 2nd half of 1998

- Hybrid operation proved difficult
- “redeployment” of key architect
 - Zayante picked up the ball, but as a background activity
- Completion of P1394a took priority for key people
- Apple patent issues diverted attention

■ Succession of proposals for hybrid operation

■ Finally (March 99) got to a stage for hybrid operation spec which would

- work in any configuration
- not impact the performance of Beta-only systems
- not require all PHYs to be bilingual
 - all the “pain” goes in the border node
- detail design not done

P1394b - DVB inputs

Presentation at the 29th April 99 meeting by Tom Suters

- the DVB IHDN requirement needs the P1394b long-haul solution asap
- most anticipated devices will be P1394a
- key application is the “long-haul” repeater
- physical media and signaling specs for this are stable
- no opportunity for Boss enhancements in this application
- border functionality looks too complicated
- hybrid operation specification is taking too long

Propose “sharing the pain”

- simplify the border node
- accept some extra Boss complexity and/or loss of performance

Proposal accepted

P1394b - Boss goals

Beta-only sought to cure some of the non-scalability of P1394a

- occurs when scaling upwards in speed and/or topology breadth.

Some of the items contributing to the lack of scalability which were targeted:

- Fixed-time delimiters (DATA_PREFIX and DATA_END, ~440ns/packet)
- Arbitration time (average of 5 μ s per packet in large topology?)
- Gaps (10-20 μ s per instance, large topology)

Beta-only packet formats addressed #1 while BOSS tackled #2 and #3.

Hybrid networks of 1394-95/P1394a and P1394b

Goal is (was) to retain the improvements in all three areas.

- allow arbitration within D/S clouds to happen concurrently with beta-only traffic in the Beta cloud.

- – Drive IDLE into the D/S cloud whenever a beta-only packet transmission occurred

- Needed to take care in synchronizing the starts and ends of isochronous periods and fairness intervals

- – results from the loose coupling between the two arbitration domains

- – Additionally, Firewire core time outs (bus reset if 20 μ s without hearing a response to arbitration) required additional effort to enable smooth interoperability.

The state of the proposed border node with concurrent arbitration was (April 99) that it has been described and scrutinised at a high level, but not reduced to implementation.

P1394b - simplified border node

In response to DVB input, to move forward with border node definition, we've made a fundamental shift in our target.

Hybrid network will no longer attempt to retain concurrent arbitration for the D/S cloud as a solution for #2 and #3 above.

- No longer drive IDLE into the D/S clouds during beta-only transfers.
- Data packets will be sent bus wide as is done today, with DATA_PREFIX substituting for the payload where speed filtering is necessary.

Need to take care as the smallest sized Beta-only packet is smaller than the shortest null packet we can send into a D/S cloud (DP + DE + MIN_IDLE = ~440 ns).

- Several possible solutions

Simplified border - gap synchronization

D/S and Beta clouds synchronized around packet transmissions.

An additional and necessary step is to get the start/ends of asynch and isoch intervals synchronized.

- – A beta-only network has no need for gaps; a hybrid network w/o concurrent arbitration has to slow down and make sure the "gaps" are synchronized in all clouds. Would like to accomplish this without crippling a beta-only bus.
- – When present, border nodes consistently send out a(new) low priority request
- – When a beta-only phy begins to issue a grant, it's prevented from advancing out of the current isoch or asynch interval due to the pending request from the border(s)
- – So after all of the beta nodes have completed transmission within the current asynch/isoch interval, BOSS will land on one of the borders present.
- – After sending the null packet noted above, the border will allow the bus to go IDLE so that D/S nodes can arbitrate. BOSS will grant any legacy or beta request per P1394a timing rules and will also generate the asynch_start or arbrst_odd/even tokens if the IDLE duration crosses the subaction/arb_reset gap boundary.

Benefits of simplified border node

Effectively lays 1394a style arbitration over the Beta domain

- familiar to 1394a designers
- simpler to develop the standard
- simpler to develop implementations

Removes need for bus management magic for detecting imminent gaps

Removes timer required for dealing with Firewire cores (1394-1995)

No configuration management software necessary

No major modifications to beta-only PHY operation

- have to deal with new “legacy” request and give it priority

When no borders are active, the bus automatically runs optimally

- A border node advertises its presence with the low level arbitration

Loop healing

Loops prevent operation of 1394

- always a known “feature”

Is it a real problem?

- P1394b is currently addressing this issue
- Solution now proposed
- – English description, not the detail
- – quarantines a putative new connection if a loop would result
- – but subsequently allows the connection if another disconnection would otherwise break the bus
- – isolates 1394 clouds with loops

Should we do the work to put it in?

Should it be optional?

- Current solution requires cooperation of both nodes involved in a connection

Schedule for P1394b arbitration

What	Who	When
Boss state machine	Mike Teener	(already in the spec)
Conf call and working meeting on Boss	Mike Teener, Jerry Hauck, Colin W-S, Maarten DeVries	18/19 May
Update Boss C code	Mike Teener	21 May
Border State machines and C code	Jerry Hauck/Mike Teener	28 May
Reset, Tree-ID, Self-ID	Colin W-S	28 May
Border description	Maarten DeVries	31 May
Loop healing description	D Wooten	31 May
Next draft published	E Hannah	1 June
Loop healing implementation	Colin W-S	7 June
P1394b meeting (Houston)	All	7/8/9 June
Review, simulate, stabilise	All	10 June - 9 July
Proposed stable draft published	E Hannah	12 July
Draft accepted at P1394b meeting (Briarcliff)	All	19/20/21 July

Conclusions

P1394b is nearing completion

Key issue of improved arbitration mechanisms

- has been worked on for 1.5 years

- Hybrid bus operation had taken a long time to finish

- DVB suggested simplification to speed up completion of spec

- CE companies now committing full-time resource to achieving this

- – very welcome, thank you!!!

P1394b has a revised schedule for completing the critical issue by July