FireWire Design Guide

Abstract: Guidelines for implementing FireWire (IEEE 1394) ports on both high level complex devices such as personal computers and simple devices such as inexpensive consumer electronics products. Includes references to normative specifications as well as full reference designs that can be used as is or, with the careful consideration of the trade-offs, as a starting point for a more optimized design.

1394 Trade Association Specifications are developed within Working Groups of the 1394 Trade Association, a non-profit industry association devoted to the promotion of and growth of the market for IEEE 1394-compliant products. Participants in Working Groups serve voluntarily and without compensation from the Trade Association. Most participants represent member organizations of the 1394 Trade Association. The specifications developed within the working groups represent a consensus of the expertise represented by the participants.

Use of a 1394 Trade Association Specification is wholly voluntary. The existence of a 1394 Trade Association Specification is not meant to imply that there are not other ways to produce, test, measure, purchase, market or provide other goods and services related to the scope of the 1394 Trade Association Specification. Furthermore, the viewpoint expressed at the time a specification is accepted and issued is subject to change brought about through developments in the state of the art and comments received from users of the specification. Users are cautioned to check to determine that they have the latest revision of any 1394 Trade Association Specification.

Comments for revision of 1394 Trade Association Specifications are welcome from any interested party, regardless of membership affiliation with the 1394 Trade Association. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments.

Interpretations: Occasionally, questions may arise about the meaning of specifications in relationship to specific applications. When the need for interpretations is brought to the attention of the 1394 Trade Association, the Association will initiate action to prepare appropriate responses.

Comments on specifications and requests for interpretations should be addressed to:

Editor, 1394 Trade Association
315 Lincoln, Suite E
Mukilteo, WA 98275
USA

1394 Trade Association Specifications are adopted by the 1394 Trade Association without regard to patents which may exist on articles, materials or processes or to other proprietary intellectual property which may exist within a specification. Adoption of a specification by the 1394 Trade Association does not assume any liability to any patent owner or any obligation whatsoever to those parties who rely on the specification documents. Readers of this document are advised to make an independent determination regarding the existence of intellectual property rights, which may be infringed by conformance to this specification.

Published by

1394 Trade Association
315 Lincoln, Suite E
Mukilteo, WA 98275 USA

Printed in the United States of America

IEEE Copyright

Portions of this specification are copied from published IEEE standards, by permission.

The source documents are:

IEEE Std 1394-1995, Standard for a High Performance Serial Bus

IEEE Std 1394a-2000, Standard for a High Performance Serial Bus – Amendment 1

The IEEE copyright policy at http://standards.ieee.org/IPR/copyrightpolicy.html states, in part:

Royalty Free Permission

IEEE-SA policy holds that anyone may excerpt and publish up to, but not more than, ten percent (10%) of the entirety of an IEEE-SA Document (excluding IEEE SIN books) on a royalty-free basis, so long as:

1) proper acknowledgment is provided;

2) the 'heart' of the standard is not entirely contained within the portion being excerpted.

This included the use of tables, graphs, abstracts and scope statements from IEEE Documents

Checklist

All designers of devices with FireWire ports should look for the following:

Contents

Figures

Tables

2. Mechanical

2.1 Connectors

Item #2.1: Connectors (both plugs and sockets) shall meet the requirements laid out in the IEEE standards.

Sockets must strongly resist “backwards” insertion

This means that the “narrow” part of the socket shall not have a split unless there is a mechanical reinforcement that prevents the narrow end from spreading when a plug is inserted backwards. It further means that the mechanical tolerences specified in [2], [3], and [4] *must* be met. In particular, the socket should meet the following insertion force requirements:

Table 2-1—Insertion and removal force *(tentative)*

Condition	Insertion/removal force	Comment
6-pin plug correctly aligned inserted into 6-pin socket 9-pin plug correctly aligned inserted into 9-pin socket	< 1 N	measured using a plug that is at the maximum dimensions allowed by 1394, and maximum value before any of the VG, VP, or TPx contacts mate, and must have a noticeable “click” or “snap” when fully inserted
6-pin plug correctly aligned removed from 6-pin socket 9-pin plug correctly aligned removed into 9-pin socket	> 1 N	measured using a plug that is at the minimum dimensions allowed by 1394, and minimum value before any of the VG, VP, or TPx contacts separate, and must have a noticeable “click” or “snap” when removed
6-pin plug reverse aligned inserted into 6-pin socket 9-pin plug reverse aligned inserted into 9-pin socket	> 10 N	measured using a plug that is at the minimum dimensions allowed by 1394, and minimum force before any of the VG, VP, or TPx contacts mate, and socket must be noticeably damaged by reverse insertion (allows diagnosis of improper use)

Item #2.2: Plugs shall have an overmold that includes strong tactile cues for orientation.

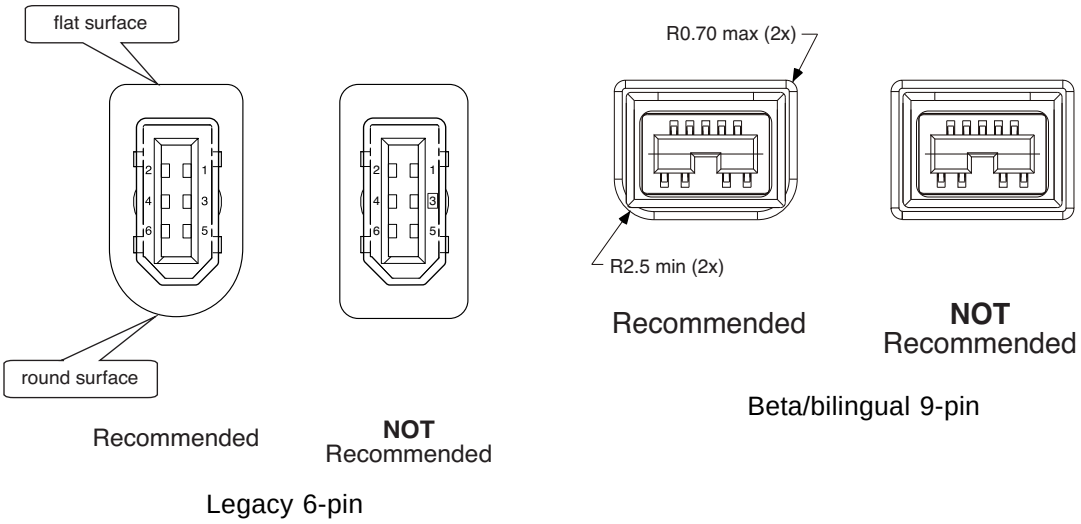

Figure 2-1—Plug overmold

Overmold cues are to provide hints for where the user should put the thumb.

a) Flat surface corresponds with “flat” surface on narrow axis
b) Round surface corresponds with “angled” surface on narrow axis

This will assist in preventing attempted backwards connection (particularly if the connector placement rules described below are followed).

Item #2.3: Sockets should be oriented so that “thumb” part of plug is on top or on the left

Sockets should have a standard orentation to aid the “blind” insertion of the plug. Since the plug overmold has features to encourage the placement of the thumb in a particular spot, and most people prefer to put their thumb on top of a plug as they insert it, the socket itself should be oriented so that the “thumb” feature is on top as viewed looking at the bulkhead from the outside. And since most people are right handed,. the alternate positioning should have the “thumb” feature on the left. Figures 2-2 and 2-3 illustrate the concept.

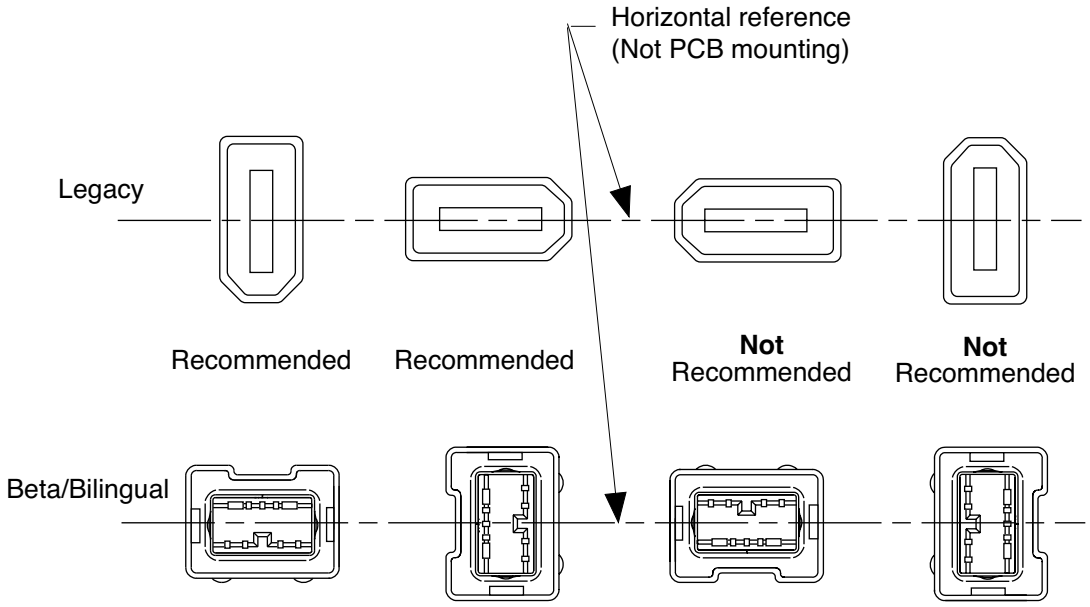

Figure 2-2—Socket orientation (as viewed from outside of bulkhead)

... or as viewed from the front of a hypothetical product:

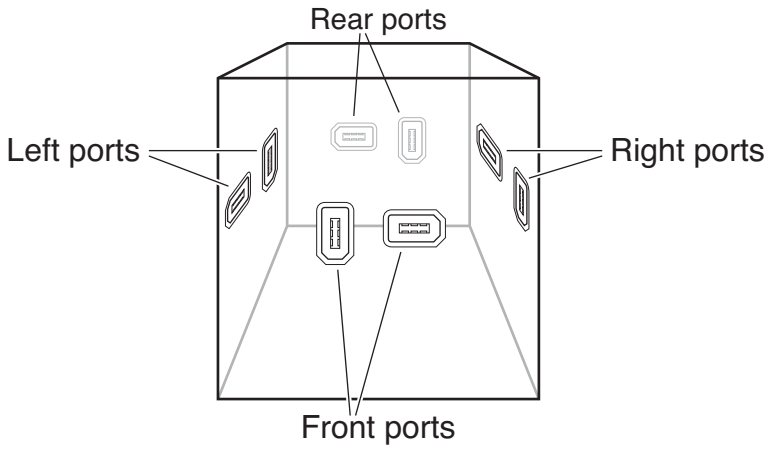

Figure 2-3—Legacy socket orientation (as viewed from front of product)

Item #2.4: Sockets attached to a single PHY should be close together.

Connectors, PHY, protection and termination components for a single port are best seen as a lumped circuit. The min. rise time for 1394a is 0.5 ns which implies 1 GHz waveforms, while the min. rise time for 1394b is 0.080 ns, over 6 GHz!

If both front and back panel sockets are desired, then there are two choices:

a) separate PHYs for front and back (two PHYs and one Link), or
b) a very carefully designed “remote” socket where the entire path from the PHY termination network (see clause 3.2.2) to the socket is within 1394 requirements (110 ± 6 Ω differential mode and 33 ± 6 Ω common mode, more details in [2]). *Note that this is quite unlikely to be successful for a S800 or faster port.*

2.2 End-to-end connections

Item #2.5: FireWire 800 9-to-9 (1394b type 1) cable assemblies shields and ground shall not be shorted together.

Since the new 1394b connections are unfamiliar with designers, the end-to-end connections for the various 1394b cables are described in figure 2-4, table 2-2, figure 2-5, table 2-3, figure 2-6 and table 2-4.

NOTE—The legacy interface cables (type 2 9-to-6 and type 1 9-to-4) do short some of the shields and ground together. It is important that this be done correctly for each cable type.

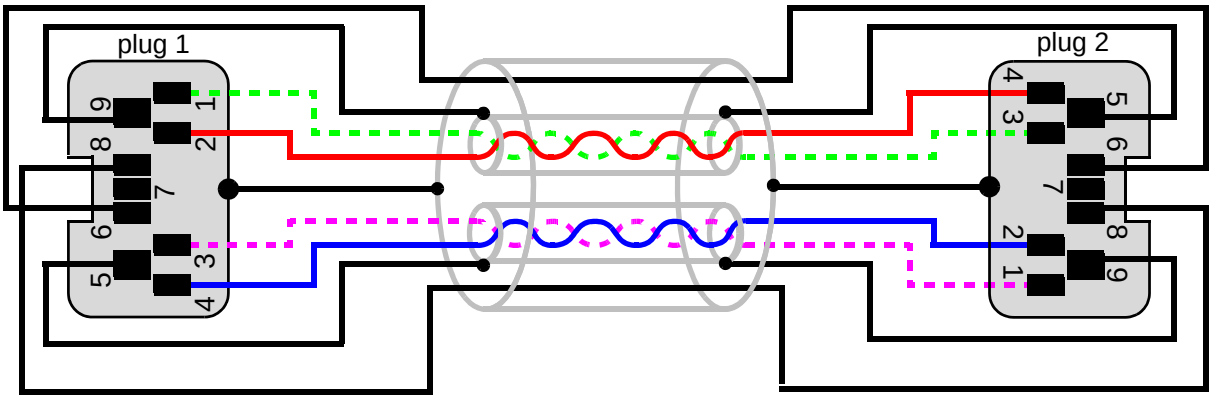

NOTE—Cable as defined in 1394b. Connectors are viewed as looking at the front plug face.

Figure 2-4—1394b type 1 cable assembly and schematic (Beta plug to Beta plug)

Table 2-2—1394b type 1 (Beta to beta) end-to-end connections

Signal	PCB pad	Socket/ plug name	Socket/ plug connection	Cable	Socket/ plug connection	Socket/ plug name	PCB pad	Signal
Chassis Ground	13	Chassis Ground	Outer shell	(no connect)	Outer shell	Chassis Ground	12	Chassis Ground
Chassis Ground, HF to Logic Ground	11	Cable Shield Ground	Inner shell	Outer shield	Inner shell	Cable Shield Ground	10	Chassis Ground, HF to Logic Ground
TPA	4	TPA	4	Signal pair #1 red	2	TPB	2	TPB
High frequency to Logic Ground	5	TPA(R)	5	Signal pair #1 shield	9	TPB(R)	9	Logic Ground
TPA*	3	TPA*	3	Signal pair #1 green	1	TPB*	1	TPB*
Logic Ground	6	VG	6	Power pair #1 white	6	VG	6	Logic Ground
(no connect)	7	SC	7	(no connect)	7	SC	7	(no connect)
FW PWR	8	VP	8	Power pair #1 black	8	VP	8	FW PWR
TPB	2	TPB	2	Signal pair #2 blue	4	TPA	4	TPA
Logic Ground	9	TPB(R)	9	Signal pair #2 shield	5	TPA(R)	5	High frequency to Logic Ground
TPB*	1	TPB*	1	Signal pair #2 orange	3	TPA*	3	TPA*
Chassis Ground, HF to Logic Ground	10	Cable Shield Ground	Inner shell	Outer shield	Inner shell	Cable Shield Ground	11	Chassis Ground, HF to Logic Ground
Chassis Ground	12	Chassis Ground	Outer shell	(no connect)	Outer shell	Chassis Ground	13	Chassis Ground

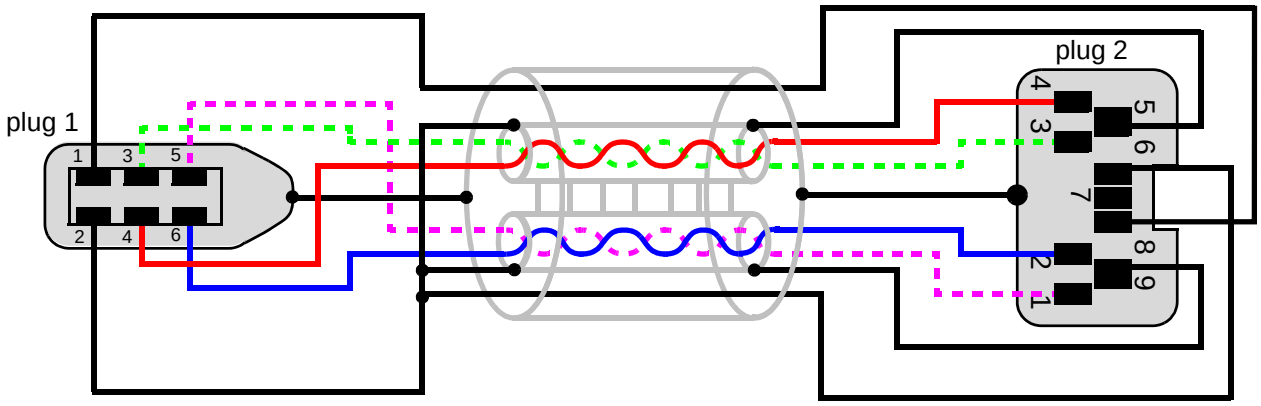

NOTE—Cable (reference) IEEE Std 1394-1995. Connectors are viewed as looking at the front plug face.

Figure 2-5—1394b type 2 cable assembly and schematic (Legacy 6 circuit plug to Bilingual plug)

Table 2-3—Legacy 6 circuit to 1394b Bilingual end-to-end connections

Signal	PCB pad	Socket/ plug name	Socket/ plug connection	Cable	Socket/ plug connection	Socket/ plug name	PCB pad	Signal
				(no connect)	Outer shell	Chassis Ground	12	Chassis Ground
Chassis Ground, HF to Logic Ground		Cable Shield Ground	Plug shell	Outer shield	Inner shell	Cable Shield Ground	10	Chassis Ground, HF to Logic Ground
TPA	6	TPA	6	Signal pair #1 red	2	TPB	2	TPB
Logic Ground	2	VG	2	Signal pair #1 shield	9	TPB(R)	9	Logic Ground
TPA*	5	TPA*	5	Signal pair #1 green	1	TPB*	1	TPB*
Logic Ground	2	VG	2	Power pair #1 white	6	VG	6	Logic Ground
				(no connect)	7	SC	7	(no connect)
FW PWR	1	VP	1	Power pair #1 black	8	VP	8	FW PWR
TPB	4	TPB	4	Signal pair #2 blue	4	TPA	4	TPA
Logic Ground	2	VG	2	Signal pair #2 shield	5	TPA(R)	5	High frequency to Logic Ground
TPB*	3	TPB*	3	Signal pair #2 orange	3	TPA*	3	TPA*
Chassis Ground, HF to Logic Ground		Cable Shield Ground	Plug shell	Outer shield	Inner shell	Cable Shield Ground	11	Chassis Ground, HF to Logic Ground
				(no connect)	Outer shell	Chassis Ground	13	Chassis Ground

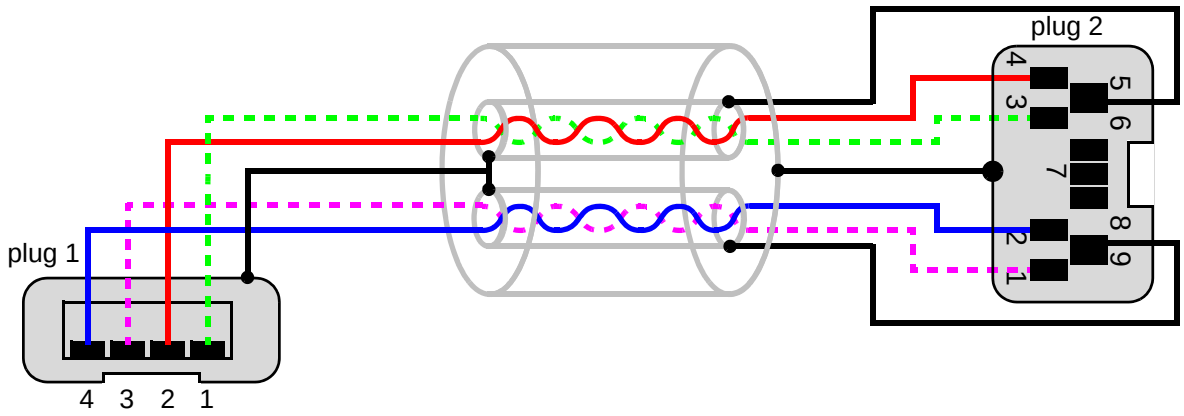

NOTE—Cable (reference) IEEE Std 1394a-2000. Connectors are viewed as looking at the front plug face.

Figure 2-6—1934b type 3 cable assembly and schematic (Legacy 4-circuit plug to Bilingual plug)

Table 2-4—Legacy 4 circuit to 1394b Bilingual end-to-end connections

Signal	PCB pad	Socket/ plug name	Socket/ plug connection	Cable	Socket/ plug connection	Socket/ plug name	PCB pad	Signal
				(no connect)	Outer shell	Chassis Ground	12	Chassis Ground
			(no connect)	Outer shield	Inner shell	Cable Shield Ground	10	Chassis Ground, HF to logic Ground
TPA	4	TPA	4	Signal pair #1 red	2	TPB	2	TPB
Logic Ground		Shell	Plug shell	Signal pair #1 shield	9	TPB(R)	9	Logic Ground
TPA*	3	TPA*	3	Signal pair #1 green	1	TPB*	1	TPB*
				(no connect)	6	VG	6	Logic Ground
				(no connect)	7	SC	7	(no connect)
				(no connect)	8	VP	8	FW PWR
TPB	2	TPB	2	Signal pair #2 blue	4	TPA	4	TPA
Logic Ground		Shell	Plug shell	Signal pair #2 shield	5	TPA(R)	5	High frequency to Logic Ground
TPB*	1	TPB*	1	Signal pair #2 orange	3	TPA*	3	TPA*
			(no connect)	Outer shield	Inner shell	Cable Shield Ground	11	Chassis Ground, HF to logic Ground
				(no connect)	Outer shell	Chassis Ground	13	Chassis Ground

3. Port design

There are a number of good 1394 port design application notes. The TI notes on EMI [11] and layout [12] as well as the LSI design guide for the FW323 [13] are particularly good. Every designer of FireWire systems should be familiar with these documents. Much of the following is derived from those application notes.

3.1 Interfaces

All interface specifications apply at the points of entry and exit from the equipment. The interface specifications may be "valid" at other places. These points are identified as points TP2 and TP3 as shown in Figure 3-1. The specifications assume that all measurements are made after a mated connector pair, relative to the source or destination. TP1 and TP4 are reference points for use by implementors to specify vendor components. In particular, PHY IC specs apply at TP1 for transmit and TP4 for receive.

The reference points for all connections are those points TP2 and TP3 at the transitions between the cabinet and the cable shield. If sections of transmission line exist within the cabinet shield, they are considered to be part of the associated transmit or receive network, and not part of the cable plant.

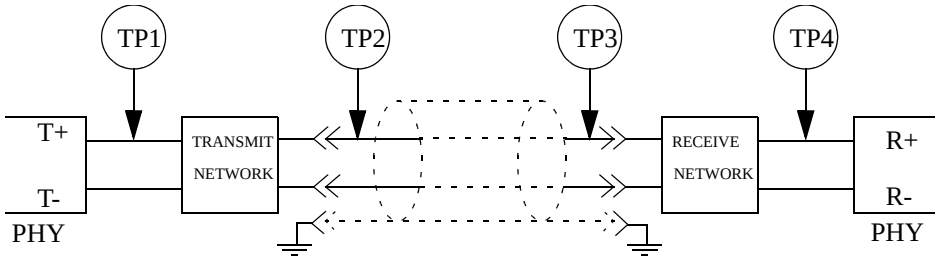

Figure 3-1—Measurement points (half connection is shown)

NOTE—Do not confuse "TPn" where "n" is the number one through four with the "TPx" where "x" is "A" or "B". TPn refers to a test point and TPx refers to a twisted pair signal.

3.2 Connector/PHY wiring

Item #3.1: The wiring between the connector and the PHY (TP1 to TP2 and TP3 to TP4) shall be as short as possible.

The TPA and TPB pairs carry very high speed signals, and so must be imlemented using good design practices for high speed circuits and interfaces to external devices. The impedance requirements for the connector/PHY wiring shall meet 1394b requirements (included below for reference).

Table 3-1—1394b receiver characteristics

Parameter	S400β	S800β	S1600β	Units
Input impedance test conditions:				
TDR rise time	100	100	50	ps
Exception window[a]	700	700	700	ps
Input impedance @ TP3:				
Through connection[b]	110 ± 20	110 ± 20	110 ± 20	Ω
At termination[c]	110 ± 10	110 ± 10	110 ± 10	Ω

Table 3-1—1394b receiver characteristics

Parameter	S400β	S800β	S1600β	Units
Differential skew	5%	5%	5%	UI
Common mode input impedance	> 550			Ω

[a.] Within the Exception-window no single impedance excursion shall exceed the Through-connection impedance tolerance for a period of twice the TDR rise time specification.

[b.] Through connection impedance describes the impedance tolerance through a mated connector. This tolerance is greater than the termination or cable impedance due to limits in the technology of the connectors.

[c.] The input impedance at TP3, for the termination, shall be recorded 4.0ns following the electrical reference plane determined by the receptacle on the receiver bulkhead

NOTE—1394 bilingual and FW400 ports have transmitter and receiver on both the TPA and TPB signals, so the receiver input impedance spec applies for both pairs.

3.2.0.1 Termination

The TPA pair has the required 110 ¾ differential termination as close as possible to the PHY, with the center point of the termination attached to the PHY's bias output for that port and capacitively to signal ground (each port on the PHY shall have an independent bias output – this is required to prevent a short on one port from causing all the other ports to not operate).

Item #3.2: Ensure correct value for TPBIAS decoupling capacitor.

The PHY's tpbias output should be decoupled with a 0.33 μF capacitor to ground, unless the PHY is a TI PHY. TI PHYs need a 1.0 μF to ground, so as to guarantee a minimum ripple on the tpbias voltage under a worst-case speed-signaling scenario, which is necessary in the TI design to maintain stability of the tpbias driver. The lower value is needed when using non-TI PHYs in order to meet the 1394a timing specification for tpbias assertion and deassertion.

3.2.0.2 TP EMC/EMI protection

Item #3.3: Limit FW400 common mode choke to the minimum necessary to pass EMC

FW400 ports may need a high speed common mode choke, particularly if the PHY-to-connector layout is long or runs near possible signal sources. Individual components must be used for each signal pair to limit crosstalk between TPA and TPB. The choke must allow the lower fequency common mode speed signals to get through (while removing high frequency common mode signals). In general, this means that the common mode impendence must be less than 165 Ohms at 100 MHz, and the differential impedence less than 15 ohm at 100 MHz.

Products should use the minimum choke necessary to pass EMC. The design should anticipate evaluation of EMC compliance with several values (for example, the Murata DLP11SN range starts at 67 Ohms, with variants at 90 Ohms, 120 Ohms and 160 Ohms), and with not using a common mode choke at all. This can be achieved by designing the layout to allow a common mode choke or a 0 Ohm resistor.

3.2.0.3 Transient protection

Item #3.4: Implement a transient protection circuit in a power provider to protect the PHY from late VG events.

Power providers should have a diode clamping network between the PHY and the socket to provide two functions:

a) ESD protection so that the high voltage is drained to the chassis, and

b) “Late VG” protection (so called because the power provided on VP can appear on the TPx signals if the VG connection takes place after the VP and at least one of the TPx connections is completed). Note that this last function requires that chassis ground and signal ground have a low resistance connection somewhere (see Item #3.8: below).

The high side clamp diodes use a 3.6V nominal zener diode prebiased with a 3.3V source via a 330 Ohm current liminting resistor (other sources may be used with an appropriate resistor value). This provides a nominal 2.4V, worst case 2.1V ESD and late-VG rail to the BAV99 diodes. In turn, allowing for a 0.6V drop through the diodes, this will permit FireWire signaling at up to 2.7V without compromising signal integrity. A late-VG event on the signal pairs will be drained back to ground via the Zener diode. Late VG events can be of a long duration, and they represent a return path for the VP supply, so the current must be directed back to VG. Despite its rating, the recommended Zener has been measured as being capable of sustaining a current of 300mA, up to 2W, keeping the voltage level on the signal pins to less than 5.5V, which the PHY is also able to tolerate. Under more stress, either the PHY fails or the Zener. The Zener tends to fail short, which results in the TPA/TPB pairs being held to GND and preventing FireWire operation. Lower voltage Zeners, which have been used in some designs, tend to fail before the PHY, increasing the system return rate unnecessarily.

ESD events are of a short duration and will be drained from digital ground to chassis ground and from there to earth ground. It is acceptable to use signal ground instead of chassis ground for the transient protection network since signal ground is capacitively connected to chassis ground at the connector. This can only be done if there is no ferrite on the VG path (see Item #3.6: below). The galvanic connection between chassis ground and signal ground must be low impedance at the frequencies not handled by the capacitive connection at the connector.

If a common mode choke is used, then the transient supression network should be placed on the PHY side of the choke. The advantage of having the diodes on the PHY side is that the common mode choke can be placed closer to the connector for better EMI performance.

Late-VG does not occur on self powered devices.Transient protection iodes must have a very low capacitance (less than 0.5 pF).

3.2.0.4 Power and ground interface (VG/VP)

Item #3.5: FireWire power should be current limited and filtered.

VP should be current limited 1.5 A and EMI filtered to minimize system noise getting out to peripherals. Current limit can be done using a polyfuse, although their slow trip speed requires careful thought put towards the effect of momentary excess current draw. A choke between 50 ohm and 1000 ohm (at 100 MHz) is acceptable in the path for EMI filtering and should be located close to the sockets. A 0.010µf capacitor should connect to chassis ground immediately next to the socket.

Item #3.6: VG must be capacitively connected to chassis ground and directly connected to PHY signal ground with minimal (preferably no) filtering.

VG should AC connect to chassis ground immediately next to the socket using a 0.010µf capacitor. There must not be a significant impedance in the path between VG and the PHY signal ground since VG is the return path for the common mode speed signalling sent out on TPB and received on TPA. If a filter is implemented, it should be designed to pass the speed signal (100-120ns pulse of about 20ma on VG). Usually this means any ferrite in the VG path must be 50 ohm or less at 100 MHz.

3.2.0.5 Shield interface

Item #3.7: The socket shield shall be directly connected to chassis ground for all unisolated ports.

There should be a chassis ground plane underneath the socket for this purpose. The galvanic isolation options defined in 1394a and 1394-1995 should no longer be used.

Item #3.8: The socket shield(s) shall be directly connected to VG at one place in the system.

This is usually achieved by connecting chassis ground to VG at one place in the system. Whether any further connection is required, and if so the optimal placing of this is a matter for system design when considering EMC compliance and noise immunity. In general, a DC connection between the socket shield and VG should not be made at the connector, though an allowance for this for experimental purposes may be made by using a 0 Ohm no-stuff. In addition, all unisolated systems must have the connector shield AC coupled (frequency characteristics TBD) to VG near the connector (see Item #3.6: above). The DC connection between shield and VG is needed for the return path of VP when it sneaks through TPx in the case of a "late VG" connection,

Item #3.9: Systems that require isolated interfaces shall use a beta-only 1394b connection.

The 1394b beta-only connection is a much more robust interface for isolated ports. All use of the 1394a and 1394-1995 "floating PHY" implementation with a galvanically isolated PHY-Link interface is discouraged.

3.2.0.6 Layout guidelines

The TPA/TPB pairs carry both high speed differential and lower speed common mode signals (with VG as the common mode return). Skew between the pairs is important. Some general rules:

a) maximize routing symmetry
b) maximize via transitioning symmetry
c) no 90 degree corners
d) a maximum of two vias in the path for FW400 ports and one via in the path for FW800 ports, and all signals in the port TPA+, TPA-, TPB+ and TPB- must have the same number of vias
e) differential intra-pair skew less than 10ps for the differential trace on the PCB
f) inter-pair skew less than 25ps on the PCB for ports supporting FW400
g) minimize distance between connector and PHY (system-dependent, but 25 mm is a good choice for the maximum for a 1394a port and 12mm for a bi-lingual port as these are roughly the distance traveled in the minimum rise time for the appropriate signals).
h) termination as near PHY as possible (< 10 mm, closer is better)
i) minimize the stub length for termination resistors and diodes
j) provide a semi-isolated ground plane for the 1394 port signals (connect with system ground plane only at the PHY)
k) no other signals should be routed across or near the TPA/TPB pairs, particularly not clock lines or **any** signals or ground/power planes associated with switching power supplies.
l) the digital ground plane should extend to the point where the digital signals are connected to the connector. A chassis ground plane should be provided under the connector, but should not extend into the board area used for the digital signal routes.
m) read and learn about high speed circuit design!

3.2.1 Termination

Same as FW400, see clause 3.2.0.1. Layout is critical.

3.2.1.1 TP EMC/EMI protection

Item #3.10: FW800 ports should avoid common mode chokes.

FW800 PHYs are extremely sensitive to signal distortion caused by the extra impedence and capacitance added by common-mode chokes. Fortunately, FW800 PHYs also generate an exceeding clean differential signal, so a choke is not needed to eliminate residual common-mode currents. This places requirements on the designer to be exceedingly careful that the TPA/B pairs do not add common mode noise (very close in length, isolated from noise sources, etc.). The 1394b connector also provides better emissions control when transmitting 1394a signaling.

Item #3.11: FW800 connector socket should internally bond the inner and outer shells.

The need for a common mode choke may be avoided by ensuring that the connector socket internally bonds the inner and outer shells together. This is appropriate for all designs that do not use galvanic PHY/Link isolation. (Note that the inner shell on the connector plug is connected to the outer shield of the cable.)

If previous experience with similar designs indicates that a common mode choke may be required, then this should be anticipated by designing the layout to allow a common mode choke or a 0 Ohm resistor. If a common mode choke proves to be necessary, then one designed for either DVI or S800 1394b is required. A suitable device is the TDK ACM2012H-900-2P.

3.2.1.2 Shield interface

Item #3.12: FW800 TPA and TPB shields have different termination requirements than each other and for FW400 shields.

The individual pair shields have their own connections in the 9-pin socket. The shield for TPA must connect to digital ground via a high value resistor (1 MΩ or greater) in parallel with a 0.1 μF capacitor and to the chassis ground via a 0.001 μF capacitor. The shield for TPB must connect directly to digital ground and to chassis ground via a 0.001 μF capacitor. Same as FW400 for VP and VG, see clause 3.2.0.5.

3.2.2 Front panel (“remote”) connector wiring guide.

Item #3.13: Ensure signal integrity of long traces to support front panel connectors.

When a front panel connection is required to be supported by a PHY primarily intended for real panel connections, then the traces on the circuit board must be effectively shielded from the noise produced by the rest of the system and the lengths very carefully controlled (see the layout guide above.). In addition, if the signal pairs travel farther than a rise time equivalent (0.5ns or about 25mm), the 110±6 Ω impedance must be maintained for the entire length. Note that additional design, layout, and test time may be necessary to ensure that the extra EMC/EMI, signal degradation, and impedance matching requirements are satisfied.

Front panel connections using a single PHY that primarily serves back panel connections are discouraged for several reasons:

a) It is much easier to get a reliable system if the connector-to-PHY distance is minimized.
b) The cable length budget calculated for cable assemblies is always done assuming there is no significant length adder inside the system (TP1-TP2 and TP3-TP4). This can affect total signal attenuation budgets.
c) A front panel connection is only useful for a system that is always right in front of the user. A desktop type system could provide this, but a tower system is typically under a desk. A better solution for a tower system would be to include a FireWire hub in the package and allow the user to place the hub wherever it is most useful.

When front panel access is required then the following possible solutions should also be considered:

a) Compromise on the position of the connector. Perhaps move all the FireWire sockets to the side (like old iMacs) or take advantage of the size of the system to make backside connectors easily accessible (like the new iMacs).

b) Consider adding a second PHY to the design that can be placed directly next to the front panel connector. The routing of the TPA/B pairs between the back and front PHYs is less critical since it will be short and terminated at both ends. In addition, there are no special consideration for external cable design since almost all of the attenuation budget can be used by the cable assembly. ***One important note:***

Item #3.14: If a system has multiple PHYs and at least one of the PHYs has S800 or faster ports, then all the PHYs should be S800 or faster.

This is because a system that ships with mixed S800 (or faster) and FW400 PHYs will always be a “hybrid” bus and will not be able to run in the more efficient “pure beta” mode. Even if a PHY only has FW400 external connectors, it should be an S800 (or faster) PHY with an S800β (or faster) connection to the other system PHY(s). The user of such a system will then always have the advantages of pure beta bus as long as nothing is connected to the FW400 port(s).

This page left intentionally blank.

4. System design

4.1 PHY selection

All PHYs should be reevaluated with the vendors at each design cycle. In particular, the following checklist should be used:

a) Determine the **actual** parameters of the PHY (measured at TP1 and TP4 in figure 3-1). They should be better than the 1394 specifications, since the 1394 specifications are measured **at the connector** (TP2 and TP3 in figure 3-1), and your design will add extra circuit elements between the PHY and the connector pins which will degrade the signal.
b) Protection features outlined previously are best done if incorporated in the PHY itself. Encourage PHY vendors to do this.
c) The PHY/Link interface (and any other bidirectional or output pins that are potentially connected to CMOS inputs) should be pulled down by a high resistance path. This is to avoid excess current draw by the CMOS receiver on the other end of the signal when the PHY is powered down, or when the interface is disabled by the PD signal.
d) more TBD.

4.2 PHY configuration recommendations

4.2.1 General

Item #4.1: Ensure that Power_class is set correctly.

Item #4.2: Ensure CONTENDER is deasserted

It is important that the node does not contend to be Isochronous Resource Manager until and unless software is ready. The PHY's contender pin must always be deasserted.

4.3 Link layer recommendations

4.3.1 PCIe recommendations

Item #4.3: Connect PCIe CLKREQ*

CLKREQ* should be connected to the enable of the clock chip associated with the PCIe REFCLK+/-. The clock chip in turn should be programmed either to ignore the enable (and always provide the clock) or programmed to accept the active-low request depending on system power management. Typically this signal is open drain and requires a 10K pull-up.

4.4 Specific device recommendations

4.4.1 TI TSB81BA3

See clause 5. for details on power class determination and implementation.

Item #4.4: TI PHY core voltage is 1.95V

The nominal 1.8V power rails DVDD-1.8 and PLLVDD-1.8 require 1.95V +-0.1V. If a filtering resistor is used, then a filtering resistor of 1 Ohm is recommended. Due to the voltage drop through the filtering resistor, the 1.95V LDO should target a voltage of 1.96V.

Item #4.5: TI PHY core voltage filtering uses 1uF caps

TI recommends one 1uF capacitor on each DVDD_CORE power terminal places as close to the device as possible (directly underneath if possible) to provide filtering.

Item #4.6: TI PHY oscillator voltage can be 1.95V

Although the XI input is documented as a 1.8V input, TI confirm that it is acceptable to run this as a 1.95V input, powered from the same rail as the core voltage.

Item #4.7: Ensure correct PHY/LINK pin termination when using TI PHYs in repeater mode

If a TI PHY is used in repeater mode, i.e. a mode in which it does not have an attached link layer (for example, on a front panel), then the PHY link interface pins should be terminated as follows:-

Unconnected: PINT, CTRL0, CTRL1, D0:D7, PCLK (1394b PHY), SCLK (1394a PHY)

Pulled low through a 1K resistor: LPS

Pulled directly to GND: LREQ, LCLK (1394b PHY)

Pulled high through a 1K resistor: BMODE (1394b PHY)

Item #4.8: Ensure configuration and termination of unused port(s) on TI PHYs

An unused port 0 or port 1 should be configured into DS-only mode by tying DS0 or DS1 respectively to VDD via 1K Ohm resistor. This prevents toning and saves power.

An unused port 2 should be configured into DS-only mode by tying LKON/DS2 to VDD via 470 Ohm resistor. This prevents toning and saves power.

The TPA+ and TPA- pins should be left as no-connects (tying them to GND can cause a false connection to be detected). The TPB+ and TPB- should be tied together and tied directly to GND. The TPBIAS terminal should be left as a no-connect.

Note: If a design allows for a connector as an option (i.e. in product variants with the same PCB layout), then TPBIAS can be connected to GND via the normal 1uF filtering capacitor (this would normally be a stuff option) and the TPB+ and TPB- pins connected via the normal 55 Ohms to a common mode point and thence via a 5K Ohms to GND (these would normally be stuff options, with 0 Ohm stuff options to connect TPB+/- directly to GND when the port is not used).

Item #4.9: Ensure LKON/DS2 is correctly biased high or low

If port 2 is to operate in DS mode or is unused then LKON/DS2 should be tied high via a 470 Ohm resistor and a series 1K Ohm resistor be used to connect to the link's PHY_LINKON pin. This is required to ensure correct biasing of the configuration setting if the link chip is unpowered (the internal ESD protection circuitry in the link acts as a pull-down), and correct signaling of LinkOn to the link. If port 2 is to operate in bi-lingual mode then LKON/DS2 should be tied low via a 1K Ohm resistor and directly connected to the link's PHY_LINKON pin (no series termination).

Note that this updates earlier recommendations to use other values for pull-up/down resistors or series termination resistors.

Item #4.10: Caution when TI 1394b PHY port 2 is unused or set to DS-only mode

Port 2 on a TI 1394b PHY is set to DS only mode by tying LKON/DS2 high via a 470 Ohm resistor. During power reset (normally this occurs only when power is first applied to the PHY) this will cause the link layer to see LKON high. In turn this may cause a software interrupt. This requires the link layer to be active before the PHY layer completes reset, and is an unlikely scenario, but is noted for reference in the case of initialization issues in new system designs.

Item #4.11: Note TI PHY max port speed is reported incorrectly for DS-mode only ports

If a port is forced to DS-only mode then the TI PHY correctly limits connections to S400 (DS), However, register 10 in the PHY register map for the port reports max_port_speed as 3 (S800) rather than 7 (DS-only).

Test and production software needs to work around this. This issue is fixed in Rev D.

Item #4.12: Use peaking inductors on bi-lingual ports on TI PHYs

Peaking inductors are recommended for TI PHY ports that are connected to bilingual connectors to open the transmit eye when operating in Beta mode. Peaking inductors are strongly recommended for signal integrity with PHYs included in MCM packages, and are suggested for all PHYs. A value of 18 nH is recommended, and the inductors should be in series with the 55 Ohm termination resistors (between the resistor and the common mode bias / termination point).

Item #4.13: Tie PLLVDD_33 to the AVDD_3_3 power rail on TI PHYs.

They are tied together inside the device.

4.4.2 TI XIO2213 (Cheetah Express, aka Chex)

Item #4.14: Ensure that TI PHY recommendations are met

The TI XIO2213 incorporates a TI PHY. All the recommendations for the TI PHY apply to Chex.

Item #4.15: Connect Chex OHCI_PME# to a GPIO for optimal power management

By connecting this signal to a GPIO (for example, on a Southbsridge), a software interrupt can be given on a PHY event (such as a new connection) even if Chex and its PCIe connection are in a low power mode with clocks off. Software can then restore full power operation and take appropriate action. This signal has an open drain driver and requires a 10K pull-up.

Item #4.16: Ensure that Chex GRST# asserts only on link layer power cycle.

This is a platform reset that should be applied only when the link layer power rails are power cycled (and, in particular, only shortly before the time when EFI will initialize the device). In particular, it should not be applied over a sleep/wake cycle.

Item #4.17: Leave Chex GPIO's as no-connects

This applies to GPIO0-GPIO7.

Item #4.18: Bring Chex CYCLEOUT to a test point

CYCLEOUT is a valuable signal for scope triggering when carrying out diagnostic investigations. Bring to a convenient test point for scope triggering with a 47K pull-down.

Item #4.19: Connect VDD_33_AUX to 3.3V

VAUX isn't supported from D3cold and the AUX power detect bit is hardwired to 0. However, connecting VDD_33_AUX to 3.3V provides minimum power consumption.

4.4.3 Texas Instruments TSB82AA2 1394b link device

4.4.3.1 Texas Instruments Link implementation requirements

Item #4.20: Ensure TI Link devices implementation requirements are met

Use or otherwise of the internal 1.8V regulator makes negligable difference to overall power consumption. If it is not used, then 1.8V should be provided on the REG18 terminal and the REG_EN* terminal tied high via a 1K Ohm resistor. An external PCI/OHCI configuration ROM is not required, and the SCL and SDA pins should be tied to ground via 220 Ohm resistors. PCI_RST* should be connected directly to PCI_RST_L, it does not need to be gated by the SMC.

4.4.3.2 Sleep mode power required

Item #4.21: Ensure power is provided during sleep mode for TI Link devices.

The TI link family of devices do not preserve the GUID setting if power is removed from the device. It is important for software that the GUID is set once on system power reset, and then never needs to be updated, otherwise a security hole is opened up. Consequently it is necessary to preserve power to the device during sleep, even for single port designs. Equally G_RST* must be used once when AC power is first supplied and then never again.

Item #4.22: Ensure that G_RST* timing requirement is met

The timing requirements on G_RST* require that the signal be asserted (held low) for at least 2ms and that the rise time be less than 4.3ms. In general, this prevents the use of a RC circuit to the power rail. It should be connected to a Power OK or similar signal that will have a duration of at least 2ms..

Item #4.23: G_RST* is asynchronous

There is a datasheet error with respect to the timing requirements of G_RST* and the PCI clock. G_RST* is an asynchronous signal, and, in contrast to the requirements specified in the datasheet, it may be deasserted before the PCI clock is provided.

4.4.4 LSI FW643

Item #4.24: Ensure VAUX_DETECT is pulled high on FW643

The package pin VAUX_DETECT must be pulled high to allow sticky reset to be a separated from non-sticky reset. Otherwise PERSTN will reset the sticky domain as well which will result in, for example, the GUID being lost. VAUX_DETECT is pulled down by default at the pad, and must be pulled high externally with a 10K (??) pull-up.

Item #4.25: Ensure FW643 power reset is stuffed consistently with regulation of 1.0V supply

The FW643 has an internal Power Up Reset (PUR) cell to insure logic is properly initialized and the crystal oscillator circuit has stabilized on power up. The internal PUR cell has no voltage ramp requirement. The PUR cell monitors the VDD10 ramp and generates an internal reset signal that is used to reset its internal flops and deactivate its internal counter until VDD10 reaches about 40 – 70% of its peak value (rising threshold). The output of the PUR stays low during this time and continues in the low state for the PUR period. The output switches to logic high once that time is reached and the internal counter is also disabled at that point. The front end has a built-in hysteresis and a filter capacitor to reject VDD10 noise once VDD10 is up. The cell also detects the falling ramp of VDD10. When VDD10 falls beyond a certain level (designated by falling threshold), the output of the PUR cell goes LOW.

If the internal regulator controller of the FW643 (and associated external circuitry) is used to generate the 1V power supply, no circuitry is necessary on the FW-RESET_N input. It has an internal pull-up and can be left unconnected.

However, if the VDD10 and VDD33 power supplies are controlled independently, the customer needs to insure that the VDD33 supply has reached its required voltage level (3.3V +/- 10%) before the VDD10 power supply has reached 40% of it's required voltage level to guarantee sufficient time for the crystal oscillator to stabilize. Alternatively, the active low FW_RESET_N input can be asserted until the 3.3V power supply has reached its required voltage level. This will extend the length of the power-up reset in insure that the crystal oscillator is stable.

4.4.5 LSI FW323

4.4.5.1 LSI integrated PHY/Link implementation requirements

Item #4.26: Ensure LSI integrated PHY/Link devices implementation requirements are met

The CNA and LPS outputs are not used and should be not connected. An external OHCI configuration ROM is not required, and the ROM_CLK and ROM_AD pins should be tied to ground. Configuration pins that are asserted (typically CARDBUSN and none or more of PC0, PC1 and PC2) should be tied to VDD via a 10K Ohm resistor. PCI_RST* should be connected directly to PCI_RST_L, it does not need to be gated by the SMC.

4.4.5.2 Sleep mode power required

Item #4.27: Ensure power is provided during sleep mode for LSI integrated PHY/Link devices.

The LSI FW323 family of devices do not preserve the GUID setting if power is removed from the device. It is important for software that the GUID is set once on system power reset, and then never needs to be updated, otherwise a security hole is opened up. Consequently it is necessary to preserve power to the device during sleep, even for single port designs. Equally RESETN must be used once when AC power is first supplied and then never again.

Item #4.28: Ensure PME is uniquely identifiable to software for LSI FW323 v129.

As LinkOn is not pinned out on this device, the only way when the device is in low power mode that software finds out that there is a new connection, or disconnection, or similar significant event, is via PCI PME. It is important that software is able to distinguish this PME interrupt from other PME interrupts, so that it does not needlessly poll the device on irrelevant PME interrupts (this involved waking the PCI clock domain parts of the device, initializing it, reading registers, discovering that there has been no change, and putting the device back to sleep again). This can be accomplished by

connecting the PME interrupt to a dedicated GPIO, or ensuring that the PCI bus is not shared by any other devices and is bridged in such a way that the FW323's PME interrupt is readily distinguishable from PME interrupts from othert devices. A dedicated GPIO is the preferred solution, so that the driver can handle a range of system devices in a similar way.

4.5 PHY/Link interface

The PHY/Link interface shall follow the specifications in 1394a Annex J1 or 1394b Clause 17. If the 1394b PIL/FOP interface is used, then 1394b Clause 18 shall apply.

Series termination (typically 22 ohms) may be beneficial in some designs, but has also been observed to aggravate problems with reflections (see Item #4.30: below).

4.5.1 Additional requirements

Item #4.29: Internal or external pull-downs required on PHY/Link interface

All signals between the Link, PHY and other system components shall be pulled down to signal ground with a high resistance. This pull down must be active even when the Link, PHY (or other system component) is powered down. If this is not done, signals could drift into the CMOS switching region resulting in excess power consumption or unreliable operation. The suggested method is for both the PHY and Link to include internal pull-downs; if this is not done, then external pull-downs are needed.

Item #4.30: Special termination may be needed if PHY/Link propogation delay is > 1 ns

If the PHY/Link interface propogation time is greater than about 1 ns then reflections at the clock cycle when a device hands back the interface to its peer can be seen by this device at the next clock. In particular, the device can mis-interpret either or both of CTRL[1:0] and some or all of DATA[7:0]. This effect has been seen on systems where the PHY/Link interface is > 1.4ns. The effect has also been seen to be worse when 22 ohm series termination is used at the peer device. In addition, when the link hands back the interface to the PHY, the effect also depends on the delay from PClk received at the link to when CTRL and DATA are clocked out at the link interface. The link times when it drives these signals with respect to LClk, but after handing back the interface to the PHY, the link samples the same signals using PClk.

The effect is most likely to occur when the link hands back the interface to the PHY after transmitting a MORE_INFORMATION cycle. The misread can result in the link seeing a phantom RECEIVE cycle, GRANT cycle or STATUS cycle. A phantom RECEIVE cycle will normally be benign (a 1 cycle runt packet). A phantom GRANT cycle could cause the link to transmit another packet while the PHY has control of the interface - causing the PHY to miss some or all of the packet. A phantom STATUS cycle can cause the link to misbehave simply because there may be more than one of D[0:7] set (only one bit is set in a valid status cycle), but in particular cause unfairness (delayed asynchronous packets) or isochronous phase confusion (link delays one cycle before transmitting an isochronous packet). The following table summarizes the possible phantom cycles as seen by the link.

Table 4-1—Phantom cycles caused by MI cycle reflection

D lines	Transmitted MORE_INFO cycle (CTL = 11)	Sampled Status cycle (CTL = 01)	Sampled Grant cycle (CTL = 11)	Sampled RX cycle (CTL = 10)
[0]	Format (Beta = 1)	PH_BUS_RESET_START	Grant Format (Beta = 1)	-
[1]	PH_NEXT_EVEN/PH_NEXT_ODD/ PH_CYCLE_START_REQ	PH_ARB_RESET_ODD	Grant Type = Async/CS/Immediate	-
[2]	PH_ISOCH_REQ_EVEN/ PH_CURRENT/PH_NEXT_ODD	PH_ARB_RESET_EVEN	Grant Type = Isoch/CS/Immediate	-
[3]	PH_ISOCH_REQ_ODD/ PH_CURRENT/PH_NEXT_ODD	PH_ISOCH_ODD	Grant Type = Async/Immediate	-
[4]	EOS	PH_ISOCH_EVEN	-	-
[5]	Speed = S400/S800	PH_SUBACTION_GAP	Grant Speed = S400/S800	-
[6]	Speed = S200/S800	-	Grant Speed = S200/S800	-
[7]	-	-	-	-

4.5.2 Layout guide

Care must be taken to avoid crosstalk between PCI signals and FireWire high speed signals and/or ground-bounce from PCI signaling. Ensure that there is sufficient grounding for the PCI interface, and keep separate from FireWire grounding if possible.

5. Cable power

5.1 Introduction

A major benefit of FireWire is the ability to power (and recharge, for that matter) devices from the bus. When designing a device, it is important to consider whether and when the device will be a power consumer (power itself from the bus, possibly charging its batteries), a power provider (provide power to other devices on the bus) and/or a power repeater (repeat power received on one port to other ports on the device), and whether it might provide more than one of these functions in different power states.

A significant concern on a bus is to ensure that there is sufficient power available for all the power consumers. Every device is allowed to power its PHY from the bus at all times, and indeed consume up to 3W. A device may not consume more power than this until permitted by a power manager. The “LinkOn” PHY command packet is sent from the power manager to a power consumer when the power manager determines that sufficient power is available. A limited amount of status is provided on bus reset to enable a power manager to evaluate power needs and availability.

Currently, power managers are not implemented!

The main reference for power distribution is the 1394 Trade Association Document “Power Specification Part 1: Cable Power Distribution” TA 1999001-1 (October 5th 1999) and the main reference for power management is the 1394 Trade Association Document “Power Specification Part 3: Power Distribution Management” TA1999001-3 (January 15, 2000). However, there are significant differences in the recommendations in this document. These differences are summarized below.

NOTE—a device normally incorporates a single PHY, and appears on the bus as one node. If a device incorporates more than one PHY, then these rules apply separately to each PHY (each node).

5.2 FireWire power checklistrecommendations

Item #5.1: All components and traces between the protection diode and the connector need to support 33V or higher.

External power providers may provide power at up to 30V.

Item #5.2: PHY operation should not be affected by power provider current limiter trips

Designs should be such that the PHY will continue operating if the regulatory current limiter trips. Note that the lack of provision of power will often cause bus resets, and the status of “PS” (Cable Power Status) will change as the current limiter repeatedly trips and resets.

Item #5.3: System operation of power providers should not be affected by power consumer inrush or shorts.

When a FireWire device is connected to a power provider, or suddenly consumer extra power (for example a disk drive starting to spin up) ,the inrush current can cause a dip in the voltage on the power rail used to provide FireWire power. This in turn can affect operation of other parts of the system dependent on the same rail. Some devices have been seen not to respect the inrush limits, and a similar but worse problem can occur if a faulty shorted FireWire device is attached.

Protection against this is normally implemented within the system power architecture. In addition, it may be helpful to include current-limiting series resistors in the FireWire power rail. The port fuse is basically a short for a few seconds after a shorted FW device is attached, and the current limit resistor limits current so that the AC/DC can source current long enough to blow the fuse. Note that these resistors are special, untrimmed surge resistors, so system designers have to be sure to get the right kind of resistors. The need for such resistors, and the appropriate value to use, willl vary according to system design.

Without such resistors, a short between VP and VG will trip the power adapter and possibly bring the system down. However, the initial value chosen for the resistors may well result in the resistors frying before the polyfuse blows. The final solution in this case may be to use relatively low value resistors (0.025 Ohm 1W) with a 0.5A polyfuse. Note that the rating of the polyfuse gives the maximum current at which the fuse is guaranteed not to open. Polyfuses typically will sustain a load of up to 2x their nominal rating, and the resistors need to be able to sustain this load.

Item #5.4: Power Class 4 power providers should implement appropriate CSRs.

Any node declaring itself as Power Class 4 and which provides power (i.e. is an Alternate Power Provider) should implement the Power Management CSRs that indicate its power providing capability. (PowerClass 4 really means “go read the CSRs to find out what I do”). These are defined in [6].

Item #5.5: Bus powered portable devices and peripherals should declare the appropriate power class.

When a device is powered from the bus, it must declare itself as power class 4, 6 or 7 as appropriate (less than 3W total, 7W total or 10W total consumption from the bus). In power classes 6 or 7, it must support LinkOn, and respect the rules for consuming no more than 3W until LinkOn is received. A node which requires more than 10W power, for example to charge its batteries, should either (a) declare itself as Class 4, use less than 3W for PHY, link and management purposes, or (b) use Class 6, use less than 7W for PHY, link and management purposes and support LinkOn. In both cases it should then use Power Management CSRs to declare its requirement for up to 45W and to allow its consumption to be controlled by a Power Manager.

See the examples in clause 5.4.4.

Item #5.6: Power consumers should implement the power management CSRs.

These are defined in [6].

5.3 FireWire device guidelines

PA1) When a device is connected to mains power, it provides power to the bus either

a) as a Primary Power Provider, providing 20W at greater than or equal to 20V (recommended 24-26V) declaring itself as power class 1, or
b) as an Alternate Power provider, providing 8W at 12-15 V, declaring itself as power class 4 (for multiple port devices) or power class 0 (for single port devices), and powers its PHY from the bus if possible.

PA2) When a device is connected to mains power, it provides power to the bus as an Alternate Power provider, providing 8-16W at 12-15V, declaring itself as power class 4 (for multiple port devices) or power class 0 (for single port devices).

Higher power provision (between 20W and 45W) is left to external hubs.

The device provides power to the bus when connected to mains power even if in “sleep” or “off”.

PA3) When a single port device is powered from battery power, it either

a) (portable device) when running provides power to the bus as an Alternate Power Provider (power class 0), providing 8W at a protected battery rail (normally between 9.25 V and 15 V), or
b) (mobile device, never connected to mains power) never provides power to the bus.

PA4) When a multiple port device is powered from battery power, it either

a) (portable device) when running provides power to the bus as an Alternate Power Provider (power class 4), providing 8W at a protected battery rail (normally between 9.25 V and 15 V) and when in sleep or off powers its PHY from the bus if possible, or

b) (mobile device, never connected to mains power) never provides power to the bus.

PA5) When a device is powered from the bus, it must declare itself as power class 4, 6 or 7 as appropriate (less than 3W total, 7W total or 10W total consumption). In power classes 6 or 7, it must support LinkOn, and respect the rules for consuming no more than 3W until LinkOn is received. A node which requires more than 10W power, for example to charge its batteries, should either (a) declare itself as Class 4, use less than 3W for PHY, link and management purposes, or (b) use Class 6, use less than 7W for PHY, link and management purposes and support LinkOn. In both cases it should then use Power Management CSRs to declare its requirement for up to 45W and to allow its consumption to be controlled by a Power Manager.

PA6) (Future requirement) A node which may consume more than 3W from the bus should include Power Management Software. This is to avoid being dependent on non-existent or inadequate power managers! Without power management there is a possibility of a bad user experience.

PA7) At any time that a node changes power class, then it advertises this with a short (arbitrated) bus reset. Care needs to be taken with power provision sequencing (provide power before changing power class to advertise its availability, remove power only after changing power class advertising its lack of availability).

PA8) The protection circuitry for Primary Power Providers shall include per-port diode protection to (a) prevent consumption of bus-supplied voltages and (b) alleviate the problem of being overridden by a higher voltage power provider, and shall include per-port resettable fuses to limit the current drawn through any port. The protection circuitry for all multiple-port power providers shall include per-port resettable fuses to limit the current drawn through any port. Each port shall support but shall not exceed 1.5A. Note that this protection is needed to protect against current summing from one or more ports and/or the internal power provider to the power provided on another port. In addition, current limiters shall be included for all power providers so that the provision of internal power is not affected by a short on any VP wire. See reference circuits below.

PA9) Any node declaring itself as Power Class 4 and which provides power (i.e. is an Alternate Power Provider) should implement the CSRs which indicate its power providing capability. (PowerClass 4 really means “go read the CSRs to find out what I do”).

PA10) Designs should be such that the PHY will continue operating if the regulatory current limiter trips. Note that the lack of provision of power will often cause bus resets, and the status of “PS” (Cable Power Status) will change as the current limiter repeatedly trips and resets.

PA11) A power brick is detected by the presence of power but the absence of a connection on a port. A node may consume as much power as it can extract from a power brick.

PA12) Any multiport device without per-port power protection diodes (typically a Power Class 4 node) must be designed at least to power the PHY from the bus in cases where no local power is available (unplugged, etc.). In no case can power “flow through” a device without data also being repeated. System may be allowed to violate this rule when not connected to mains power where this is considered not a normal mode of operation.

5.4 Example Circuits

This clause shows example circuits (including power provision, PHY Vdd from bus power, and all protection circuits) for a variety of applications. Note that these circuits are only shown as functional examples, they are not exact implementations. In particular, the use of so many diodes in a high current system is discouraged because of the power dissipation due to forward current losses. Switches with a polarity detect might be a better choice.

Annex 10. briefly describes a hypothetical integrated circuit that could be used to implement all the following examples.

Item #5.7: Ensure that CPS is correctly connected

Multiple port systems that may be powered from cable power must take the CPS tap at the point at which the VP connections to the various ports are OR'd together in order to ensure that the presence of cable power is reported correctly. Systems that may in some situations not power the PHY when cable power is present should take this input from the power provision side of the isolating diode in order to prevent pump-up current entering the PHY when it is not powered (although any such current, necessarily via the high value series resistor, will be very small).

Desktop, server or hub primary power provider

A normal stationary device or a hub may be a primary power provider. This allows power to be much more easily managed (see above), and up to 45W output per port, if desired. Power domain isolation is provided by means of the per-port diodes Power is provided to the PHY and to the ports as log as main power is available (while the device is plugged in), even if in "sleep"..

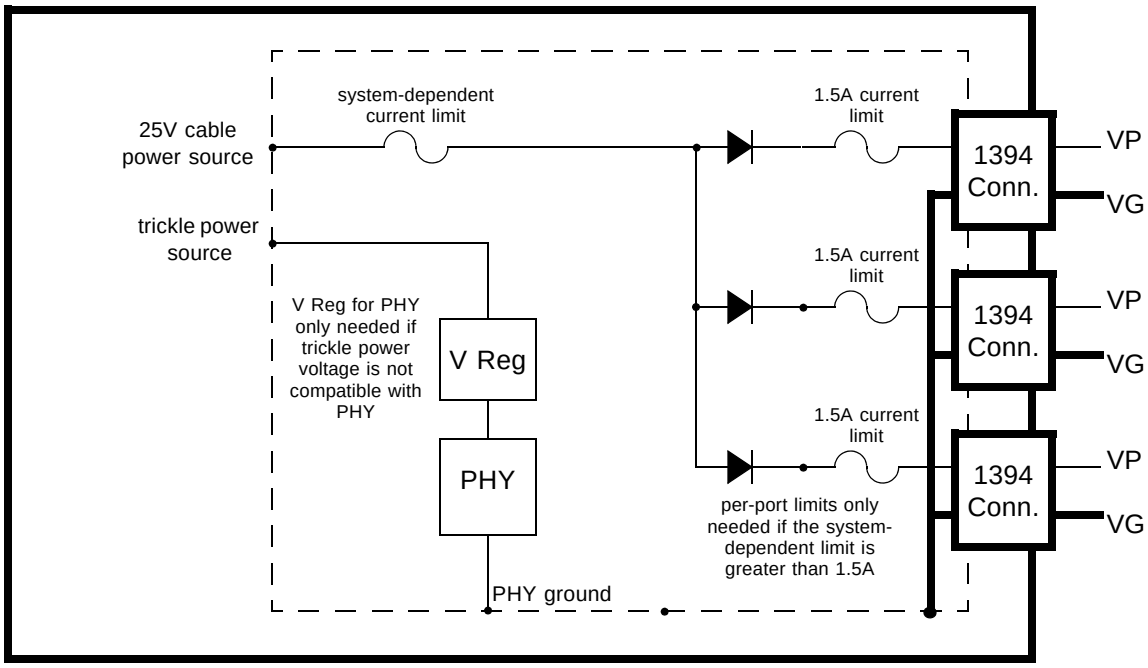

Figure 5-1—A multiport Power Provider class 1/2/3 node

1.5A current limiters are only required on each port if the device can provide more than 1.5A since 1394 requires that no more than 1.5A be consumed via any port (1394/1394a speed signaling fails if more than 1.5A is consumed via a port)..

5.4.1 Desktop alternate power provider

Another way to build a desktop sytsem is as an “alternate power provider” which does not have per-port diode isolation. This device will provide power when plugged in, even if in “sleep” or “off”. Since it does not prevent power from flowing from one port to another, it must also guarantee that the PHY repeats data whenever there is bus power. This means that the PHY must get its power from the bus as shown below in figure 5-2 .

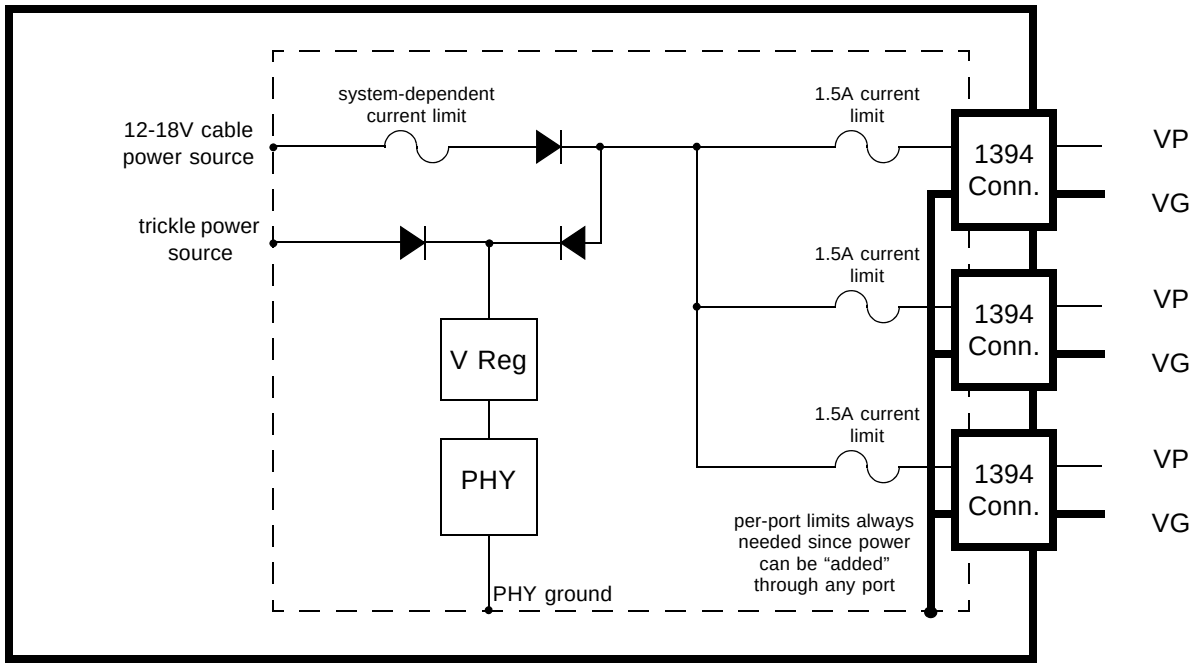

Figure 5-2—A multiport Power Provider class 4 node

When local power is provided, then power is provided to the PHY and to the three ports. There is no power domain isolation: power flows freely between the ports.

When local power is not provided the PHY operates using bus power.

1.5A (min.) current limiters are required on each port because the example device has two or more ports. They provide a measure of regulatory power protection (without them, a faulty device on one port could consume excessive power from power providers upstream on the other port(s) and/or the system). The current limiters must allow at least 1.5A, but also 1394 requires that no more than 1.5A be consumed via any port (1394/1394a speed signaling fails if more than 1.5A is consumed via a port).

A node with this design must ensure that port power pass through is not set for a port in suspend.

5.4.2 Desktop primary power provider with backup power pass-through

It is also possible to combine the two previous designs, allowing a node to be a primary power provider whenever main power is available, yet passing power and data through when main power is turned off.

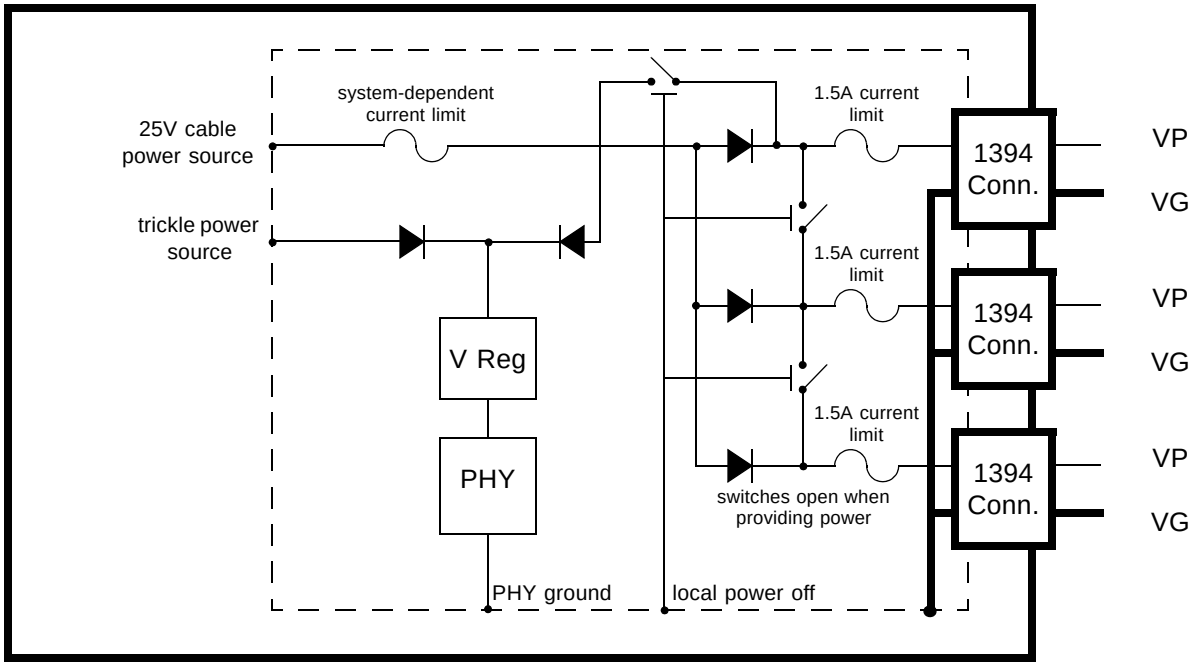

Figure 5-3—A multiport Power Provider class 1/2/3 node, operating as Class 4 when power is not available

This device provides power to 1394 when plugged in, even if in “sleep”.

When local power is provided, then the three switches are held open. Power is provided to the PHY and to the three ports. Power domain isolation is provided by means of the three per-port diodes.

When local power is not provided, then the three switches are held closed. Power is repeated between the ports and the PHY operates using bus power. Note that operation should be phased so that any power consumers connected via any of the ports do not see an interruption in power supply when power is applied or withdrawn locally. The switches are opened after local power becomes available (but before the bus reset declaring power class 1, 2 or 3), and are closed before local power is withdrawn (but after the bus reset declaring power class 4).

1.5A (min.) current limiters are required on each port because the device has three or more ports. They provide a measure of regulatory power protection (without them, a faulty device on one port could consume excessive power from power providers upstream on the other two ports). The current limiters must allow at least 1.5A, but also 1394 requires that no more than 1.5A be consumed via any port (1394/1394a speed signaling fails if more than 1.5A is consumed via a port).

A node with this design must ensure that port power pass through is not set for a port in suspend.

5.4.3 Portable computer

Portable computers will normally be Alternate Power Providers, which are intended for applications with only a limited amount of power..

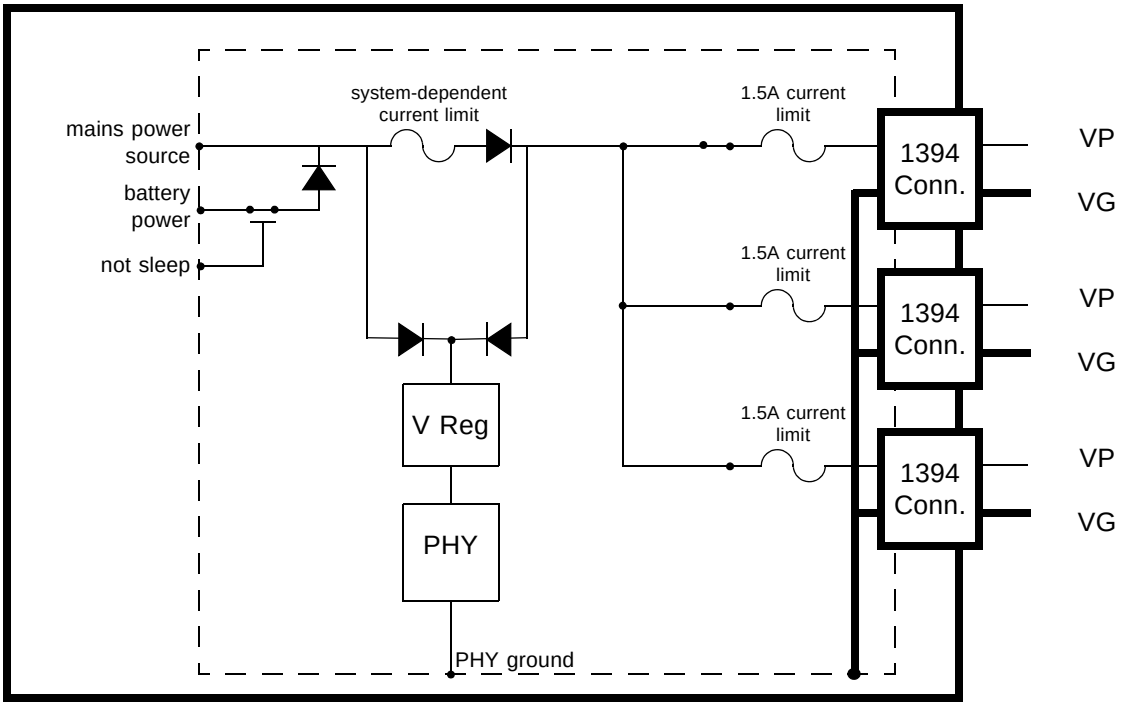

Figure 5-4—A multiport Alternate Power Provider class 4 node

This device provides bus power whenever the device is plugged into the wall, or the device is operating on battery and not sleeping. The PHY is powered from the bus when bus power voltage exceeds system-provided cable power voltage, otherwise it is powered from the local supply.

1.5A (min.) current limiters are required on each port if a device has two or more ports. They provide a measure of regulatory power protection (without them, a faulty device on one port could consume excessive power from power providers upstream on the other port(s) and the internal supply). The current limiters must allow at least 1.5A, but also 1394 requires that no more than 1.5A be consumed via any port (1394/1394a speed signaling fails if more than 1.5A is consumed via a port).

A system-dependent current limit is provided to protect the local supply against excessive demand by external devices and for regulatory purposes.

5.4.4 Peripherals or mobile devices

A single port mobile device may provide power to its PHY using the main supply, if appropriate. This is acceptable because the device does not repeat the FireWire signal

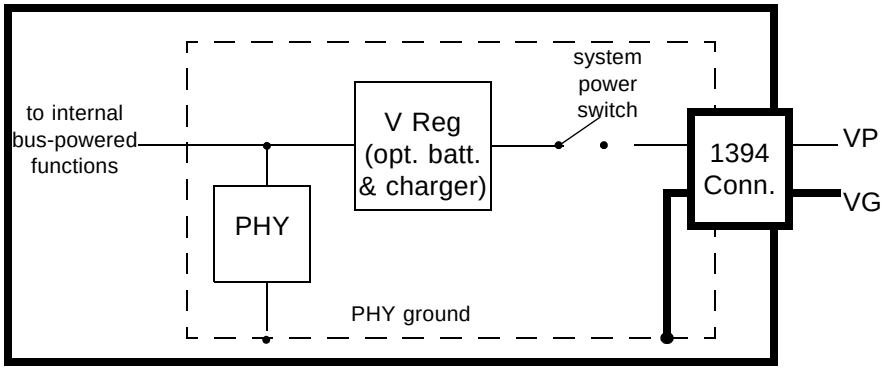

Figure 5-5—A single port bus-powered (class 4) node

A multiport mobile device is required to repeat the FireWire signal, so its PHY should always be powered whenever there is bus power available. It can optionally use power from the main device power supply, if the device itself is always powered when there is bus power.

A mobile bus power device with 3 or more ports must have 1.5A current limiters on each port. They provide a measure of regulatory power protection (without them, a faulty device on one port could consume excessive power from power providers upstream on the other two ports). The current limiters must allow at least 1.5A, but also 1394 requires that no more than 1.5A be consumed via any port (1394/1394a speed signaling fails if more than 1.5A is consumed via a port).

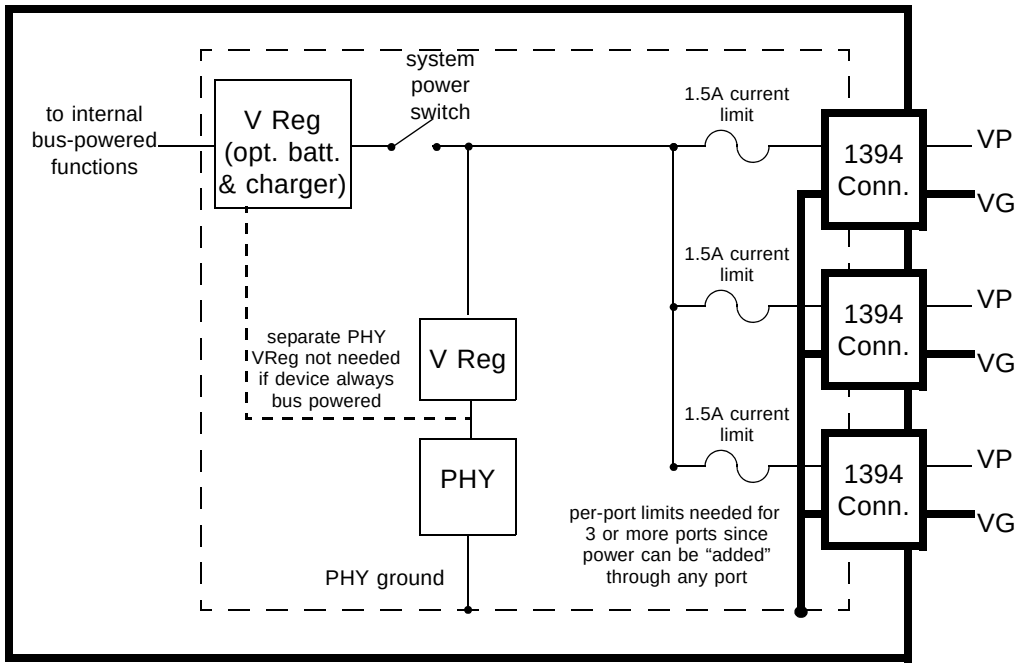

Figure 5-6—A multiport bus-powered (class 4) node

A node with this design cannot use 1394 port-suspend.

5.5 Notes

5.5.1 Power Classes (for reference)

Power consumption and source characteristics are shown in table 5-1:

Table 5-1—Power class

Power class (binary)	Power class	Description
000_2	0	Node does not need power and does not repeat power.
001_2	1	Node is self-powered and provides a minimum of 15 W to the bus.
010_2	2	Node is self-powered and provides a minimum of 30 W to the bus.
011_2	3	Node is self-powered and provides a minimum of 45 W to the bus.
100_2	4	Node may be powered from the bus and is using up to 3 W. No additional power is needed to enable the link[a].
101_2	5	Reserved for future standardization.
110_2	6	Node is powered from the bus and is using up to 3 W. An additional 3 W is needed to enable the link[a].
111_2	7	Node is powered from the bus and is using up to 3 W. An additional 7 W is needed to enable the link[a].

5.5.2 Note on diode protection.

A power provider node uses a diode to protect its internal circuitry from drawing power from the bus when another power providing node is providing power at a higher voltage than the local node. This can be achieved using a single diode on the power supply.

However, a primary power provider is required to implement a per port diode. This assists in the creation of separate power domains, each powered by a single power provider. The result is to reduce the number of scenarios in which the connection of another power provider can disrupt existing devices (see example below in figure 5-7), to reduce the number of scenarios in which an individual cable may be overloaded, and to assist power management software do its job.

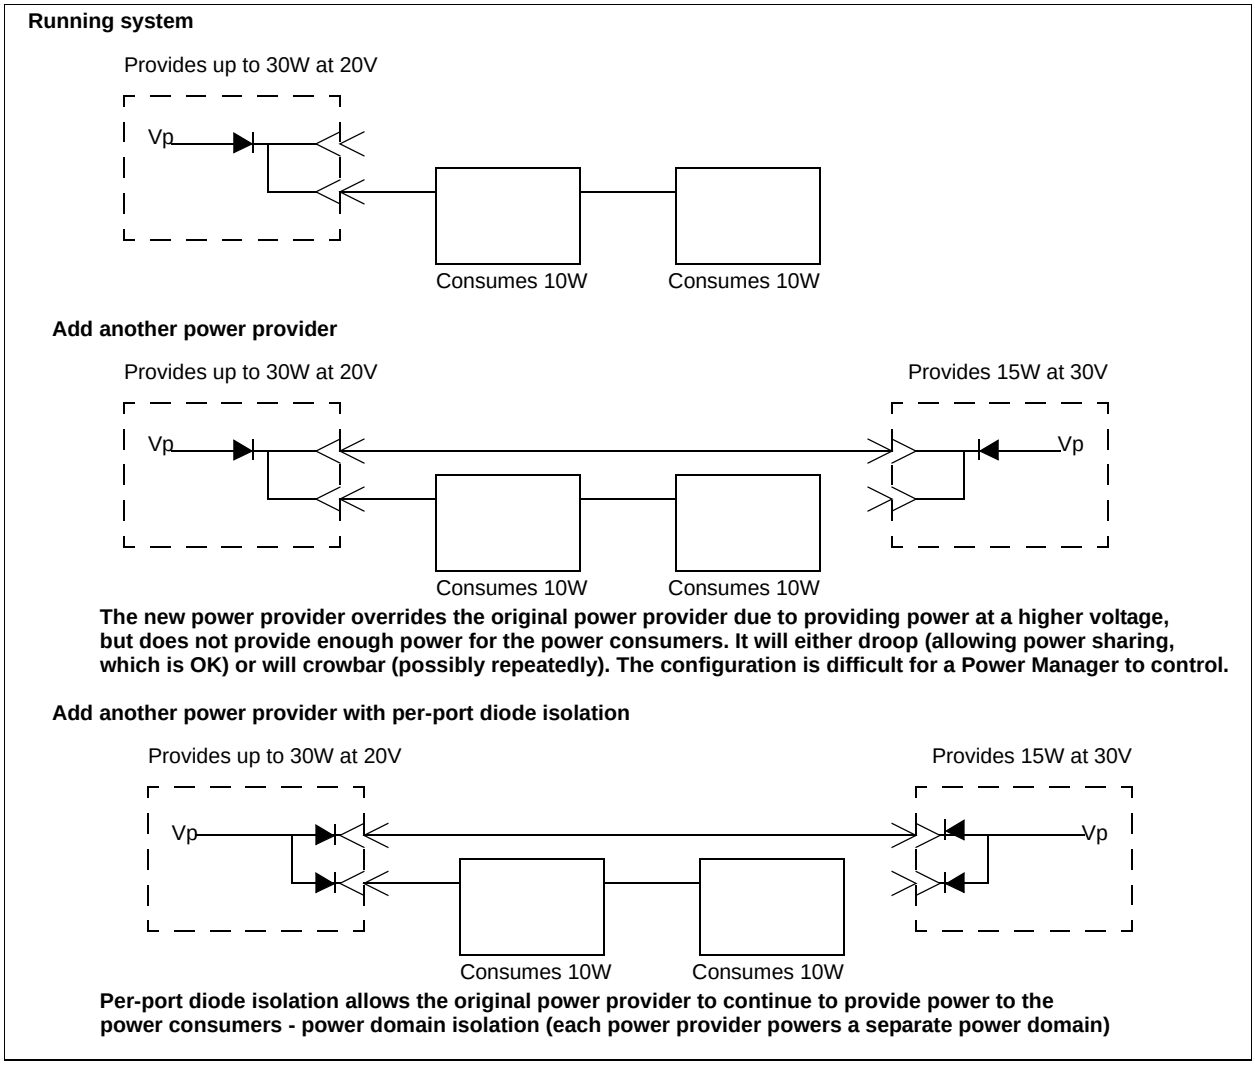

Figure 5-7—Diodes to provide power management domains

Per-port diode protection is required for Primary Power Providers.

5.5.3 Use of Power Down and Cable Not Active

Some devices provide pins on the PHY to report no activity on the cable (CNA - Cable Not Active) and/or to allow Power Down (PD). These may be used in 1394a systems to provide power savings on the PHY.

More TBD

Some devices may use Suspend or Standby if there is no connection/active port to save power.

5.5.4 Trade Association Cable Power Distribution Specification - proposed variations

a) Multi-port power consumers are permitted. They repeat power.

<more observations to be added>

6. Link selection

All host-level devices (CPUs) shall use an Open Host Controller Interface design for the Link layer, devices for embedded applications in computer peripherals or consumer electronics have more freedom.

6.1 OHCI requirements

a) The OHCI for FW400 applications shall follow the OHCI 1.1 specification. See [9]
b) The OHCI for FW800 applications shall follow the OHCI 1.2 specification. See [10]

Item #6.1: OHCI link must reliably meet real-time requirements under anticipated load.

OHCI design must be matched with the encompassing system so that FireWire real-time requirements can be reliably met under anticipated load, such as Final Cut Pro; the OHCI FIFO depth and isochronous transmit packet workahead must be sufficient to cope with worst-case host memory contention, or stated inversely, host memory performance to FireWire must be fast and frequent enough to satisfy the OHCI's workahead and FIFO depth limitations. Stated either way, one must consider DMA descriptor fetch and status overhead in addition to actual payload bandwidth.

7. Firmware/higher layers

7.0.0.1 IRM not root

If an Isochronous Resource Manager discovers that it is not root (and therefore not cyclemaster), then it has some responsibilities to ensure that the bus is configured both correctly and optimally.

Given that there is no incumbent bus manager, it should proceed immediately to reading the Bus_Info_Block of the root.

If the root is cycle master capable then it should immediately set its force root flag and clear the flag on all other nodes (by sending out a PHY configuration packet with the R bit set and the phy_ID of the root), and set the cmstr bit in the STATE_CLEAR register of the root node. This is because the current root may have become root “by chance” with no force_root bit set in any node at the time of the bus reset. It is necessary to make sure that the next time there is a bus reset for any reason, the root stays root (unless the IRM or the BM is explicitly deciding that someone else should be root). Sending out the PHY configuration packet does this. It is the IRM’s responsibility in the absence of a bus manager, to do this. No one else will.

If the root is not cycle master capable, then the IRM have to hunt round for a node that is (The IRM itself is probably a good candidate!) and make that node root (by setting its force_root flag and invoking a new arbitrated bus reset). It is important to get an active cyclemaster going within 125 usec of a bus reset – so it is not practicable to wait for the election of a new bus manager, for example.

The IRM is allowed to detect the presence of root as cyclemaster by the detection of cycle start packets, should the root start sending these quickly. This may save the necessaity of reading the Bus_Info_Block, but the IRM must not wait long to detect the absence of cycle start packets!

7.0.0.2 1394a IRM-capable node

If the node is 1394a IRM-capable (i.e. implements the BROADCAST_CHANNEL register, with channel 31 allocated as the default broadcast channel)? If so, then it should check whether the current IRM (when not itself) has this capability. If it does not, then it should make itself root and IRM.

There are four tests – use one or more of them:-

a) receipt of a write request addressed to your BROADCAST_CHANNEL register that sets the valid bit to 1 (must have come from a 1394a IRM)
b) successful completion of a read request addressed to the IRM’s B_C register with the most significant bit set in the response
c) successful read of the IRM’s bus information block that shows the generation field to be non-zero
d) analysis of self_ID’s shows that the IRM has not changed, and previously was known to be 13954a compliant

See 8.4.2.3 in 1394a (page 169).

This page left intentionally blank.

Annex 8. B reference design

NOTE—As of 10/14/2009, this document includes the latest schematic drawings provided by Texas Instruments.

This FW800 PCI reference design has been contributed by Texas Instruments. The power circuitry and PCI connector portions of the schematic are provided only to show a complete design and may not be applicable to all PCI applications. For more schematic and layout information, see appropriate TI application note.

NOTE—This design needs minor updates. See the main text for correct components, etc.

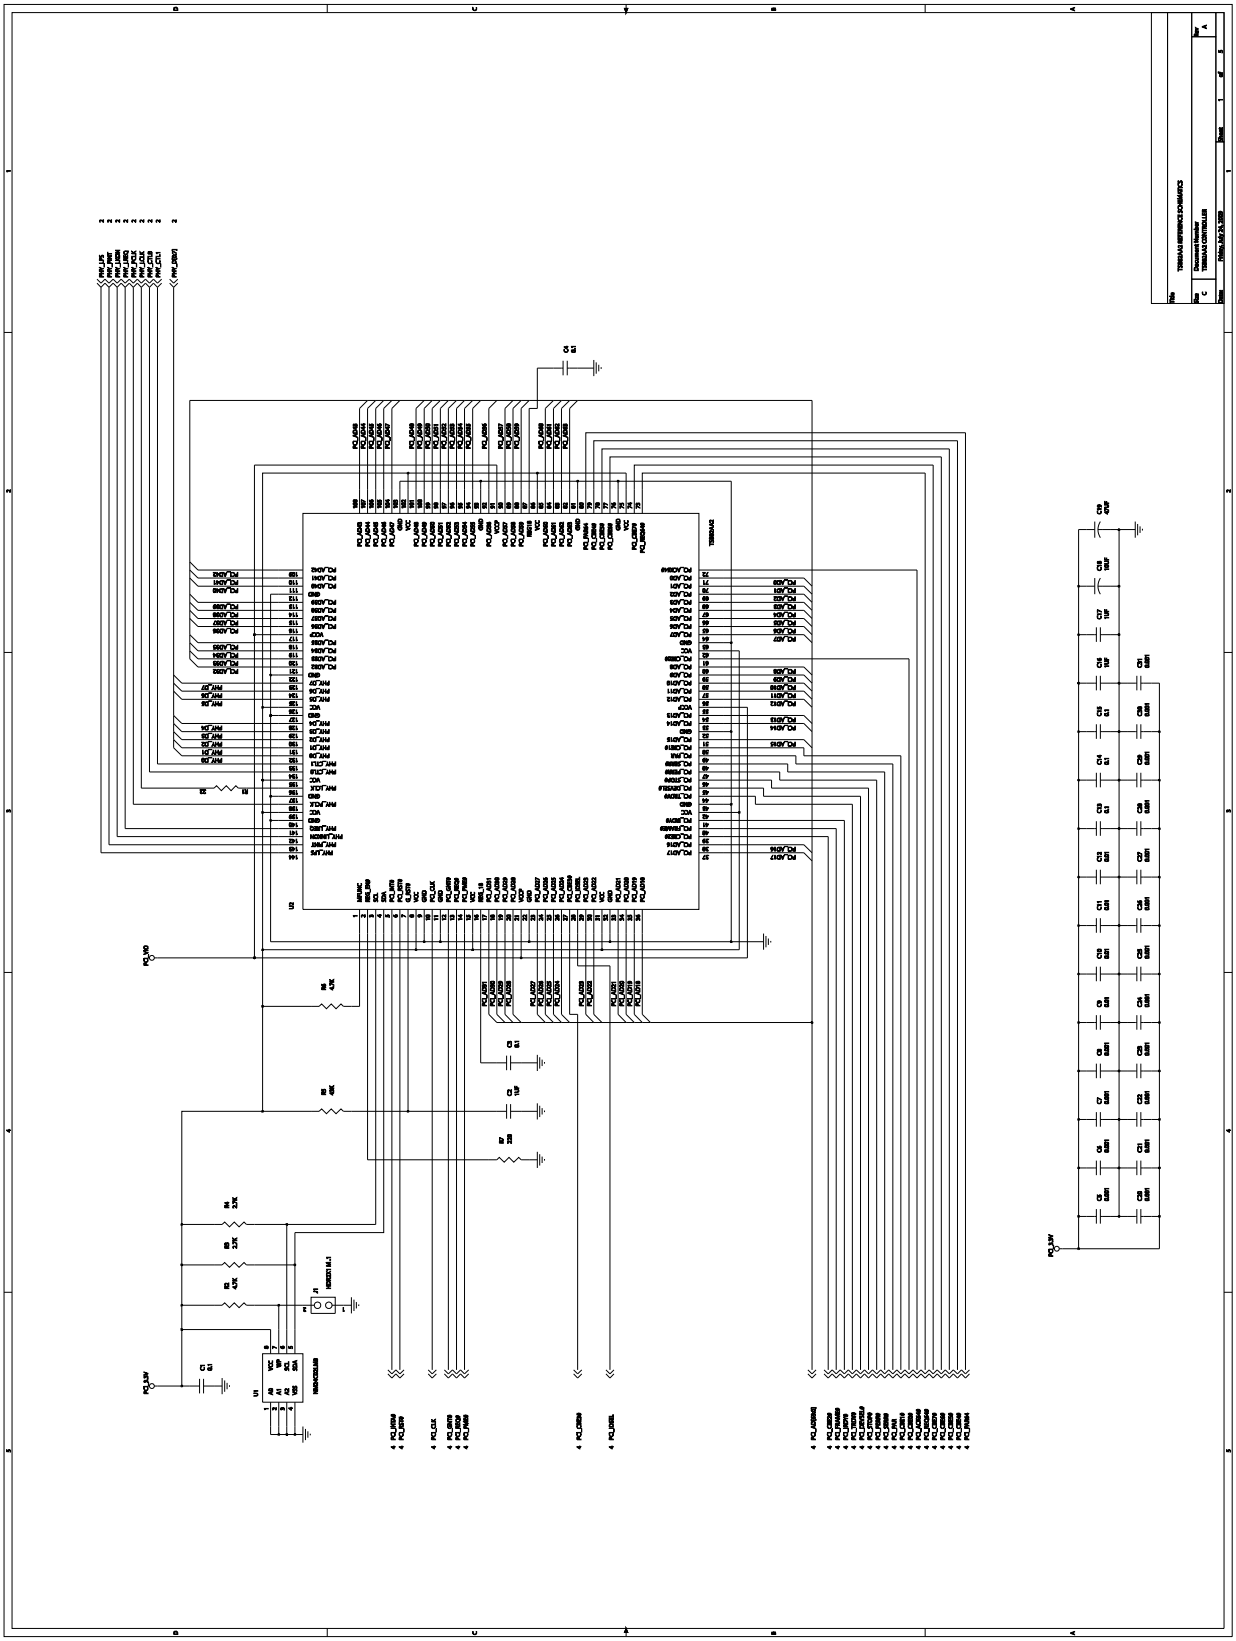

Figure 8-1—Example FW800 OHCI Controller

Notes on Example FW800 OHCI Controller (figure 8-1.)

1) For cost reduction, input signal terminals may be tied directly to GND if required to set their state. However for more conservative (better) ESD performance, signal pins that need to be set to a ground state should be pulled to ground through a resistor valued from 220 Ohms to 1.2 Kohms (general guideline for values).
2) The oscillator in the reference design has a 3.3V output. The PHY XI pin requires a 1.8V input signal. Therefore there is a voltage divider on the output of the oscillator to shift the oscillator 3.3V output to a 1.8V level for the input to the PHY. It is also possible to use a 1.8V oscillator and avoid the use of a voltage divider. The oscillator circuits must be laid out to minimize noise and induced jitter. The oscillator used must be an extremely low jitter oscillator. Using a separate Vcc power plane for the oscillator can help minimize jitter and noise.
3) LCLK is a 98.304MHz input to the PHY that should be series terminated to reduce standing waves at its source (the 1394b link layer).
4) CPS is a connection from the PHY, which has internal signals of 500MHz, to the cable, which is a large antenna. The 390K resistor should help prevent EMI, but if required a ferrite may be placed in series with the CPS signal to reduce the chance of PHY signals unintentionally getting on the 1394b cable.
5) LPS determines if the PHY-link interface is active, a pulldown is included so that even if the PHY is powered, but the link is not powered, that the LPS pin will be in the correct state.
6) Pin 26 is not used, but will have a 98.304MHz signal present on it. To help prevent EMI the customer may terminate this signal with an approximately 10K Ohm resistor to ground.
7) To pull up a signal to the power rail, an approximately 1 K Ohm pullup should be used. Pins should not be tied directly to a power rail.
8) It is recommended that pins 2, 32, 33, 73, 66, 67, and 68 (PC0, PC1 and PC2) be capable of being pulled to either the power rail or to ground based on the desired configuration. These pins can control aspects of PHY behavior that a customer may want to change without having to create a new board or write new software. It is recommended that for all cases but pin 2 this be done by having the following option resistors: a 1K pullup to the power rail and a 220 Ohm pulldown to ground. To program a pin either one or the other of the resistors can be populated to set the desired functionality. Since the 220 Ohm pull-down will dominate, even if both are installed the signal will be pulled down. Pins 2, 32 and 33 (DS2/LKON, DS1 and DS0) also control aspects of PHY behavior that can be changed without having to create a new board or write new software, however, they are set appropriately for this reference design. See the datasheet for more information on the function and appropriate setting for these pins. Can also be an output, therefore it should not have a value lower than 1K Ohm to either rail. Therefore pin 2 should have a pullup and pulldown of 1Kohm and one or the other is populated, depending on the value to be programmed to the pin.
9) The PCLK signal is a 98.304 MHz clock signal. To reduce standing waves it is recommended a small series resistor be included close to the PHY pin 5. The value used on the reference board is 22 Ohms.
10) Cable Not Active (CNA pin 79) output pin is just brought out to the test header pin on the reference design. It may be left unconnected on the reference design. However it goes logic high when ever either bias (if 1394a connection is made), tone, or a trained connection is made to any of the device 1394 ports. This signal may be used by HW circuitry external to the PHY to indicate when a connection has been made to the PHY. Note that the PHY can be programmed to generate an interrupt to the link layer for the same circumstances, if SW control is desired.

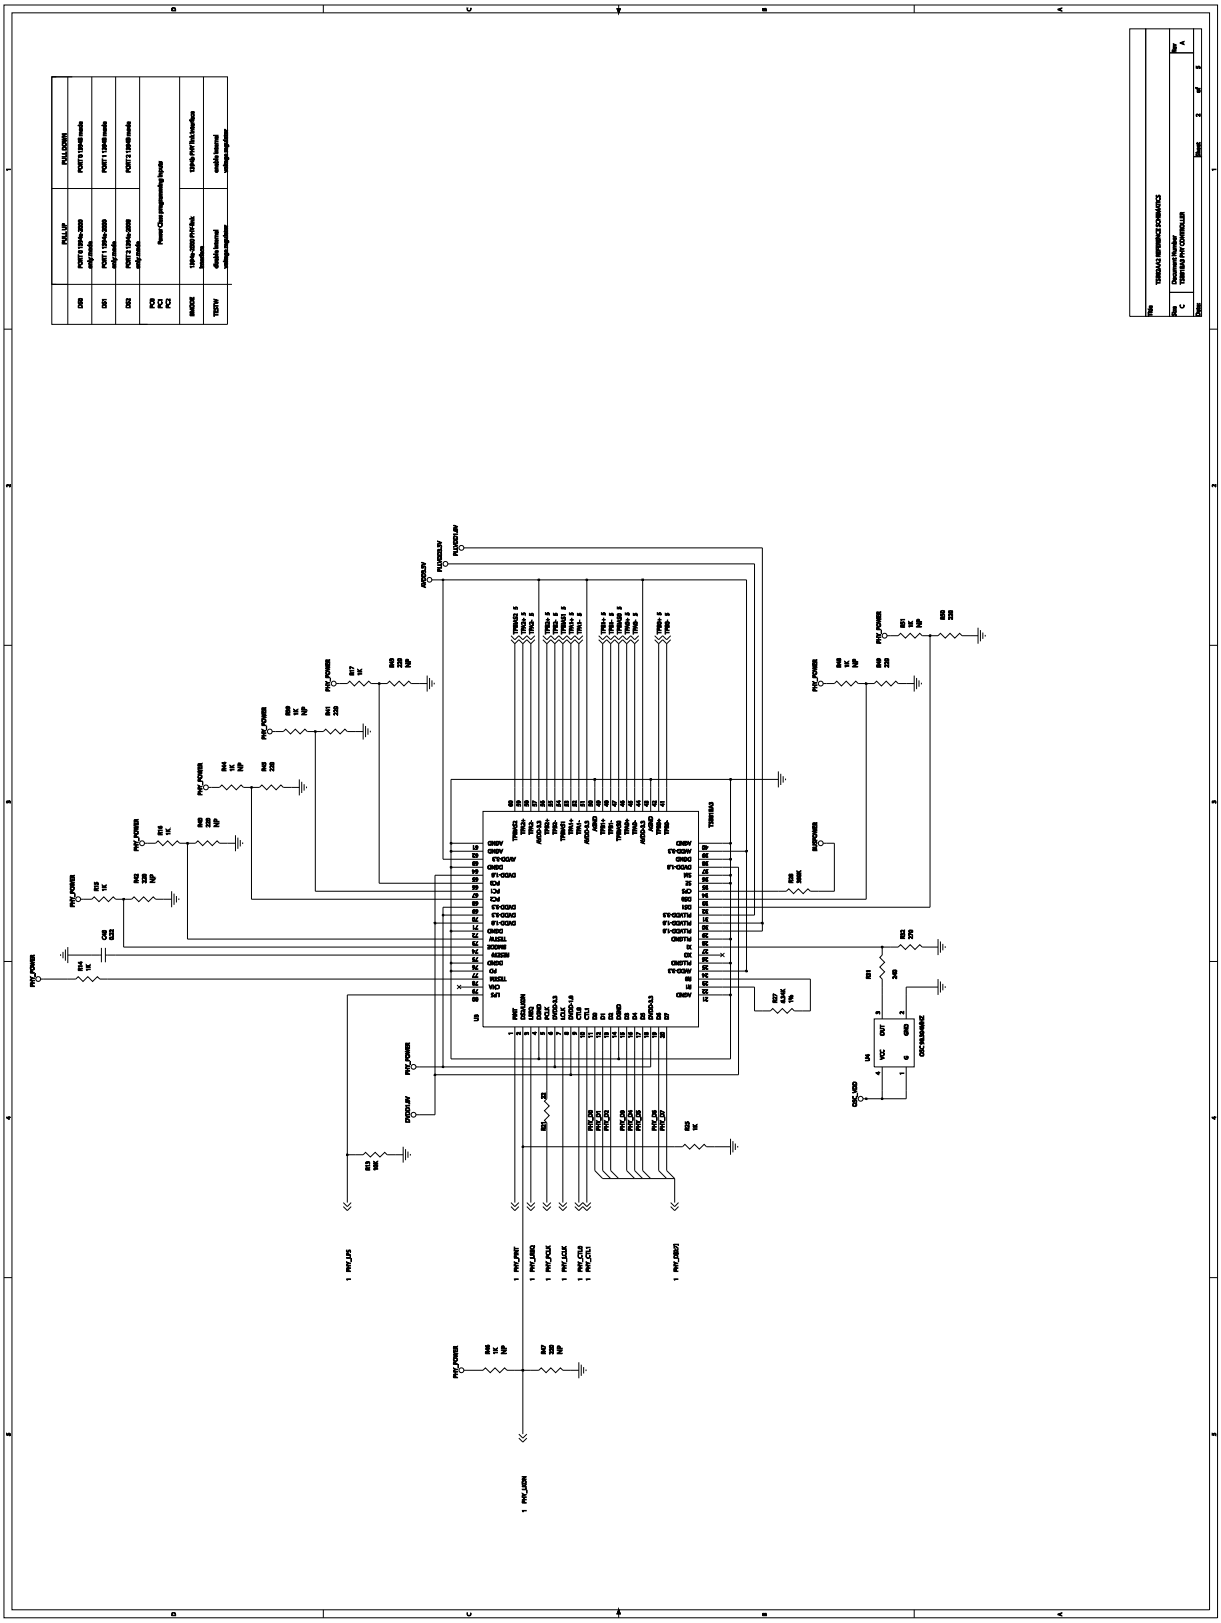

Figure 8-2—Example FW800 design (ports 1 and 2)

Figure 8-2 Continued

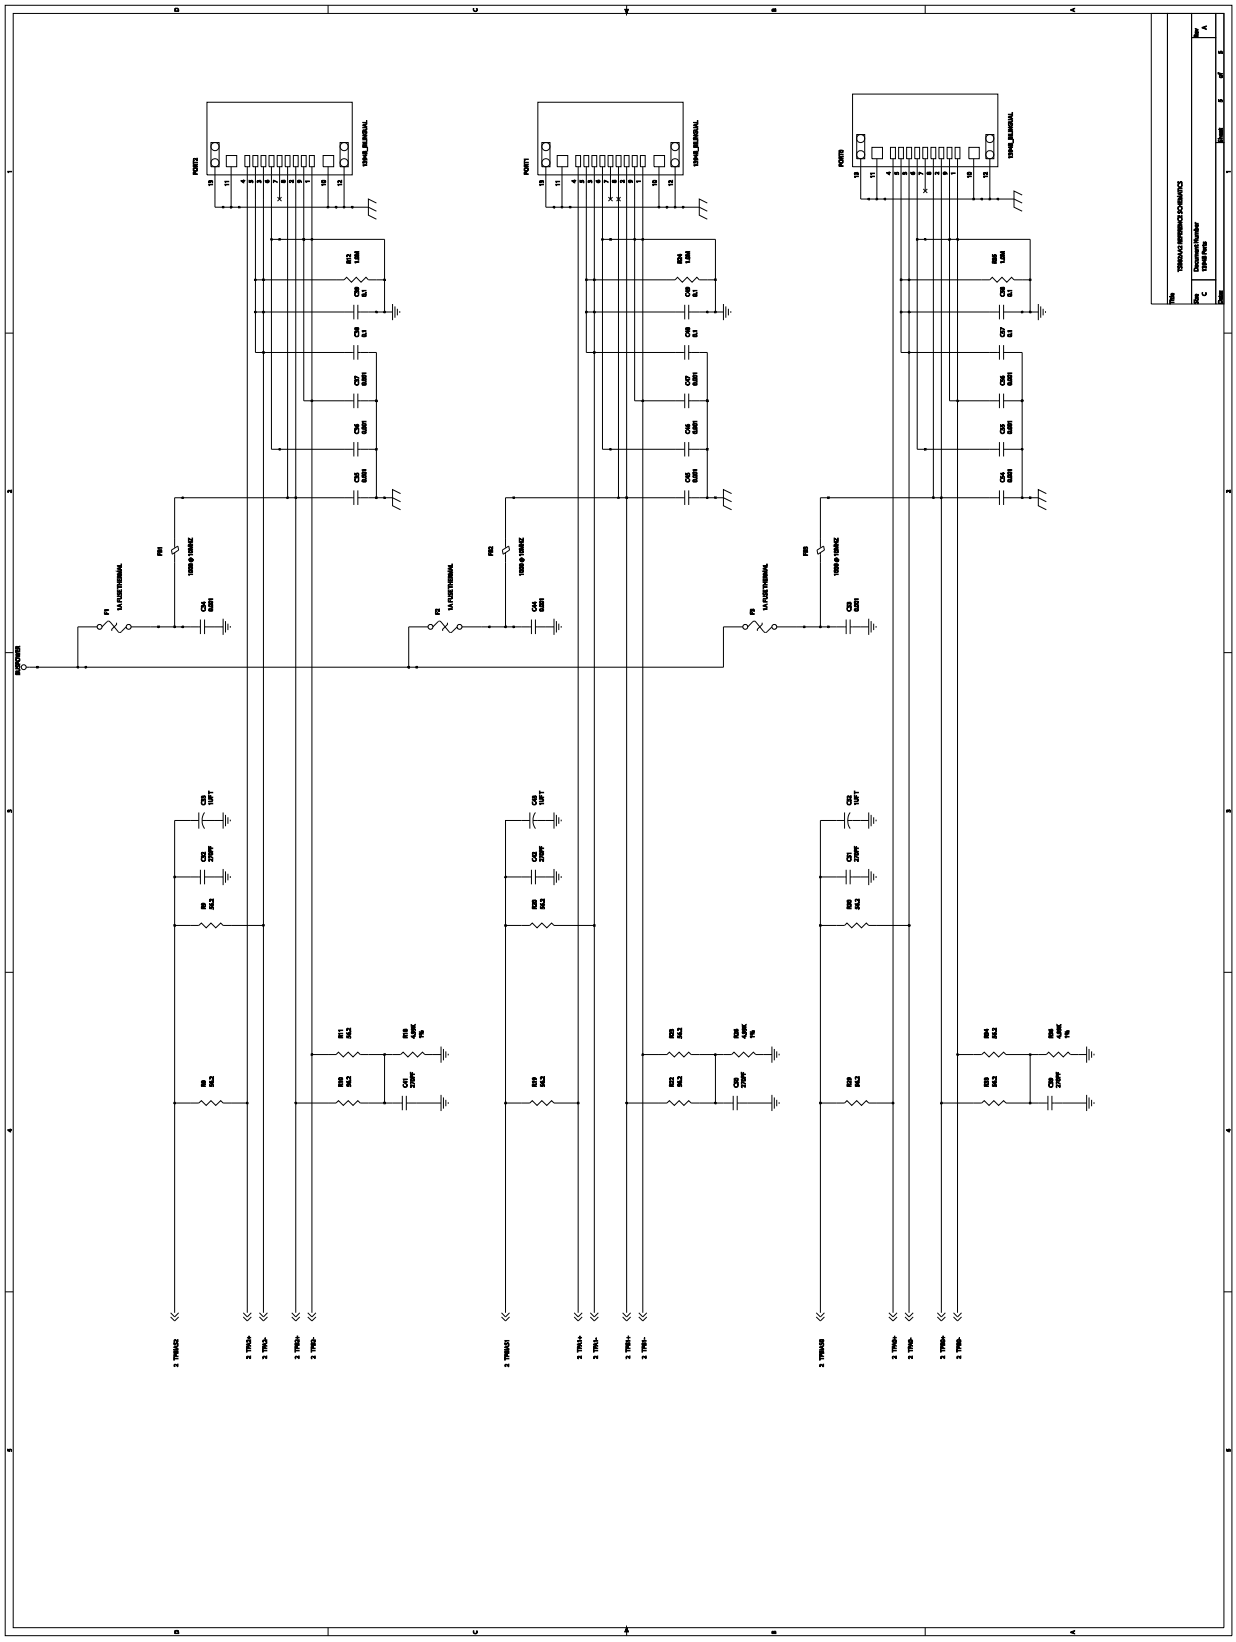

Notes on figure 8-2.

This reference design is for a PC add-in card, where the system has a chassis ground for ESD and EMI control.

1) This design uses 1394b bi-lingual connectors. These connectors are used when the PHY port may be either a 1394a port or a 1394b port. What mode the port talks is determined by what type of device is plugged into the port from the other side of the connection. It is recommended that customers use the surface mount connectors with through-hole mounting pegs that are soldered into the board for more rigid connections.
2) The outside shield of the connector should be tied to chassis ground to provide a low impedance path to the chassis ground for ESD currents.
3) Pin 6 is the VG ground return for cable power and speed signal current. Pin 9 is the ground return for TPB (the transmitting pair of a 1394b port). Both of these pins should be via-ed to the board ground plane as soon as possible. Pin 5 is the ground return for TPA (the receiving pair of a 1394b port), it is isolated from local ground for future definition in the standards. See 1394b paragraph 5.5.1.
4) Connector pins 5, 6, 8, and 9 each have an ESD capacitor connecting them to chassis ground. This is in an attempt to dump as much ESD energy as possible to chassis ground before it gets distributed onto the local board ground plane. Therefore these capacitors should be placed as close as possible to the connector and before the pins are via-ed to the board ground plane. This must be traded off for better ground connection (make traces fat).
5) The ferrite bead on the VP pin (pin 8) of the connector is placed for both ESD and EMI reasons. The 0.001 uF capacitor on the PHY side of the ferrite is meant to provide a low impedance path to local ground for any high frequency EMI currents that might exist on the VP power trace. The 0.001uF capacitor on the 1394b connector side of the ferrite is meant to provide a low impedance path to chassis ground for fast rise time ESD currents that might have coupled onto the VP trace from the cable. The ferrite is meant to present a high impedance to both of these currents, forming a “PI” filter.
6) The transmission line termination networks are required and should be placed close to the PHY as possible (see layout recommendations). The 1uF capacitor is to provide charge for 1394a speed signaling. The 270pf capacitors at the mid-points of both the TPA and TPB pairs termination network is to provide common mode filtering for noise reduction and EMI reduction.

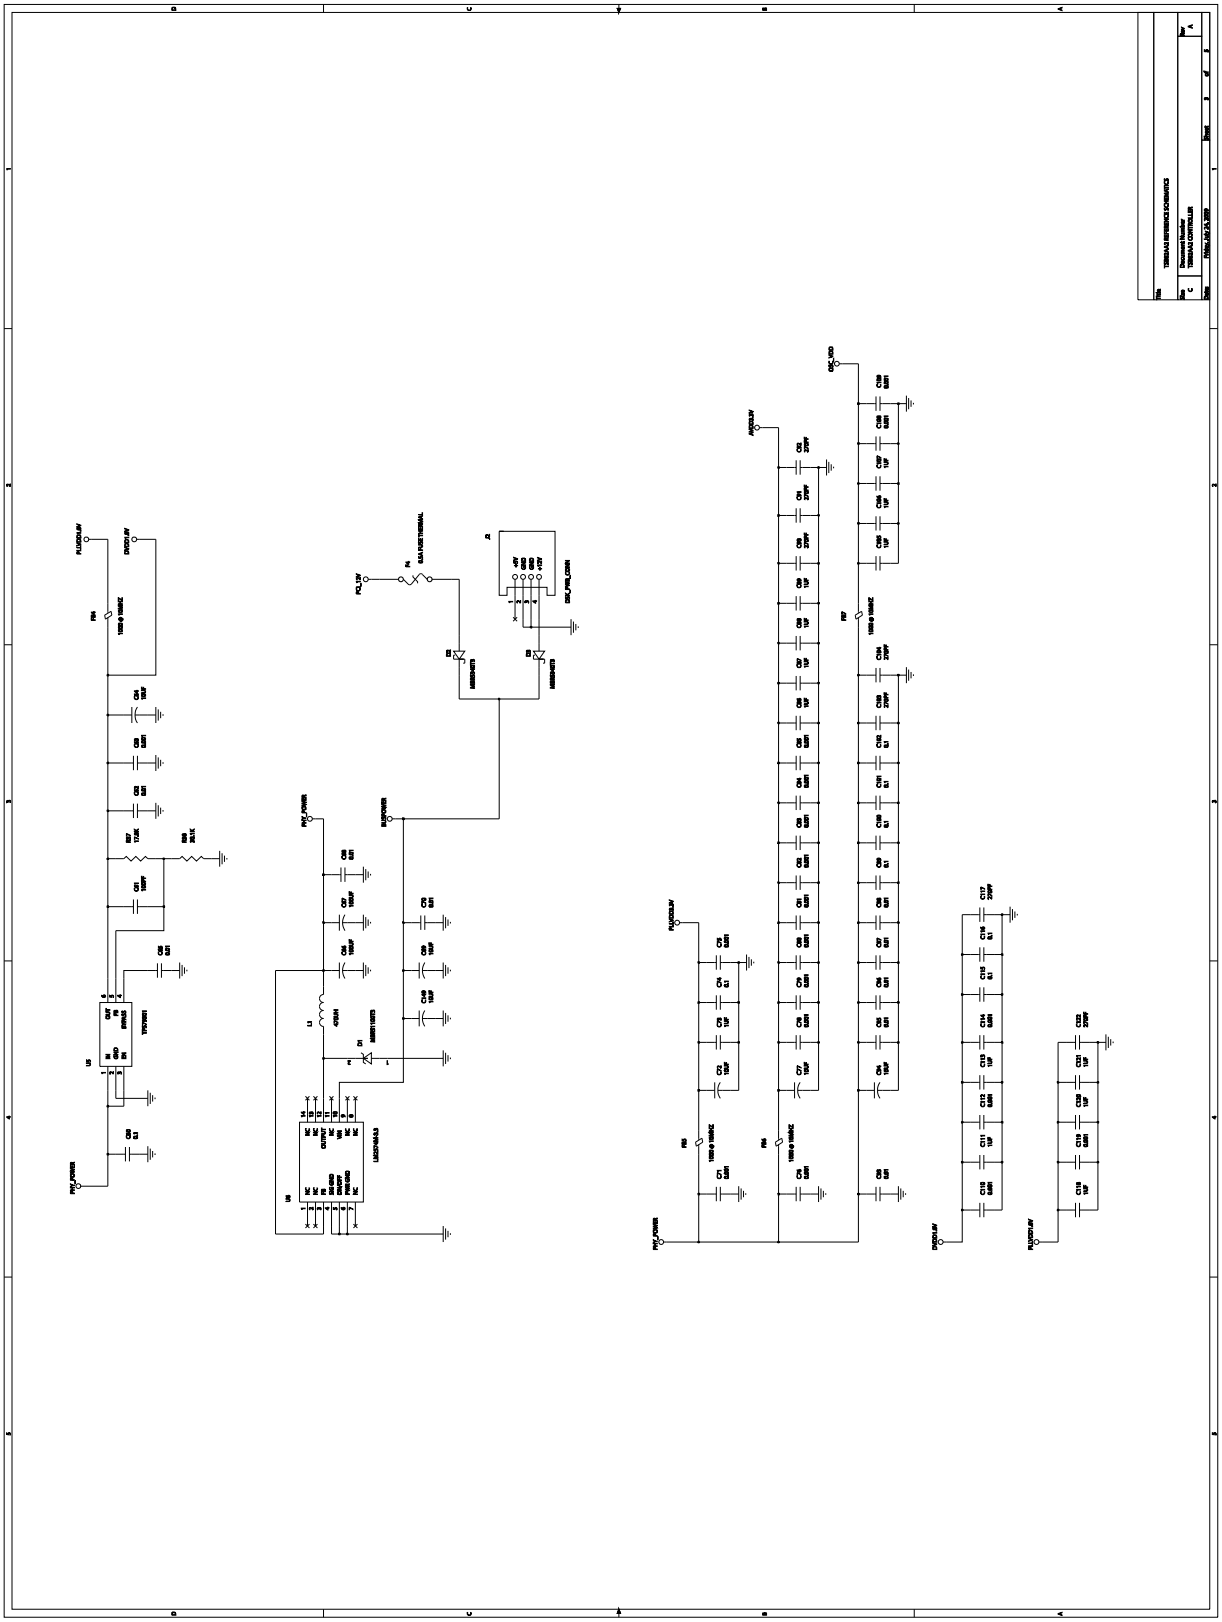

Figure 8-3—Example FW800 design (port 3 and PHY power)

Notes on figure 8-3.

The regulator provides power to the low voltage core circuitry of the PHY. The ferrite bead is meant to provide noise filtering to prevent noise generated from the digital circuitry from contaminating the PLL circuitry. The regulator chosen provides 200mA and is a low noise regulator.

8.1 Layout recommendations for S800 1394b TSB81BA3 Physical Layer:

In order of priority

a) Twisted Pair Transmission Lines
   1) The twisted pair lines must be very clean. The bit rate is 1Gbps (raw with out coding loss), this implies very tight tolerances for jitter and edge placement. The minimum rise time for 1394b is 80ps. An 80ps rise time gives a critical etch length of 4.7mm. What this is saying is that the twisted pair lines MUST be laid out as controlled impedance (110 Ohm differential) transmission lines. The layout for the twisted pair transmission lines should be as short as possible. An individual twisted pair should have less than 10 mils difference between the lengths of the lines. The lines need to come together and be run as a pair until they connect to the PHY terminals.
   2) The termination network must have a minimal effect on the transmission lines. For the reference layout, “fly-by” termination was used. This means the transmission lines were run directly from the connector to the PHY without disruption. Then the termination resistor network was placed on the back side of the board to allow the signals to “fly-by” the PHY terminals before being terminated. This also reduces the stub length from the termination resistors to the end of the transmission line.
   3) To ensure that there is a good ground return path from the PHY terminals to the connector, at a minimum it is recommended that a solid ground plane, without cuts, be implemented. This plane should extend immediately below the PHY, the PHY twisted pair terminals, the twisted pair transmission lines and the terminals of the 1394b connector to the point that the connector ground return terminals (pins XX) are via-ed to this solid ground plane. Via the connector ground return pins and the power gnd connector pin to the solid gnd plane as quickly as possible.
   4) It is required that at least a 4 layer PWB be used to ensure the good ground return path for the 500MHz signaling.
b) Oscillator
   1) The oscillator must be laid out to minimize noise introduced into the PHY PLL circuitry. The oscillator must be located as close as physically possible to the PHY XI input terminal. A wide low impedance etch should be used to connect the oscillator to the PHY XI terminal. This connection should be short and direct. The TSB81BA3 requires an input voltage of 1.8V. The selected oscillator outputs 3.3V. This requires a voltage shift to the lower levels required by the PHY. The means selected in the reference design was a resistor divider. Note that this also must be done with care. The location of the resistors was done to minimize the disruption to the 98.304MHz transmission line from oscillator to PHY terminals.
   2) The oscillator was laid out with a mini-power plane just for the oscillator. This power plane was isolated by a filter from the other power planes both to reduce the noise to the oscillator, but also to reduce the noise introduced from the oscillator to the rest of the board. This plane also helped minimize the impedance from the decoupling capacitors to the power inputs of the oscillator
c) Power Plane Decoupling
   1) PLL power plane decoupling - tbd
   2) Analog Power plane decoupling - tbd
   3) Digital power plane decoupling - tbd
d) PHY Link interface
   1) There are two 98.304MHz clocks on the PHY link interface. These traces need to be spaced well away from the other PHY-link interface pins to reduce coupling onto the other lines and those lines acting as antennas to radiate the noise to other part of the board. There are source terminating resistors close to the source of each of these clocks. Making room for the resistors pushed the other traces away. The other traces were routed with more than minimum spacing to comfortably route from PHY to link.

2) The switching power supply was placed well away from the other active circuitry on the board. This is to isolate the switching noise away from active and especially analog circuitry.

e) Thermal Land underneath PHY

It is strongly recommended, though not required, that a thermal land be placed underneath the PHY.

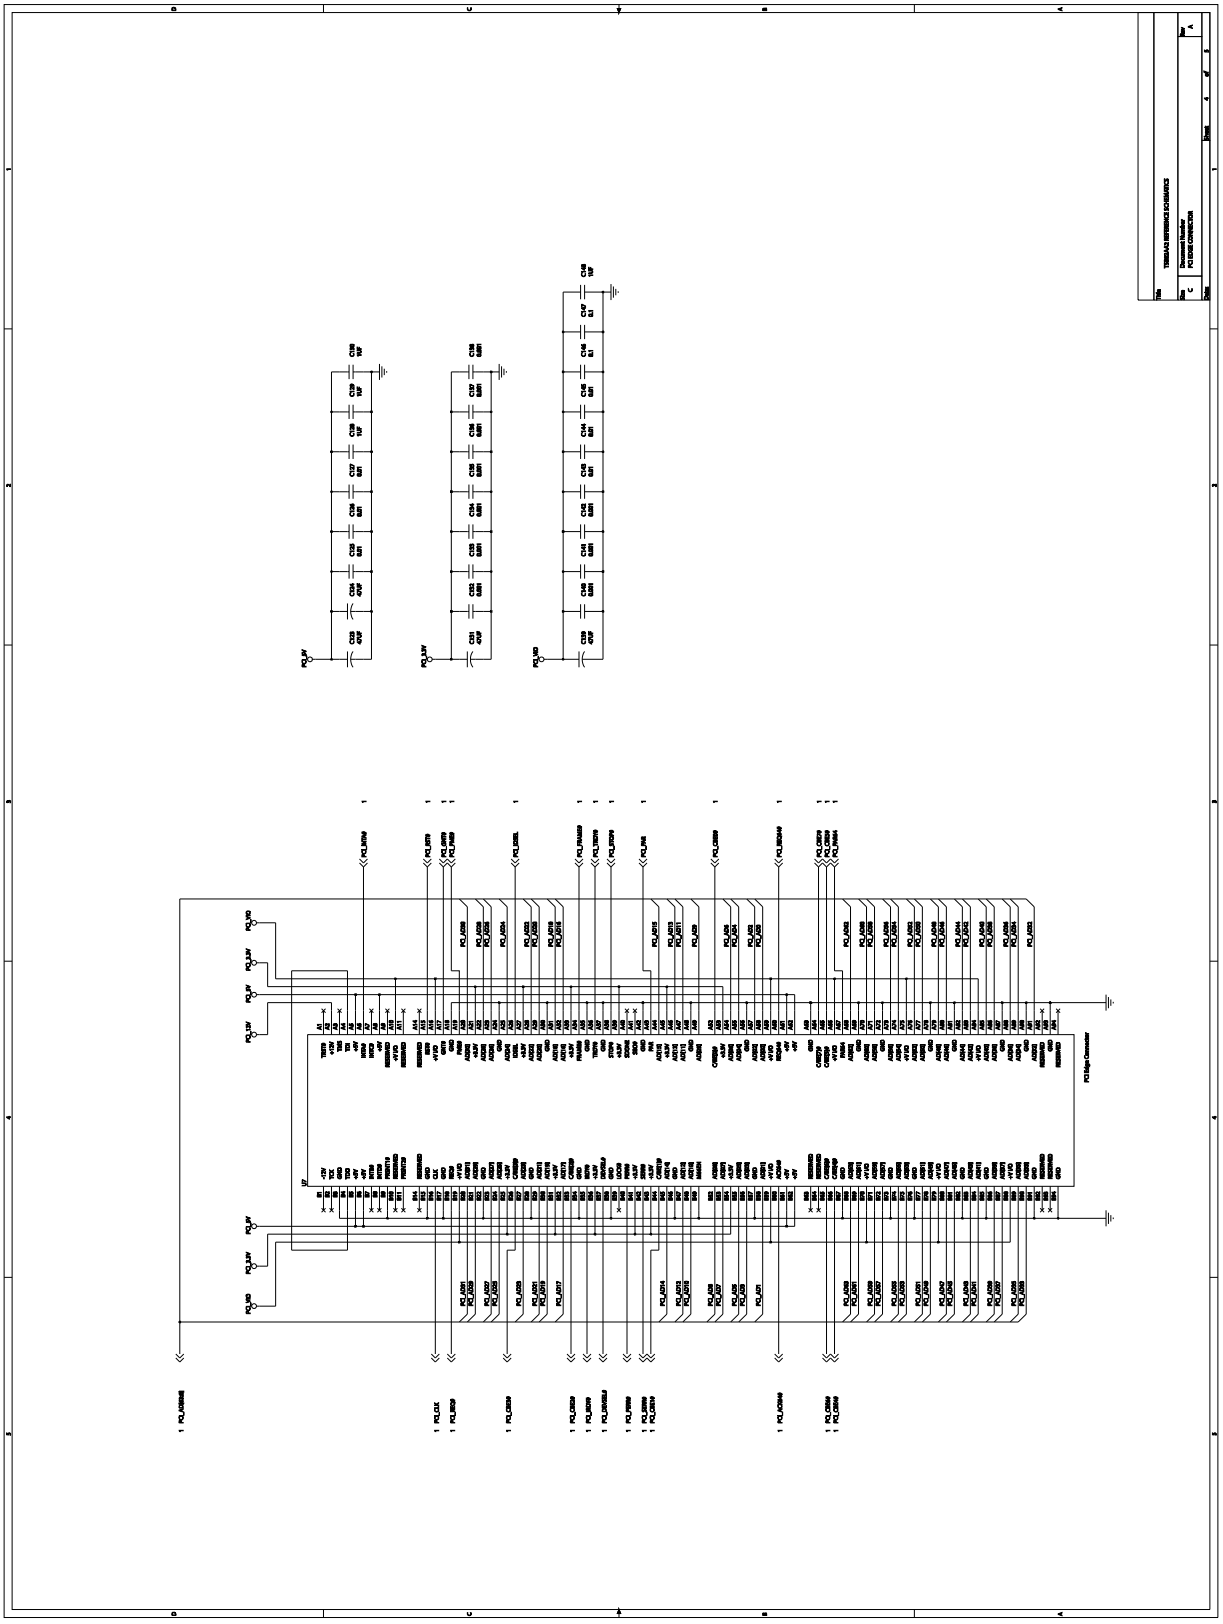

Figure 8-4—PCI Connector Reference

9. Robust Port Design

9.1 Circuit Protection Needs for Firewire Devices – Introduction

At various places in this guide, there are recommendations for adding circuit protection when implementing FireWire (IEEE 1394) ports on electronic devices to ensure the robustness of the application while in field service. Two major areas are considered: power circuits and signal circuits. This section provides additional information on circuit protection, including a discussion of events that can cause circuit damage, industry standards on circuit protection, placement of protection devices in a circuit, and typical device types that are used.

In order to ensure high yields in the manufacturing process, basic EOS (Electrical Overstress) protection features are implemented within a 1394 PHY (physical layer device). However, it should be noted that they are optimized for the manufacturing environment and designed for those respective test standards (HBM, MM, and CDM). Application-level testing to ensure EMC compliance uses a different model that includes more severe current surges at higher voltages. The on-chip protection devices often cannot protect against these events and supplemental protection can be added at the board level to ensure reliable performance of the application. In general, for robust port design, all pins of a 1394 port should be protected as described herein. This includes ESD, EOS and overcurrent protection.

In Figure 1, a notebook computer that implements a number of interface ports is shown. Specifically, it includes 1394 ports. The occurrence of user-originated electrical threats should be expected as these ports are used by the consumer to connect the notebook to various peripherals (printer, external storage, digital camcorder, etc.). During connect and disconnect actions, incidents such as ESD, Late VG or fault current may occur.

Electrostatic Discharge (ESD) occurs when electrical charge is transferred from one object to another. In this case, the user or the cable can become electrically charged, and then discharged through the 1394 port. Late VG occurs when the ground connection is lost, but the power bus is active in the cable. Power is forced onto the data pairs and damages the PHY. Fault current can occur because of a faulty cable or connector (bent pin, etc.) or due to a metal object temporarily inserted into the connector. Details on the electrical threats as well as potential solutions will be provided in the following sections of this document.

Figure 9-1—Overview of notebook computer showing implementation of IEEE 1394 ports.

9.2 “Hot” Connection Problems (also known as “Late-VG”)

9.2.1 Background

All 1394 PHY devices are susceptible to electrical overstress damage when exposed to a combination of high voltage cable power and a faulty cable or connector system that allows data (TPx) and cable power (Vp) connections to engage before the cable ground (Vg) connection. This set of circumstances is known as a “Late Vg” event. This document provides information on how to recognize a 1394 PHY device that has been damaged by a Late Vg event, an explanation of how the damage occurs and suggestions on both decreasing the frequency of Late Vg events in a system and protecting 1394 PHY devices from damage due to Late Vg events.

The IEEE 1394A-2000 specification calls for 1394 cable power (Vp) to range from 8 VDC up to 30VDC with a maximum current of 1.5A. Legacy IEEE 1394-1995 systems can support cable power of 8 VDC up to 40VDC, also with a maximum of 1.5A current. A typical 1394 PHY device has an absolute maximum input voltage rating of -0.5 VDC to Vdd + 0.5 VDC, where the maximum value of Vdd is 4.0 VDC. There is an obvious potential for catastrophic damage to a 1394 PHY device if the voltage levels of Vp are expressed on the differential pair signals of the 1394 PHY device.

If the Vp connection and any of the data (TPx) connections of two 1394 nodes are engaged before the Vg connection due to extreme angling of the cable during connection or if the Vg connection is not well connected due to distressed cables or connectors, some of the voltage on Vp of the 1394 cable can appear on TPx signals. For example, referring to Figure 2, it can be seen how pins 1 and 3 can be connected before or without the pin 2, Vg, connection.

The Vp voltage appears on the TPx signal(s) because the data connection is used as a ground return path in lieu of the Vg connection. The high voltage and current on the return ground path can damage the TPx inputs of the 1394 PHY devices on both sides of the cable. ESD protection circuitry implemented inside the 1394 PHY device cannot protect against a Late Vg event because it is of longer duration and higher current than a typical ESD event. It should be noted that only one of the TPx connections is required for a Late Vg event to occur. A node with only Vp and TPB* connected would still be susceptible to Late Vg damage.

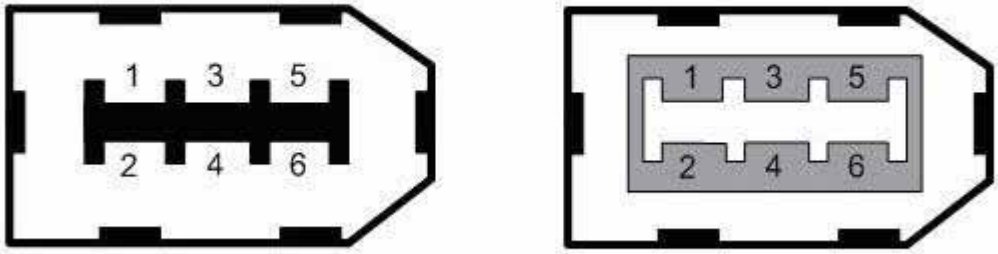

Pin Number	Contact Name	Comment
1	Vp	Cable power
2	Vg	Cable ground
3	TPB*	Differential Pair B
4	TPB	
5	TPA*	Differential Pair A
6	TPA	

Figure 9-2—1394 6-pin receptacle (connector) and plug (cable) diagrams.

9.2.2 Symptoms of 1394 PHY Device Damage due to a Late Vg Event

To the end user, a 1394 port damaged by a Late Vg event will simply fail to operate. In most cases, the other ports of the same 1394 PHY device are not affected and continue to work properly. Replacing the damaged 1394 PHY will fix the issue completely since no other portion of the application is typically affected.

When a damaged PHY device is bench tested, it is likely to show continuity failures. Most often there is an open or short on one of the TPA lines of the failing port. Alternately, damage on the TPB pair is also seen in 1 out of 4 cases. Ancillary damage to the TPBIAS circuitry has been observed in some cases as well.

A full failure analysis of a 1394 PHY device that has experienced a Late Vg event will often show a failure mechanism of damage to metal, polysilicon and die substrate resulting in the short to ground or open, consistent with electrical overstress. Excerpts from previous failure analyses of EOS damaged PHY devices are shown below. This is damage which is indicative of a Late Vg event. Not every 1394 PHY damaged by a Late Vg event will show the exact same damage pattern.

9.2.3 Characteristics of Applications Susceptible to Late Vg Events

Some 1394 applications are more susceptible than others to 1394 PHY damage due to Late Vg events. Characteristics of these applications are listed below for reference. Please note that these characteristics are not required for a Late Vg event nor do they indicate that a Late Vg event will definitely occur. This list is purely informative. It is provided to allow a designer to determine if Late Vg is likely to be a concern for their application.

Cable Powered System: The damage caused by a Late Vg event is linked to cable power that is seeking a return path, thus systems that rely heavily on 1394 cable power are more likely to see Late Vg damage. A Late Vg event cannot occur on a 1394 node where only a 4 connection (TPA, TPA*, TPB, TPB*) receptacle is used. If there is no Vp connection, by extension there can be no damage. In addition, self-powered applications (a system which never sources or sinks current from the 1394 cable) typically do not see Late Vg events. It is, however, theoretically possible for a 1394 self-powered node to experience damage due to a Late Vg event, if it is a two port application that is repeating power to a 1394 device that is cable power consumer.

High Voltage on Vp: The IEEE 1394A-2000 specification limits Vp supported on the 1394 bus to 30VDC. Typical PC applications source 10VDC-12VDC on Vp driven by the PCI bus. Damage to 1394 PHY devices due to a Late Vg event occurs most often on systems with greater than 15VDC on Vp. The increase in voltage that a PHY device is exposed to during a Late Vg event also increases the likelihood of device damage.

1394 6-Pin Connectors: Late Vg events can occur on 1394B systems implementing 9-pin connections, however the 1394 6-pin connector is much more susceptible to angled cable insertion or even cable reversal due to its geometry. This increases the likelihood of a Late Vg event.

Repetitive Cable Use: In the majority of cases of 1394 PHY devices with EOS damage investigated by Texas Instruments, the root cause has been traced back to a single cable or system in the manufacturers test process. Heavily used 1394 cables or connectors can become worn or damaged after thousands of insertions, resulting in a degraded or nonexistent Vg connection which translates to an extremely high incidence of Late Vg events.

Hot Plug: In a 1394 system with intact 1394 cables and connectors, a Late Vg event is only possible during a cable insertion where Vp is already actively powered. If all 1394 nodes are cabled before cable power is applied, damage due to a Late Vg event is not possible, unless the Vg connection is somehow subsequently lost or broken.

9.2.4 Explanation of a Late Vg Event

When the Vp connection and any of the four data (TPx) connections between two 1394 nodes, one a power provider and the other a cable power consumer, are engaged without a Vg connection, the data connection becomes the lowest impedance ground return path to the Vp provider. The voltage and current on the ground return path can damage the TPx inputs of the 1394 PHY devices on both sides of the cable.

Figure 3 shows a schematic of possible electrical connections between a power provider node (NODE A) and a cable power consumer (NODE B) during a late Vg event. For this explanation, it can be assumed that there is no direct ground connection between the two 1394 nodes anywhere in the design. The cable shield in the figure is shown as shorted to chassis ground with no connection to signal ground on either 1394 node.

The Vp connection can be assumed to occur first because it is a longer pin in the connector, thereby raising all signals on NODE B to the level of Vp. Once a data connection is made between the two nodes, the ground of NODE B is pulled down and the voltage regulator on NODE B will engage. Since the data connection is acting as the ground return path, this forces the TPA/TPA* of NODE B in the figure to be pulled below the 1394 PHY device ground. This low voltage is most likely outside of the absolute maximum input voltage rating of -0.5VDC on the 1394 PHY device and can result in damage to the 1394 PHY device on NODE B.

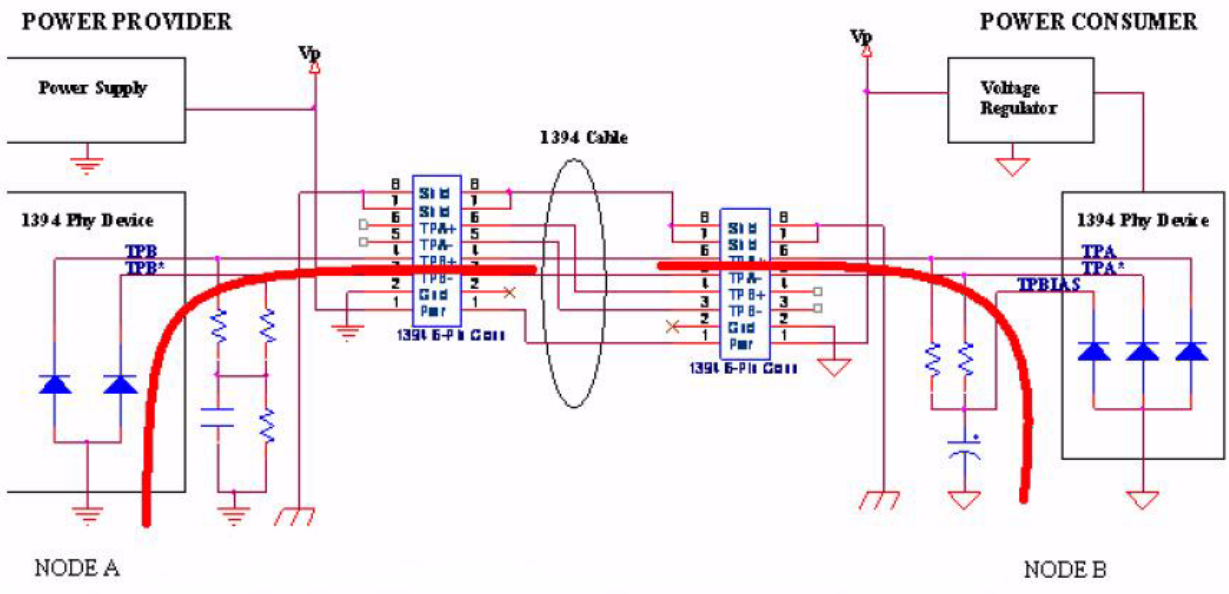

Figure 9-3—Overview of Late Vg event occurring during cable operation.

As the voltage regulator of NODE B begins driving out, additional current will be returned through the data connection raising the voltage on the TPB / TPB* signals of NODE A in the example. Depending on the cable power voltage of the system, the voltage on TPB/TPB* could be in excess of the absolute maximum input voltage rating of Vdd+0.5VDC on the 1394 PHY device and can result in damage to the 1394 PHY device on NODE A.

Thus a Late Vg event can cause damage on either the power provider node or the power consumer node. In a real world situation, the Vg connection is likely to appear after a short delay or in the case of damaged cables or connectors, the connection could be intermittent. While this may limit the exposure of the 1394 PHY devices to the effects of acting as the ground return path, it can also introduce transients that may cause additional damage to the devices.

9.2.5 1394 PHY Device Protection – Late Vg events

Passive Solution:

It is possible to decrease the frequency of Late Vg events by addressing cable and connector quality issues, however, these events remain a concern for many 1394 applications. While there is no way to conclusively protect 1394 PHY devices against all EOS damage that could be caused by a Late Vg event, it is possible to limit the effects.

One easy-to-implement suggestion is to provide an alternate low impedance ground return path by shorting the chassis ground of the cable shields to digital (signal) ground. This can be done through a low impedance filter to prevent excess noise if needed.

Another option is to add high speed switching double diode circuits like the BAV99 to each of the TPx lines to protect them from damaging voltages without affecting signal integrity. Figure 4 shows how these preventative measures can be implemented in a typical 1394a application. Refer to Figure 5 for the details of a typical 1394b application.

When the elevated voltage is impressed on a data line, the diode will break over in the forward direction and shunt the Late Vg current to the power rail. Later, when the Late Vg condition is resolved, the diode will transition to its off state and data signals can again move down the data pairs. It should be noted that the use of these diodes may be limited in future, higher-speed versions of IEEE 1394 implementations. The parasitic capacitance of the diodes may cause sufficient signal distortion that signal integrity is compromised.

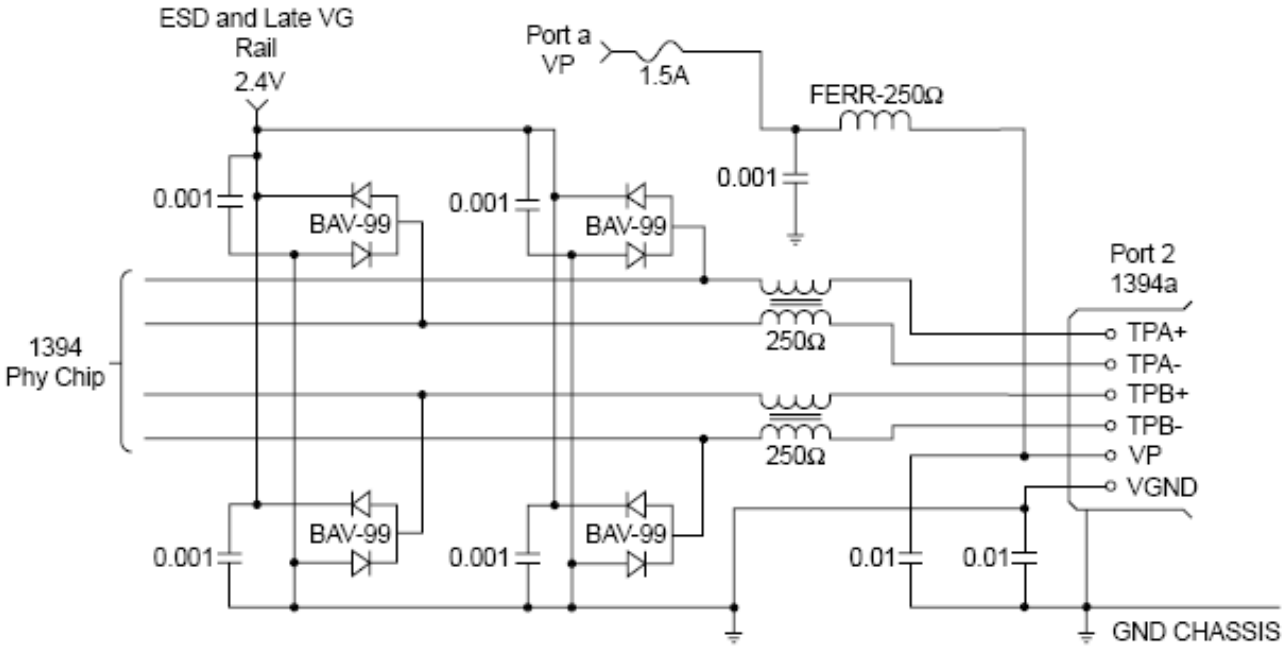

Note: All resistances in ohms and capacitances in microfarads, unless otherwise noted.

Figure 9-4—An example of Late Vg protection implementation using diodes in 1394a port

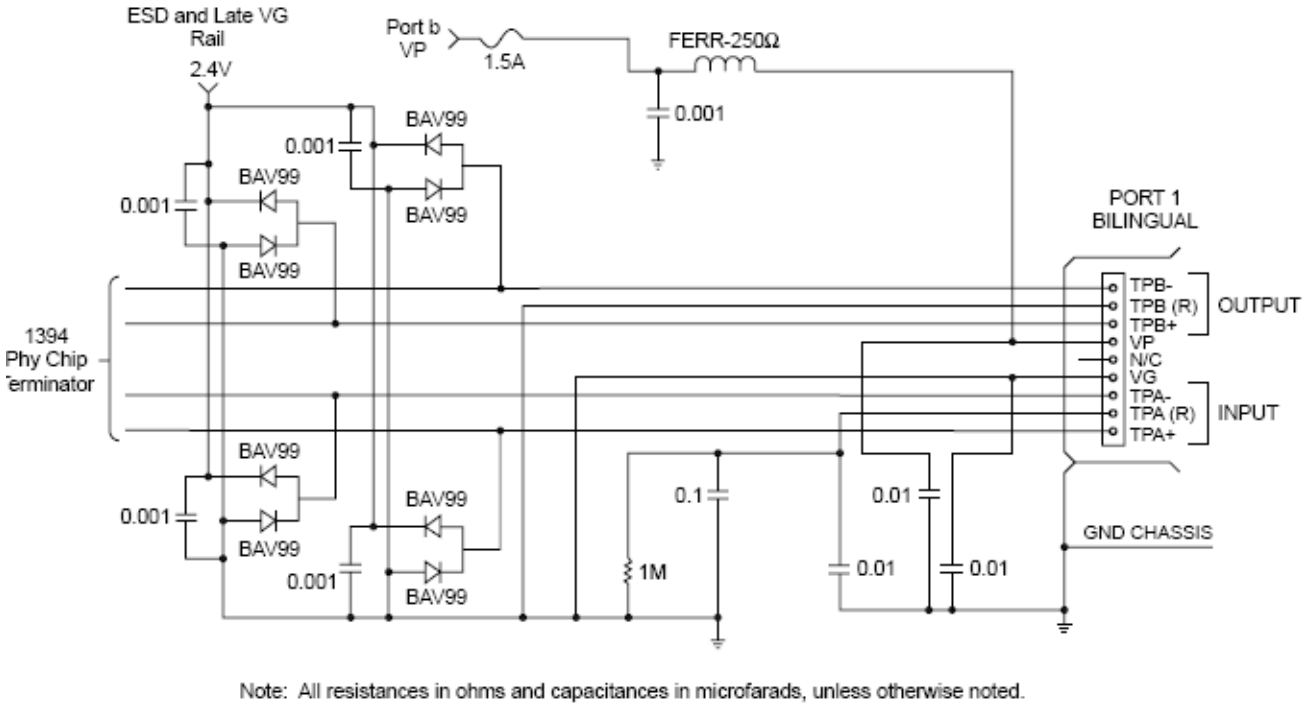

Note: All resistances in ohms and capacitances in microfarads, unless otherwise noted.

Figure 9-5—An example of Late Vg protection implementation using diodes in 1394b port

Active Solution:

Another approach to Late Vg protection is to use active circuitry to monitor the condition of the power and signal lines and to react by disabling power during a Late Vg event. Referring to Figure 6, the current limiting IC monitors the output of the comparator and if the ESD and Late Vg rail is at a higher bias than the 3.3V rail, power will be disabled. The input to the ESD and Late Vg are the BAV99 diodes that break over during the Late Vg event.

Note that this solution can also be used to provide overcurrent protection during fault events. The sense resistors (referenced below in Note 1) are used to set the trip current, and essentially provide a “programmable” feature as the trip level varies as a function of resistance value. In the below example, 0.02? resistors set the trip level at 2.4A. As the resistance value is increased, the trip level decreases.

Lastly, more information can be found by reviewing the Maxim data sheet as well as their Application Note (AN3984) which discusses the use of the MAX5944 (dual port) and MAX5943 (single port) solutions.

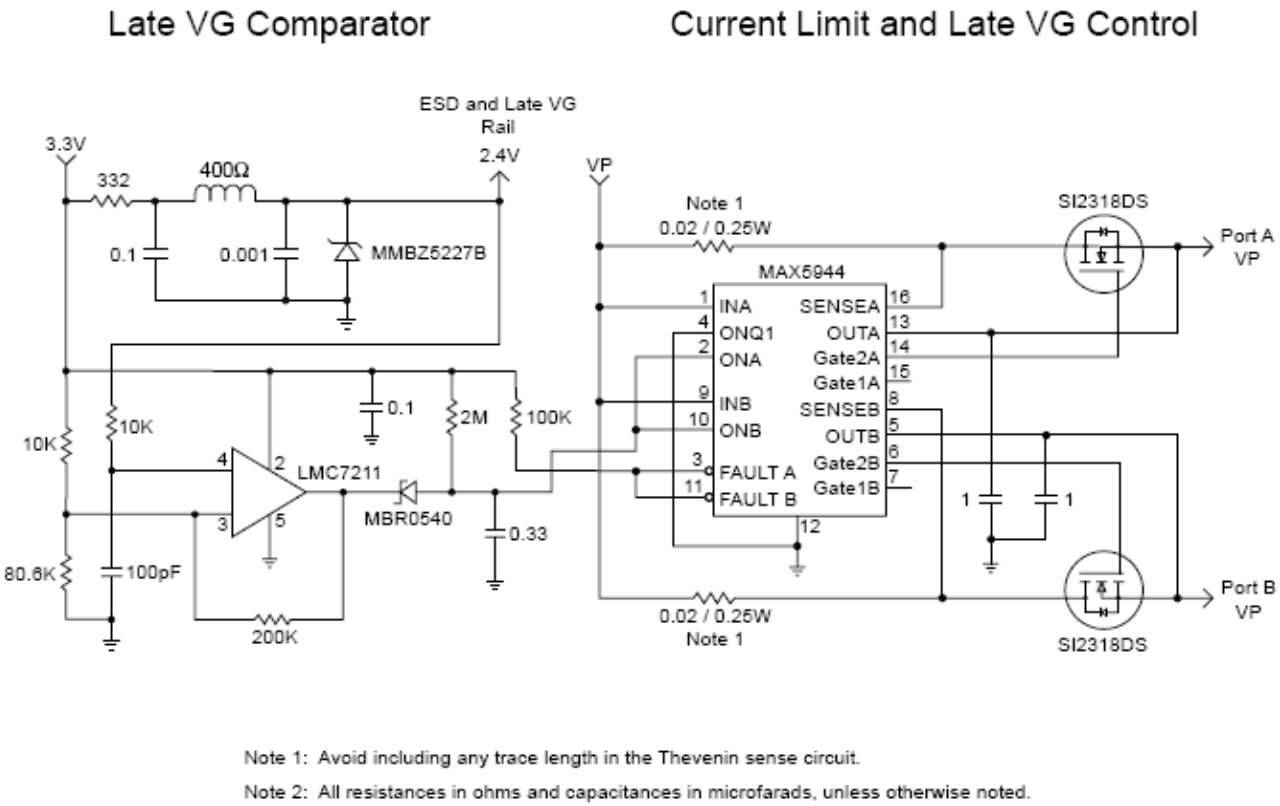

Note 1: Avoid including any trace length in the Thevenin sense circuit.

Note 2: All resistances in ohms and capacitances in microfarads, unless otherwise noted.

Figure 9-6—An example of Late Vg protection implementation using a Maxim protection IC.

9.3 VP Line Fault Currents

When a FireWire device is powered by the system power bus, the VP line from the 1394 port to the PHY should be protected from fault currents on the bus. Fault currents occur when a component connected to the power bus fails or is damaged such that its resistance value dramatically drops. Likewise, power-providing FireWire devices should be protected against user misapplications or defective connections (bent or broken cable/pins, etc.). To protect the PHY from high pass-through (fault) currents, this Guide recognizes the use of resettable fuses, because they have the ability to reset (restore power) after a fault current event is cleared/removed.

Resettable fuses can prevent damage from excessive currents, and then restore the circuit to normal operation when that level of current is no longer present. Typically, these devices are positive temperature coefficient (PTC) thermistors, whose resistance increases due to self-heating (I2R), and thereby limit current in the line on which they are installed. Their use is recommended in FireWire equipment because of its hot-plug IEEE1394 ports, which can be exposed to unpredictable or frequent fault current events on the VBUS line.

Various PTC manufacturer ratings are available, from 6VDC to 72VDC, and 100mA to 9A. Surface mount and radial lead form factors can be supplied. In keeping with FireWire power specifications, resettable fuses should be rated for at least 33VDC operation. See Figure 7 which demonstrates the use of this device. When the fault current event is cleared and current drops to a normal level, the PTC cools and its resistance also goes back to normal, restoring proper operation of the circuit. This removes the need for warranty claim that would occur if a single-shot fuse was used.

9.4 Electrostatic Discharge (ESD)

9.4.1 Origins

Damage to the FireWire port from electrostatic discharge (ESD) is generally caused by the transfer of static electrical charge from the human body to an electronic circuit. The accumulation of this charge on a body is due to frictional (triboelectric) forces, and can amount to tens of thousands of volts. It is also possible for ESD to transfer high peak voltages and currents into electronic circuits due to inductive and capacitive mechanisms in cables and ungrounded devices. An example of this phenomenon is “Cable Discharge Event” or CDE.

9.4.2 Component Sensitivity to ESD

As IC manufacturers continue to scale down device dimensions, the transistors, interconnections, and silicon layers become more susceptible to breakdown or electrical overstress damage due to ESD. Although IC manufacturers typically add some form of ESD protection to their products, this is intended to protect them during wafer fabrication and back-end assembly processes. Lately, the trend has been to reduce the level of on-chip protection in the interest of gaining higher device speeds, saving wafer space, and enhancing production processes. Therefore, user-induced ESD far surpasses the vulnerability threshold of standard semiconductor devices, and may cause any of the following effects:

Soft Failures – ESD currents can change the state of internal logic, causing data corruption, erratic operation, or system latch-up, which may require rebooting.

Latent Defects – The system may function properly for a while until the damaged component and system eventually fail (prematurely).

Catastrophic Failures – The current produced by severe ESD transients can melt silicon and semiconductor interconnects. It can also cause semiconductor junction failures due to insulating oxide breakdown. Regardless of failure mechanism, the components and the system are permanently inoperable.

ESD is characterized by very fast rise times and high peak voltages, so the protective devices within the PHY must have a similarly fast response time and voltage withstand rating. In those cases where the PHY is vulnerable to ESD, additional ESD protection devices are needed at the board level. Typically, the IEC 61000-4-2 test specification is used to verify the reliability of the system against ESD. It should be noted that this specification does not require that ESD transients be pulsed into open (unconnected) ports. This configuration is common in applications either during setup or if they are portable like Notebook computers. Passing the IEC specific my grant a false sense of security and it is recommended that after EMC testing is complete, ESD testing should be redone such that ESD is pulsed into open connectors.

9.4.3 ESD Standards

Since ESD is the most common cause of semiconductor device failures, several industry standards and specifications have been developed to test and qualify integrated circuits to determine their ESD sensitivity. These include:

a) Human Body Model (HBM) in MIL-STD-883, Method 3015
b) Machine Model (MM) in EIAJ IC121
c) Charged Device Model (CDM) in US ESD DS 5.3

All three of these models relate primarily to the manufacture and testing of an IC, but can also be applied to circuits and assemblies. The differences relate to the electrical models (charge capacitor and discharge resistor) that are used to simulate the respective transients. Details are shown below in Figure 9-7 and Table 1.

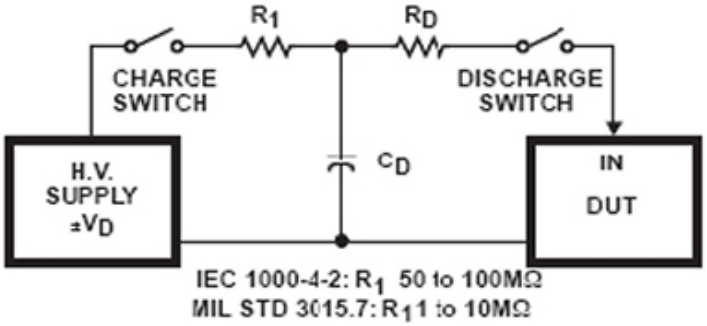

Figure 9-7—Generalized test circuit for protection devices and end products.

STANDARD	TYPE/MODE	R_D	C_D	$\pm V_D$
IEC 61000-4-2 (Level 4)	HBM, Air Discharge	330Ω	150pF	15kV
	HBM, Direct Discharge	330Ω	150pF	8kV
MIL-STD-3015.7	Modified HBM	1.5kΩ	100pF	8kV†
	Standard HBM	1.5kΩ	100pF	2kV
EIAJ IC121	Machine Model	0kΩ	200pF	400V
US ESD DS 5.3	Charged Device Model	0kΩ	NA	3kV

† Upper limit of laboratory test set.

Table 9-1—Test circuit variables for Figure 9-7

One of the most severe standards is IEC 61000-4-2 from the International Electrotechnical Commission and referenced in the EMC directive. This test specification applies to the completed system (computer, printer, etc.), as opposed to the integrated circuit as in the standards identified above. Level 4 of this test method is the highest level, subjecting the device under test to the 8kV contact discharge method (preferred), and/or a 15kV air discharge.

The designer should be aware of the ESD ratings of the semiconductors used in FireWire circuits. For example, a semiconductor device rated by a manufacturer to 2kV per MIL-STD-883 may not actually survive when subjected to the more severe IEC test method. Additionally, even if semiconductors meet some level of ESD immunity according to IEC standards, this does not imply that additional ESD suppression is not required. As mentioned earlier, real world ESD transients can exceed the peak currents and voltages as defined by the standards and can have much faster rise times. Also, a particular level of immunity may be prescribed for electromagnetic compatibility of an end product.

In a 1394 interface, protection devices should be implemented in parallel to the high-speed TPx lines to shunt ESD away from sensitive PHY input circuitry. Therefore, in addition to the appropriate level of voltage immunity, an ESD suppressor must have low capacitance to prevent the loading and distortion of data signals.

9.4.4 Protection Criteria

Generally, ESD is the greatest threat to TPx lines, having fast rise times to peak values up to tens of thousands of volts. Protection devices must limit those peaks to much lower values, and quickly clamp residual voltages to even lower levels as current is drained to ground. Furthermore, the capacitance value of an ESD suppressor must be low enough to not cause signal distortion on the TPx lines.

9.4.5 Suppressor Characteristics

Several different technologies exist for ESD suppression on TPx lines. Given the need for low parasitic capacitance, there are generally two technologies to consider: silicon protection arrays (SPAs), and polymeric suppressors. These technologies have their own unique characteristics that should be considered during ESD protection selection. Silicon Protection Arrays have the lowest turn-on and clamping voltages, so should be considered for PHYs that are particularly sensitive to ESD. Polymeric suppressors have the lowest capacitance values, so should be considered where the capacitance budget is being used up and it is necessary to minimize parasitic capacitance.

9.4.6 Representative Circuits and Products

Depending on circuit design and cost constraints, a single channel ESD suppressor might be used for each of the four TPX lines, as depicted in the lower circuit in Figure 8. Another option is to use a multichannel silicon protection array for the data lines as shown in the upper circuit in Figure 8.

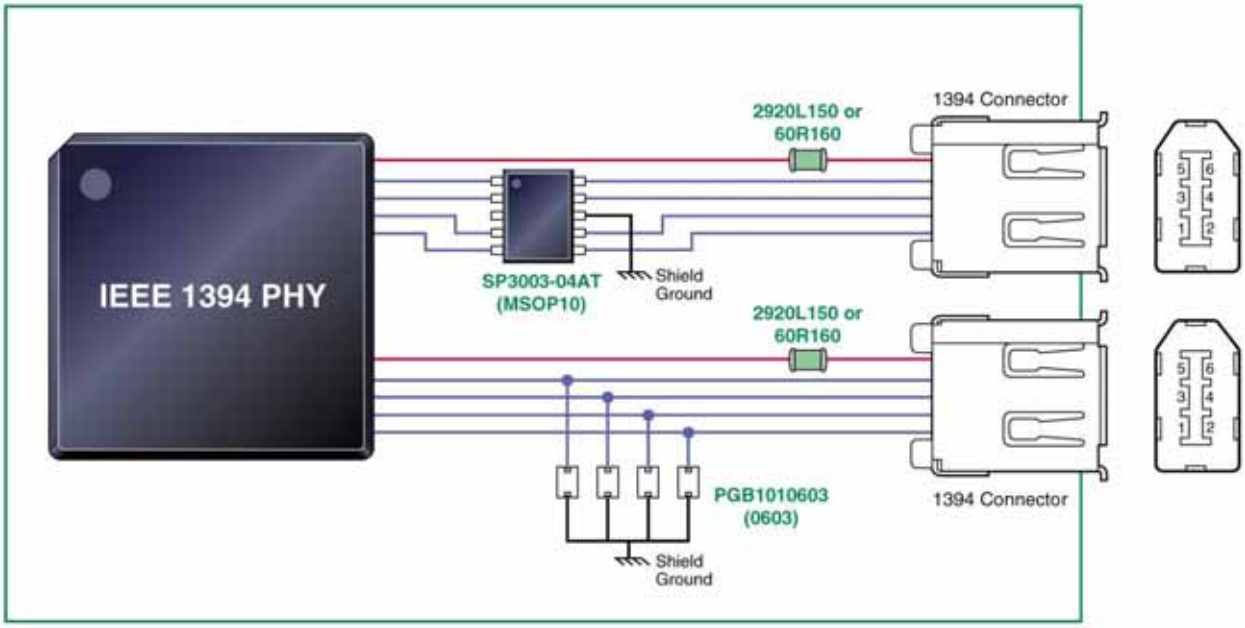

Figure 9-8—Various options for fault current and ESD protection for 1394 port.

Technology	Number of Channels	Voltage Rating, VDC	Peak ESD Voltage, V	Clamp Voltage, V	Peak Current Rating, A	Typical Cap., pF	Mounting Style
SPA[1]	4	5.5	160	20	30	0.65	SMD
Polymeric Device[2]	1	24	550	60	30	0.06	SMD

Table 9-2—Representative products for ESD protection of TPx lines.

Footnotes:

1) Littelfuse SP3003-04ATG Silicon Protection Array or equivalent.

2) Littelfuse PGB1010603 PulseGuard® suppressor or equivalent.

3) All devices tested in accordance with IEC 61000-4-2.

9.4.7 Consideration of Parasitic Capacitance Effects

Parasitic capacitance effects on signal integrity will depend on the version of 1394 that a FireWire device is designed for, and the specific data rate through that device (100MHz, 200MHz, 400MHz, etc.). As capacitance of the ESD suppressor and data rate of the port increase, the amount of distortion to leading and trailing edges of the signal pulses increases. Eventually, distortion is sufficient to interfere with data transmission.

A digital oscilloscope can be used to observe the signal with and without ESD protection on the TPx lines. Figure 8 illustrates the effects of parasitic capacitance at 100MHz and 500MHz for three different ESD suppressor technologies and a surface mount capacitor for comparison. Low capacitance devices have virtually no effect on the test signal. Its green curve is essentially superimposed on the yellow (no ESD suppressor) curve.

In order to confirm that the ESD suppressor capacitance level is appropriate, it is recommended that signal integrity testing (eye diagram, time domain reflectometry, etc.) is performed prior to the design being considered final..

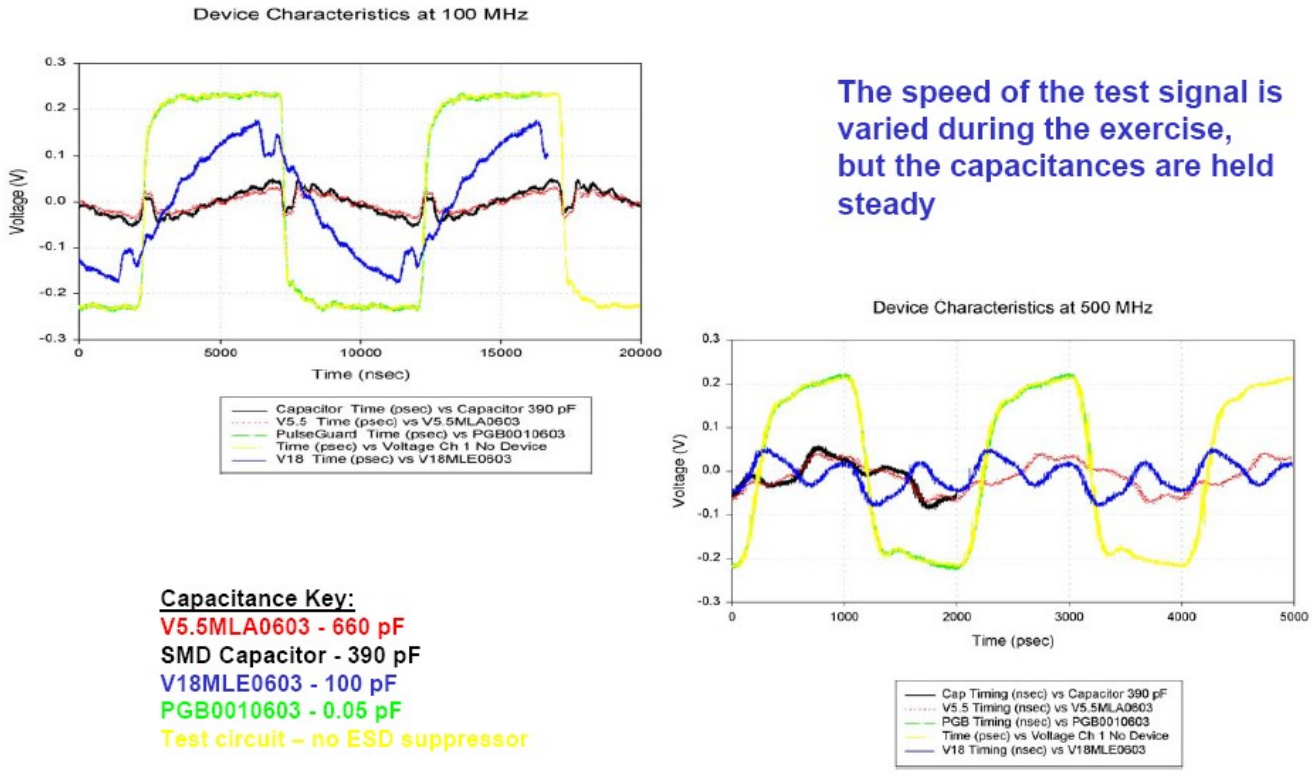

Figure 9-9—Effects of parasitic capacitance on signal integrity at 100 and 500 MHz.

Annex 10. Two port cable power distribution integrated circuit

Note, this section needs to be updated.

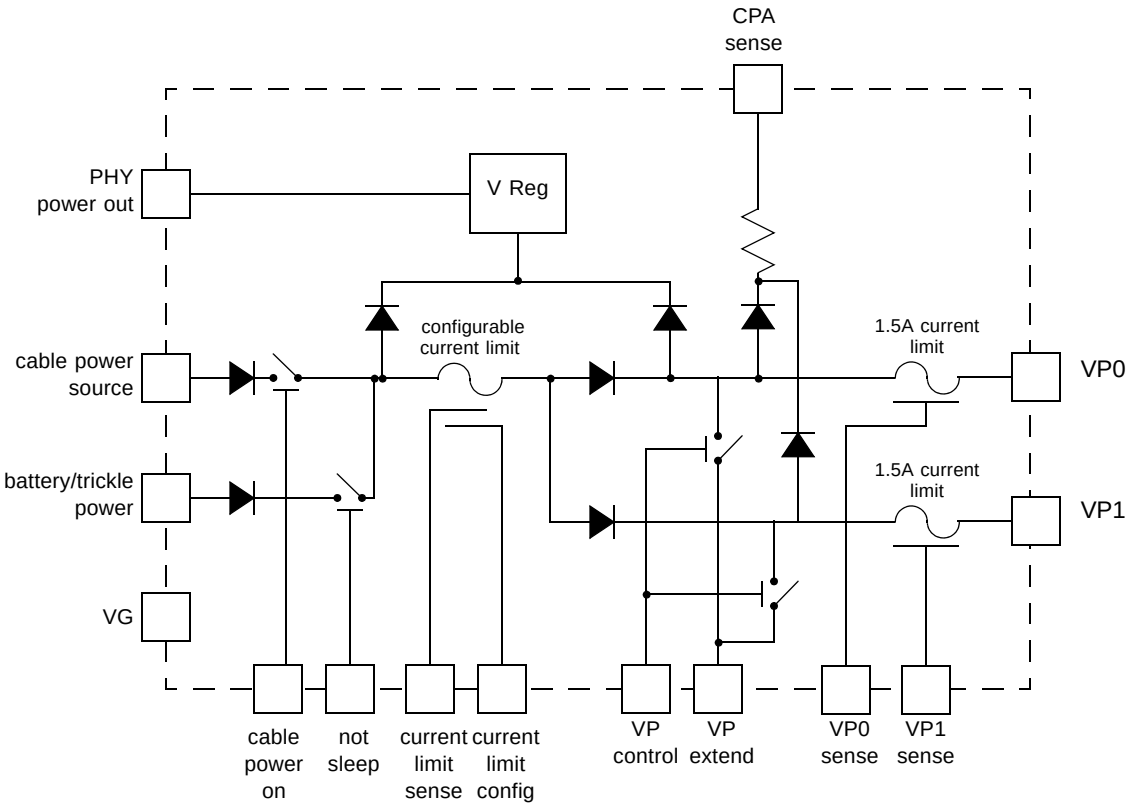

Figure 10-1—Example extendible two port power management IC

The power manager IC operates when power is applied to one or more of Cable Power Source, Battery/Trickle Power Source, VP0, VP1 and VPextend.

VP0 and VP1 switches must be able to pass current in either direction, and should operate for voltages in the range 7V - 33V.

VP control should be such that if no input is applied to it, then the switches are closed (pass current). (Suggest internal pull-down to VG)

Regulator should provide up to 1W at 3.3V to PHY Power Out if more than 7V is provided on any of the power sources.

The configurable current limit (for supply protection) should be configurable to operate between 0.5A and 3A (resistor or I2C control). The current limit sense output shall be active if the current limit has tripped.

The 1.5A current limits should respond to overload in <10ms. The VP0 and VP1 sense pins are active if the corresponding current limit has tripped.

The CPA sense output should be a simple “or” of cable power of the two ports. This signal should be conditioned to be acceptable to common PHYs.

Forward loss through diodes should be less than 0.7V... better if switches with polarity detect is used instead to minimize forward loss (otherwise the part will consume too much power).

NOTE—For 3 or more ports, two devices can be used and ganged together by connecting their VP control and VP extend pins (allowing the switched voltage to be forwarded from any VP to any other VP).

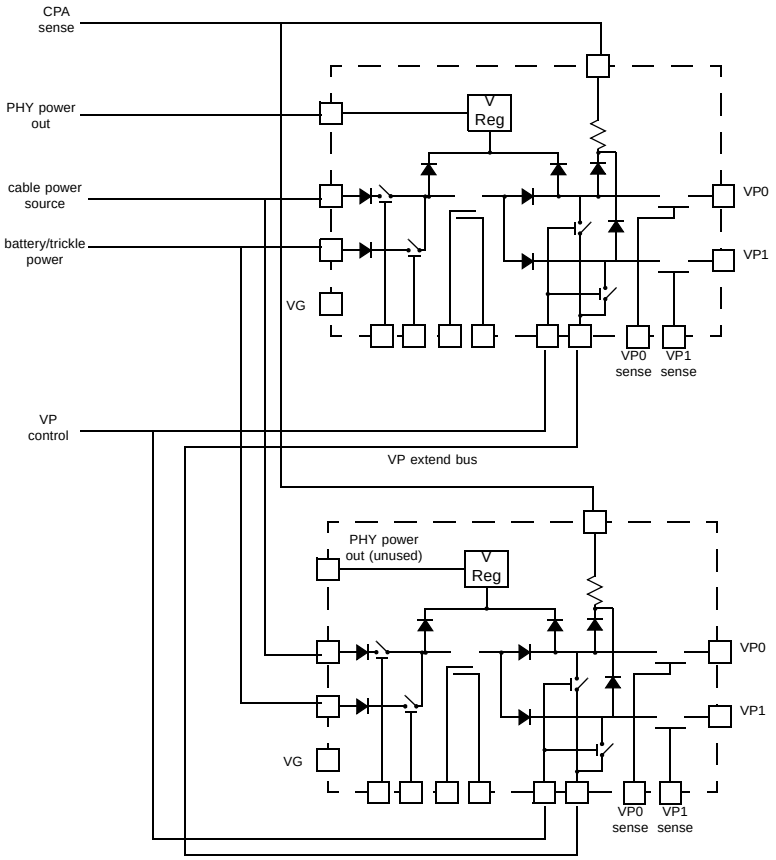

Figure 10-2—Using multiple two port power management ICs

Annex 11. References

[1] IEEE Std 1212-2001, Standard for a Control and Status Registers (CSR) Architecture for microcomputer buses

[2] IEEE Std 1394-1995, Standard for a High Performance Serial Bus

[3] IEEE Std 1394a-2000, Standard for a High Performance Serial Bus—Amendment 1

[4] IEEE Std 1394b-2002, Standard for a High Performance Serial Bus—Amendment 2

[5] 1394 TA 1999001-1, Power Specification, Part 1: Cable Power Distribution, October 5, 1999

[6] 1394 TA 1999001-3, Power Specification Part 3: Power Distribution Management, January 15, 2000

This standard shall also be used in conjunction with the following publications under development. When approved as a standard, the approved version shall apply.

[7] NCITS TR-25-1999 1-SEP-1999 Information Technology - Fibre Channel - Methodologies for Jitter Specification - MJS[1]

[8] NCITS T11/02-127v2 3-AUG-2002 Information Technology - Fibre Channel - Methodologies for Jitter and Signal Quality Specification - MJSQ

[9] 1394 Open Host Controller Interface Specification, Release 1.1, January 6, 2000

[10] 1394 Open Host Controller Interface Specification, Release 1.2.

[11] TI Application Note SLLA117, IEEE 1394 EMI Board Design and Layout Guidelines, Jose A. Cadena-Hernandez, Burke Henehan, Lee Myers, Revision 1.1.3, July 2002

[12] TI Application Note SLLA020A, Recommendations for PHY Layout, Ron Raybarman, March 1999

[13] LSI Application Note AP00041-02, FW323/FW322 Hardware Implementation Design Guideline. Revision 2, December 2002

[1] NCITS T11.2 working documents are available at http://www.t11.org.

Annex 12. Definitions and abbreviations

These definitions a liberally borrowed from the 1394b specification.

12.1 Conformance terminology

Several keywords are used to differentiate between different levels of requirements and optionally, as follows:

12.1.1 expected: A keyword used to describe the behavior of the hardware or software in the design models assumed by this standard. Other hardware and software design models may also be implemented.

12.1.2 ignored: A keyword that describes bits, bytes, quadlets, octlets or fields whose values are not checked by the recipient.

12.1.3 may: A keyword that indicates flexibility of choice with no implied preference.

12.1.4 reserved: A keyword used to describe objects—bits, bytes, quadlets, octlets and fields—or the code values assigned to these objects in cases where either the object or the code value is set aside for future standardization. Usage and interpretation may be specified by future extensions to this or other standards. A reserved object shall be zeroed or, upon development of a future standard, set to a value specified by such a standard. The recipient of a reserved object shall not check its value. The recipient of a defined object shall check its value and reject reserved code values.

12.1.5 shall: A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products conforming to this standard.

12.1.6 should: A keyword indicating flexibility of choice with a strongly preferred alternative. Equivalent to the phrase “is recommended.”

12.2 Technical glossary

The following are terms that are used within this standard:

12.2.1 8B/10B: A line code that maps 8 bit symbols to 10 bit symbols so as to achieve DC balance and bounded disparity.

12.2.2 acknowledge: An acknowledge packet.

12.2.3 acknowledge packet: An 8-bit packet that may be transmitted in response to the receipt of a primary packet. The most and least significant nibbles are the one's complement of each other.

12.2.4 acronym: A contrived reduction of nomenclature yielding mnemonics (ACRONYM).

12.2.5 active port: A connected, enabled port that is capable of detecting all Serial Bus signal states and participating in the reset, tree identify, self-identify and normal arbitration phases.

12.2.6 arbitration: The process by which nodes compete for control of the bus. Upon completion of arbitration, the winning node is able to transmit a packet or initiate a short bus reset.

12.2.7 asynchronous packet: A primary packet transmitted in accordance with asynchronous arbitration rules (outside of the isochronous period).

12.2.8 attached peer PHY: A peer cable PHY at the other end of a particular physical connection from the local PHY.

12.2.9 B cloud: A collection of B nodes and/or Border nodes in which all inter-node connections are made through Beta ports.

12.2.10 B link: A link which is capable of operating according to the specifications given in clause 14 or 15 of 1394b, and in particular issues requests appropriate to BOSS arbitration.

12.2.11 B bus: An operating bus in which all nodes are operating as B PHY's.

12.2.12 B node. A node whose PHY is operating as a B PHY.

12.2.13 B only PHY: A PHY which is only capable of B PHY mode of operation, i.e. all its ports are Beta-only ports and its link, if any, is a B link or a PIL.

12.2.14 B-parallel link: A mode of Link operation in which the PHY-Link signalling is provided to the PHY using a parallel interface.

12.2.15 B PHY: A mode of PHY operation in which all the logically connected ports are operating in Beta mode and the link, if any, is *not* configured to operate in Legacy PHY-link mode.

12.2.16 base rate: A data rate of 98.304Mbits/s ± 100 ppm. In a cable environment, all Legacy capable nodes are capable of communicating at this rate and all B capable nodes are capable of communicating at 4 * base rate.

12.2.17 BER: Bit error ratio. The ratio of the number of bits received in error to the total number of bits received.

12.2.18 Beta Mode: When a port is operating according to the specifications given in clauses 10, 11, and 13 of 1394b and, in particular, is using the 8B/10B symbol coding and obeying the BOSS arbitration protocols then the port is in Beta Mode. The speed of a port sending in Beta Mode is denoted by the β suffix (e.g., S400β)

12.2.19 Beta port: A port that is operating in Beta Mode.

12.2.20 Beta-only port: A port that is only capable of operating as a Beta port.

12.2.21 bilingual port: A port that is capable of operating both as a Beta port and as a DS port. One of the modes is selected at the time that the logical connection is made; which one is selected depends on the capabilities of the peer port.

12.2.22 Border node: A node with both
(i) either its link operating as a B link or at least one Beta port or both, and
(ii) either its link operating as a Legacy link or at least one DS port or both.

12.2.23 BOSS: An acronym for Bus Owner/Supervisor/Selector. In a B cloud, the BOSS is the node currently responsible for taking arbitration decisions. A node becomes BOSS by virtue of being the last node to transmit data in a subaction (in general, this is the node transmitting an acknowledge to a non-broadcast asynchronous packet, or the primary packet transmitter in all other cases), or by receiving a grant. The BOSS determines the end of the fairness interval and the end of isochronous intervals, notifying the other nodes. Finally, the BOSS selects the path to grant next, thereby passing the BOSS rights and responsibilities to another node.

12.2.24 BOSS arbitration: The arbitration scheme defined by this standard. The principle features of BOSS arbitration are that the node making the arbitration decision varies (see BOSS), and that arbitration requests may be overlapped with data transmission and both isochronous and asynchronous requests may be pipelined for the succeeding isochronous interval or fairness interval respectively.

12.2.25 bus ID: A 10-bit number uniquely specifying a particular bus within a system of multiple interconnected buses.

12.2.26 byte: Eight bits of data.

12.2.27 cable PHY: Abbreviation for the cable physical layer.

12.2.28 cable physical layer: The version of the physical layer applicable to the Serial Bus cable environment.

12.2.29 CAT-5: Category 5 UTP cable.

12.2.30 character: A 10-bit sequence of data sent on a connection that is operating in B mode.

12.2.31 concatenated transaction: A split transaction comprised of concatenated subactions.

12.2.32 connected PHY: A peer cable PHY at the other end of a particular physical connection from the local PHY.

12.2.33 connection: Two ports that can communicate with each other and the media between those two ports.

12.2.34 connection tone: Signal used to indicate that a port is capable of operating in Beta mode. Also confirms that a connection exists between two Beta-mode capable ports.

12.2.35 CSR Architecture: ISO/IEC 13213:1994 [ANSI/IEEE Std 1212, 1994 Edition], Information technology-Microprocessor systems-Control and Status Registers (CSR) Architecture for microcomputer buses.

12.2.36 cycle master: The node that generates the periodic cycle start packet 8000 times a second.

12.2.37 cycle start: Synonymous with cycle start packet.

12.2.38 cycle start packet: A primary packet sent by the cycle master that indicates the start of an isochronous interval.

12.2.39 data bit: The smallest signaling element used by the physical layer for transmission of packet data on the medium.

12.2.40 data character: A character used by the 8B/10B code.

12.2.41 Data Strobe (DS): A signaling method using two signals in which one signal (Data) always indicates the binary value of the data (0 or 1) and the other signal (Strobe) changes if the data stays the same during successive bit cells. This signaling method is used in IEEE Std 1394-1995 and IEEE Std 1394a-2000.

12.2.42 DC balance: The requirement that over long runs of binary symbols there be no net disparity.

12.2.43 destination: A node that is addressed by a packet. If the destination is individually addressed by a source, then it has to return an acknowledge packet.

12.2.44 disabled port: A port configured to neither transmit, receive or repeat Serial Bus signals. A disabled port shall be reported as disconnected in a PHY's self-ID packet(s).

12.2.45 disconnected port: A port whose connection detect circuitry detects no peer PHY at the other end of a cable.

12.2.46 disparity: The number of ones in a transmission character minus the number of zeros in a character.

12.2.47 doublet: Two bytes, or 16 bits of data.

12.2.48 DS mode: The form of electrical signaling and handshaking that is compatible with IEEE Std 1394-1995 and IEEE Std 1394a-2000.

12.2.49 DS only port: A port only capable of operating as a DS port.

12.2.50 DS port: A port that is operating according to Legacy specifications, in particular using DS electrical signaling and obeying the arbitration protocols defined therein. A DS only port or a bilingual port can operate as a DS port.

12.2.51 EIA: Electronic Industries Association.

12.2.52 eye diagram: Oscilloscope display of the physical layer signal which is triggered on a bit edge and shows many different bit patterns overlaid on top of each other.

12.2.53 fairness interval: A time period delimited by arbitration reset indicators. Within a fairness interval, the total number of asynchronous packets that may be transmitted by a node is limited. Each node's limit may be explicitly established by the bus manager or it may be implicit.

12.2.54 FOP: A fanout PHY. A multi-ported PHY that is attached to a PIL using the serial interface defined in clause 15 of 1394b.

12.2.55 galvanic isolation: A mechanism to avoid low frequency ground loop currents.

12.2.56 gap: A period of idle bus.

12.2.57 Hybrid bus: An operating bus that contains at least one Border node.

12.2.58 IEC: International Electrotechnical Commission.

12.2.59 initial node space: The 256 terabytes of Serial Bus address space that is available to each node. Addresses within initial node space are 48 bits and are based at zero. The initial node space includes initial memory space, private space, initial register space and initial units space. See either ISO/IEC 13213:1994 or IEEE Std 1394-1995 for more information on address spaces.

12.2.60 initial register space: A 2048 byte portion of initial node space with a base address of FFFF $F000_{16}$. This address space is reserved for resources accessible immediately after a bus reset. Core registers defined by ISO/IEC 13213:1994 are located within initial register space as are Serial Bus-dependent registers defined by IEEE Std 1394-1995.

12.2.61 initial units space: A portion of initial node space with a base address of FFFF F000 800_{16} This places initial units space adjacent to and above initial register space. The CSR's and other facilities defined by unit architectures are expected to lie within this space.

12.2.62 ISO: International Organization for Standardization.

12.2.63 isochronous: Uniform in time (i.e., having equal duration) and recurring at regular intervals.

12.2.64 isochronous gap: On a Legacy bus, the period of idle bus following an isochronous subaction that precedes asynchronous arbitration.

12.2.65 isochronous interval: A period that begins after a cycle start packet is sent and ends with a subaction indication. During an isochronous interval, only isochronous subactions may occur. An isochronous interval begins, on average, every 125 μs.

12.2.66 isochronous resource manager: A node that implements the BUS_MANAGER_ID, BANDWIDTH_AVAILABLE, CHANNELS_AVAILABLE and BROADCAST_CHANNEL registers (some of which permit the cooperative allocation of isochronous resources). Subsequent to each bus reset, one isochronous resource manager is selected from all nodes capable of this function.

12.2.67 isochronous subaction: Within the isochronous interval, either a concatenated packet or a packet and the gap that preceded it.

12.2.68 isolated node: A node without active ports; the node's ports may be disabled, disconnected or suspended in any combination.

12.2.69 jitter: Any variation in the zero-crossing time from the ideal bit pattern.

12.2.70 Legacy: Characteristics or behavior of a link, node, PHY, cable or connector defined by IEEE Std 1394-1995 or IEEE Std 1394a-2000.

12.2.71 Legacy cloud: A collection of Legacy nodes and/or Border nodes in which all inter-node connections are made through Legacy ports.

12.2.72 link: Abbreviation for the link layer. Also, the physical entity that implements the link layer. Also, one of the components connected to the PHY-link interface.

12.2.73 link layer: The Serial Bus protocol layer that provides confirmed and unconfirmed transmission or reception of primary packets.

12.2.74 logically connected port: A port whose “connected” status is TRUE. If a port is physically connected to a peer which is not powered then the port is not logically connected.

12.2.75 low-power connection signaling: Signaling, based on the exchange of very low duty cycle tones, by which the connectivity status of a port is determined. This takes place when the port is not active or is disabled.

12.2.76 module: The smallest component of physical management; i.e., a replaceable device.

12.2.77 Near-End Cross-Talk (NEXT): The noise induced in the receiving pair due to the signal on the transmitting pair on the same port. For example, the signal on the TpB pair can cause NEXT on the TpA pair of a port.

12.2.78 node: A Serial Bus device that may be addressed independently of other nodes. A minimal node consists of only a PHY without an enabled link. If the link and other layers are present and enabled they are considered part of the node.

12.2.79 node controller: A component within a node that provides a coordination point for management functions exclusively local to a given node and involving the application, transaction, link and physical elements located at that node.

12.2.80 node ID: A 16-bit number that uniquely differentiates a node from all other nodes within a group of interconnected buses. The 10 most significant bits of node ID are the same for all nodes on the same bus; this is the bus ID. The six least-significant bits of node ID are unique for each node on the same bus; this is called the physical ID. The physical ID is assigned as a consequence of bus initialization.

12.2.81 non-return to zero: (NRZ): A signaling technique in which a polarity level high represents a logical 1 (one) and a polarity level low represents a logical level 0 (zero).

12.2.82 null packet: A packet in which no data is transmitted.

12.2.83 octlet: Eight bytes, or 64 bits, of data.

12.2.84 operating speed: Nominal speed at which a port operating in Beta mode is communicating clocked information with its peer, measured in Mbits/s (before 8B/10B encoding), usually identified with the “S” notation (e.g. S100, S200, S400, etc.).

12.2.85 originating port: A transmitting port on a PHY which has no active receiving port. The source of the transmitted packet is either the PHY's local link or the PHY itself.

12.2.86 packet: a sequence of zero or more bits transmitted on Serial Bus and delimited by packet start symbol and a packet end symbol.

12.2.87 packet speed: The data rate of a packet (necessarily less than or equal to the operating speed of a Beta-mode connection used to transmit the packet).

12.2.88 path: The concatenation of all the physical connections between the link layers of two nodes.

12.2.89 payload: The portion of a primary packet that contains data defined by an application.

12.2.90 PCB: Printed circuit board.

12.2.91 peer: Service layer on a remote node at the same level. For instance a peer link layer is the link layer on a different node.

12.2.92 PIL: PHY integrated with Link. A link that uses a modified Beta port to attach to a fanout PHY using the protocols defined in clause 15.

12.2.93 PHY packet: A 64-bit packet where the most significant 32 bits are the one's complement of the least significant 32 bits.

12.2.94 physical connection: The full-duplex physical layer association between directly connected nodes. In the case of the cable physical layer, this is a pair of physical connections running in opposite directions.

12.2.95 physical ID: The least-significant six bits of the node ID. On a particular bus, each node's physical ID is unique.

12.2.96 physical layer (PHY): The Serial Bus protocol layer that translates the logical symbols used by the link layer into electrical signals on Serial Bus media. The physical layer is self-initializing. Physical layer arbitration guarantees that only one node at a time is sending data. The mechanical interface is defined as part of the physical layer. There are different physical layers for the backplane and for the cable environment.

12.2.97 physical link: In the cable physical layer, the simplex path from the transmit function of the port of one node to the receive function of a port of a directly connected node.

12.2.98 ping: A term used to describe the transmission of a PHY packet to a particular node in order to time the response packet(s) provoked.

12.2.99 PLL: Phase Locked Loop.

12.2.100 PMD: Physical Medium Dependent.

12.2.101 PMD interface: The part of an interface that is specific to single kind of interconnect.

12.2.102 Point-to-point (P2P) Packet: Special packet type that is sent on the PIL-FOP interface. This packet is used to carry data that cannot be sent as part of a normal Serial Bus packets.

12.2.103 port: The part of the PHY that allows connection to one other node.

12.2.104 primary packet: Any packet that is not an acknowledge or a PHY packet. A primary packet is an integral number of quadlets and contains a transaction code in the first quadlet.

12.2.105 primary power provider: A node that provides at least 20V and declares it’s power sourcing capabilities in its self-ID packet. Other restrictions and requirements are included in the TA Cable Power Distribution document.

12.2.106 quadlet: Four bytes, or 32 bits, of data.

12.2.107 random jitter: Jitter that comes from random sources. Characterized by gaussian statistics and unbounded variation according to the gaussian distribution function.

12.2.108 receiver eye opening: The interval in time within a bit period where the sampled data value will have a probability of error less than the specified bit error ratio (BER).

12.2.109 register: A term used to describe addresses that may be read or written by Serial Bus transactions. In the context of this standard, the use of the term register does not imply a specific hardware implementation. For example, in the case of split transactions that permit sufficient time between the request and response subactions, the behavior of the register may be emulated by a processor within the module.

12.2.110 repeating port: A transmitting port on a PHY that is repeating a packet from the PHY's receiving port.

12.2.111 request: A primary packet (with optional data) sent by one node's link (the requester) to another node's link (the responder).

12.2.112 response: A primary packet (with optional data) sent in response to a request subaction.

12.2.113 restore: The process of causing a connection in standby to return to the active state.

12.2.114 resume signal: A signal requiring the port to resume normal operations.

12.2.115 resuming port: A previously suspended port which has observed signaling other than the connection tone or has been instructed to resume. In either case, the resuming port engages in a protocol with its connected peer PHY in order to reestablish normal operations and become active.

12.2.116 run length: The length of a sequence of bits that have the same value, e.g., 1's or 0's.

12.2.117 running disparity: An estimate of the running digital sum at the end of a character sub-block, based upon the most recently transmitted (or received) character sub-block. When initialized with the same value, the running disparity and running digital sum will be equal at the end of any character sub-block. Errors in a received character stream may result in the running disparity being not equal to the running digital sum.

12.2.118 scrambler: Transmitted signals are chosen from the set of all possible symbols by using both the desired symbol and a pseudo-randomly generated offset pointer. Used to average out the spectral content of the transmitted signal to avoid strong spectral lines.

12.2.119 senior border: A unique border node in a B cloud. The senior border is the last node to originate a self-ID packet without a speed code into the cloud (i.e. repeated from a DS-mode port or generated because of a Legacy link), and is responsible for ensuring that certain Legacy gap timings are observed.

12.2.120 self-ID packet: A PHY packet transmitted by a cable PHY during the self-ID phase or in response to a PHY ping packet.

12.2.121 Serial Bus management: The set of protocols, services and operating procedures that monitors and controls the various Serial Bus layers - physical, link and transaction.

12.2.122 source: A node that initiates a bus transfer.

12.2.123 split transaction: A transaction where unrelated subactions may take place on the bus between its request and response subactions.

12.2.124 STP: Shielded, twisted pair.

12.2.125 standby: A low-power state of a Beta connection in which only low-power connection signaling takes place. No bus reset is generated as a result of a port entering or leaving the standby state.

12.2.126 standby initiator: An active port that transmits the STANDBY configuration request and engages in a protocol with its connected peer PHY to place the connection into the standby state.

12.2.127 subaction: A complete link layer operation minimally consisting of a packet transmission. The packet may be optionally preceded with bus arbitration and optionally followed by an acknowledgment.

12.2.128 subaction gap: In a Legacy cloud, period of idle bus that precedes arbitration for an asynchronous subaction.

12.2.129 suspend: To go into a low-power mode of operation while maintaining low-power connection signaling. When a port enters or leaves the suspend state a bus reset is generated.

12.2.130 suspend initiator: An active port that transmits the SUSPEND configuration request or the TX_SUSPEND signal and engages in a protocol with its connected peer PHY to place the connection in the suspend state.

12.2.131 suspend target: An active port that receives the SUSPEND configuration request or observes the RX_SUSPEND signal. A suspend target requests that all of the other active ports on the PHY become suspend initiators while the suspend target port engages in a protocol with its connected peer PHY to suspend the connection.

12.2.132 suspended node: An isolated node with at least one port that is suspended.

12.2.133 suspended port: A connected port not operational for normal Serial Bus arbitration but otherwise capable of detecting either a physical cable disconnection or a resume signal.

12.2.134 synchronization: The process of aligning the receiver's circuits to properly detect received bits and to properly detect symbol boundaries.

12.2.135 TDR: Time Domain Reflectometry.

12.2.136 TIA: Telecommunications Industry Association.

12.2.137 transaction: A request and the optional, corresponding response.

12.2.138 transaction layer: The Serial Bus protocol layer that defines a request-response protocol for read, write and lock operations.

12.2.139 unit: A component of a Serial Bus node that provides processing, memory, I/O or some other functionality. Once the node is initialized, the unit provides a CSR interface. A node may have multiple units, which normally operate independently of each other.

12.2.140 unit architecture: The specification document that describes the interface to and the behaviors of a unit implemented within a node.

12.2.141 unit interval: The nominal amount of time 1 bit takes to transmit.

12.2.142 UTP: Unshielded Twisted Pair.