Draft Standard for Low-Voltage Differential Signals (LVDS) for Scalable Coherent Interface (SCI)

Draft 1.3
November 27, 1995

Sponsor

Microprocessor and Microcomputer Standards Subcommittee of the IEEE Computer Society

Abstract: The IEEE Std 1596-1992 Scalable Coherent Interface (SCI) provides computer-bus-like services but uses a collection of fast point-to-point links instead of a physical bus in order to reach far higher speeds. The base specification defines differential ECL signals, which provide a high transfer rate (16 bits are transferred every 2 ns), but are inconvenient for some applications.

This extension to the SCI standard defines a lower-voltage differential signal (as low as 250 mV swing) that is compatible with low-voltage CMOS, BiCMOS, and GaAs circuitry. The power dissipation of the transceivers is low, since only 2.5 mA is needed to generate this differential voltage across a 100 Ω termination resistance.

Signal encoding is defined that allows transfer of SCI packets over data paths that are 4, 8, 32, 64, and 128 bits wide. Narrow data paths (4 to 8 bits) transferring data every 2 ns can provide sufficient bandwidth for many applications while reducing the physical size and cost of the interface. The wider paths may be needed for very-high-performance systems.

Keywords: backplane, bus, cable, differential, low-power, low-voltage, point-to-point, scalable, signal

---

IEEE Standards documents are developed within the Technical Committees of the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Board. Members of the committees serve voluntarily and without compensation. They are not necessarily members of the Institute. The standards developed within IEEE represent a consensus of the broad expertise on the subject within the Institute as well as those activities outside of IEEE that have expressed an interest in participating in the development of the standard.

Use of an IEEE Standard is wholly voluntary. The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five years for revision or reaffirmation. When a document is more than five years old and has not been reaffirmed, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE Standard.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. 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 regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of all concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason IEEE and the members of its technical committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration.

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

Secretary, IEEE Standards Board
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
USA

Introduction

(This introduction is not a part of IEEE P1596.3-1995, IEEE STANDARD FOR LOW-VOLTAGE DIFFERENTIAL SIGNALS FOR SCI.)

The demand for more processing power continues to increase, and apparently has no limit. One can usefully saturate the resources of any computer by merely specifying a finer mesh or higher resolution for the solution to a physical problem such as hydrodynamics or 3-D graphics. This demand leads engineers and scientists in a desperate search for more powerful and faster computers.

To economically obtain this kind of computing power, it seems necessary to use a large number of processors cooperatively. This cooperation is provided by the Scalable Coherent Interface, a high-speed packet transmission protocol that efficiently provides the functionality of bus-like transactions (read, write, lock, etc.). However, the initial physical implementations are based on ECL signal levels, which consume more power than is practical in the low-cost workstation environment. The initial specification's 1 Gbyte per second bandwidth (16-bit data path) may be overly expensive in the workstation environment. It may be more cost effective to use a narrower data path of sufficient bandwidth. The combination of a high speed transmission environment and efficient protocols can provide the link for multiple processors to cooperate in a low cost workstation environment.

The initial developers of this standard came from the IEEE Std 1596 Scalable Coherent Interface project. The ECL signal levels defined there were to get the standard implemented quickly and are practical for high-performance applications. They are less well suited, however, to using SCI in low-cost workstations. The obvious low cost solution is to integrate the transceivers into the controller and implement both in CMOS. This integration will satisfy the space and power requirements of the workstation and personal computing market.

Eventually, a lower voltage swing will be needed in order to get higher speeds than ECL signal levels can provide. This standard can provide the basis for increasing parallel signal switching frequency into the GHz range.

Stephen Kempainen, Chairman

Comments on this draft should be addressed to the chair of the working group:

Stephen Kempainen
National Semiconductor
2900 Semiconductor Drive M/S E-2595
Santa Clara, CA 95052-8090
Phone: 408-721-2836
Fax: 408-721-4785
Email: stephen@lightning.nsc.com

Questions on other Scalable Coherent Interface projects should be addressed to the Chair of the IEEE Std 1596 Scalable Coherent Interface Working Group:

David B. Gustavson
SCIzzL
1946 Fallen Leaf Lane
Los Altos, CA 94024
Phone: 415-961-0305
Fax: 415-961-3530
Email: dbg@sunrise.scu.edu

Committee Membership

The following is a list of participants in the IEEE Project 1596.3 Working Group. Voting members at the time of publication are marked with an asterisk (*).

Stephen Kempainen, *Chair**
David B. Gustavson, *Vice Chair**

Romain Agostini
Duane Anderson
Andre Bogaerts
James R. (Bob) Davis
Mike Chastain
Stein Gjessing
James Goodman
Emil N. Hahn
Craig Hansen
Mats Hedberg
David V. James
Ross Johnson
Anatol Kaganovich
Khan Kibria
Tom Knight
Michael J. Koster
Ernst Kristiansen
John Moussouris
Gary Murdock
Gurindar Sohi
William Terrell
Hans Wiggers
Danny Yeung

The following persons were on the balloting committee that approved this document for submission to the IEEE Standards Board:

Robert E. Allen
Knut Alnes
Harry Andreas
Keith D. Anthony
Wolfgang Attwenger
Harrison A. Beasley
Chris Bezirtzoglou
Janos Biri
Martin Blake
Andre Bogaerts
Charles Brill
Alice Brown
Haakon Ording Bugge
Andrew M. Cofler
Roger D. Edwards
Douglas K. Endo
Wayne Fischer
Gordon Force
Willard Graves
John Griffith
David B. Gustavson
Phillip Hughes
David V. James
Kenneth A. Jansen
Ernst H. Kristianson
Conrad Laurvick
Gerry Laws
Roland Marbot
William C. McDonald
Chris McFarland
Thanos Mentzelopoulos
Klaus Dieter Müller
J. D. Nicoud
Tadahiko Nishimukai
Michael Orlovsky
Granville Ott
Elwood Parsons
Mira Pauker
Brian Ramelson
William Ramsay
Fred U. Rosenberger
Frederick E. Sauer
V. Singer
Manu Thapar
Michael G. Thompson
Robert Tripi
Tom Vrankar
Hans A. Wiggers
David L. Wright
Yoshio Yamaguchi
Oren Yuen
Janusz Zalewski

(The following page will be replaced by the IEEE Editors after the document is approved.)
When the IEEE Standards Board approved this standard on Month dd, 199x, it had the following membership:

A. Good Person, *Chair* **Alternate B. Him**, *Vice Chair*

R. Dear Scribe, *Secretary*

Silja Theresa Spencer David Barbara Granse*

*Member Emeritus

Also included are the following nonvoting IEEE Standards Board liaisons:

How R. You
Eye M. Fine

Kristin M. Dittmann
IEEE Standards Project Editor

Contents

CLAUSE

Draft Standard for
Low-Voltage Differential Signals (LVDS) for
Scalable Coherent Interface (SCI)

1. Overview

1.1 Scope

The scope of this project, as defined by its Project Authorization Request, is as follows:

> Specify a process-technology-independent low voltage (less than 1 V swing) point-to-point signal interface optimized for IEEE Std. 1596 (SCI), using a differential driver connected to a terminated receiver through a constant-impedance transmission line. The interface will be optimized for CMOS processes, while being compatible with other IC processes, including GaAs and BiCMOS. The specification should support a transfer rate of at least 200 mega-transfers/second.
>
> In addition, the specification will define encodings for transporting SCI packets over narrow and wide data paths (4, 8, 32, 64, and 128 bits, rather than the 16 bits defined by 1596-1992) using these signals.

1.2 Objectives

The primary goal of this standard is to create a physical layer specification for drivers and receivers and signal encoding suitable for use with the IEEE Std 1596-1992 Scalable Coherent Interface in low cost workstation and personal computer applications. Other objectives include the following:

1) Technology independence. Specifications should allow designs to be implemented in a variety of integrated-circuit technologies.
2) CMOS compatible. Signal voltage levels and other specifications should be compatible with digital CMOS processes operating from 2 V through 5 V power supply levels.
3) Backplane and cable applications. Specifications should be optimized for connections between boards contained within one chassis and short (less than 5 meters) chassis to chassis interconnects. Longer connections are not prohibited, provided they meet specified signal loss and ground shift criteria for proper receiver operation. Connector and cable specifications are beyond the scope of this standard.
4) Scalable. The original 16-bit-wide SCI data path should be supplemented by 4- and 8-bit-wide data paths to support a variety of cost/performance ratios. Support for 32, 64, and 128-bit-wide data paths will also be addressed.

1.3 Strategies

The basic design strategies selected by this standard include the following:

1) Low voltage swing. To minimize power dissipation and enable operation at very high speed, low-swing (400 mV maximum) signals are specified.

2) Differential signals. Small signal swings require differential signaling for adequate noise margin in practical systems.
3) Self terminated. To minimize board real estate and costs, and to maximize clock rates, each receiver is assumed to provide its own termination resistors.
4) Uniform ground. The standard assumes that the ground potential difference between driver and receiver is kept small by the system design. The mechanism for constraining the ground potential difference is beyond the scope of this standard.

The most controversial decision was to use differential signals, which at first appears to double the number of signal lines. The pin-count overhead is actually much less than this, since reliable single-ended schemes require many more ground signals (many high-speed chips and/or backplanes provide one ground for every two signal pins) and run significantly slower. Other design benefits associated with differential signals include the following:

1) Constant driver current. The transmitter consumes a (near) constant current when driving the links; the current remains the same, but is routed in the opposite direction when the signal value changes. This simplifies the design of power-distribution wiring.
2) Constant link current. The net signaling current in a differential link is (nearly) constant, which greatly simplifies system design. The links are unidirectional and transmitters always drive a differential signal per table 3–1 or table 3–2. Reversing or stopping links would cause the net common-mode signaling current to change, creating system noise.
3) Low power. A low signal current can be used, since much of the induced noise and ground-bounce appears as a common-mode signal.
4) Simple board design. Although differential signals must be carefully routed on adjacent matched tracks, they are usually less sensitive to imperfections in the transmission line environment.
5) Low EMI. Differential signals minimize the area between the signal and the return path. In addition, the equal and opposite currents create canceling electromagnetic fields. This dramatically reduces the electromagnetic emissions.
6) Low susceptibility to externally generated noise. Though these links generate little noise, other parts of the system may. Differential signals are relatively immune to this noise.

1.4 Design models

1.4.1 Source-synchronous data

The SCI-LVDS link model assumes unidirectional operation (the driver always at one end of the link, the receiver at the other), and that a clock signal is sent along with the data as though it is just another data bit.

Both edges of the clock are used to delimit data, so the maximum transition rate of the clock is the same as the maximum transition rate of the data signals. This clock flows through the link at the same velocity as the data, and is to be used as the time reference for sampling the data.

In most applications, the received sampled data will need to be synchronized to the receiver’s local clock. If the transmitter’s clock and the receiver’s clock are independent, and thus perhaps at slightly different frequencies, occasional symbols will need to be inserted or removed from time to time in an elasticity buffer in order to maintain synchronization.

The transmission system shall ensure that the setup and hold requirements of the receiving latches are met, in order to avoid incorrect data sampling and triggering metastable states. The receiver can observe the timing of the received clock relative to its own clock in order to choose an appropriate sampling time.

By carefully adhering to these assumptions, the SCI signaling protocol becomes independent of distance or delay. The maximum distance is limited by signal skew, caused by slight differences in propagation velocity from one signal to another, and by attenuation and distortion of the signals.

Because these signals are unidirectional, it is relatively easy to reshape and time-align them in order to transmit them greater distances. However, this may introduce timing jitter which can make it impossible for the receiver to anticipate clock transitions with sufficient accuracy for reliable operation.

1.4.2 Terminated transmission lines

In addition to extending the signal encoding to parallel widths not included in IEEE 1596-1992, this standard specifies driver and receiver parameters only. However, a system must interconnect these components to be useful. The interconnect termination is specified in the receiver portion of this standard. The interconnect is beyond the scope of this standard because of the many options possible. The interconnect could include bond wires, packages and pins, printed circuit board, cables, connectors, multi-chip modules, wafer scale integration or any combination of the previous options in one driver-to-receiver signal path. This signal path is important to the correct operation of a system implementing LVDS signals and is therefore discussed in general terms in this standard.

At the high data rates this standard supports, it is important to consider the transmission line aspects of the signal path. The high frequency components of the 300 ps transition times make the parasitic reactive signal path components important. Familiar concepts such as the receiver input capacitance are overshadowed by the parasitic inductance of signal path elements that shape the waveform. If the signal delay through a signal path section is greater than the allowed minimum transition time, 300 ps, that section must be analyzed as a transmission line with associated characteristic impedance and delay. Impedance discontinuities through connectors, pins, solder pads and bond wires to the IC itself cause reflections that degrade the signal integrity.

The receiver and its package input impedance need to match the signal transmission line impedance. This serves to minimize noise-causing reflections that create data errors. Given typical CMOS process tolerances, this generally implies the use of active devices to adjust the terminating resistance until it matches an external reference. Integrating the terminating impedance onto the receiver chip complicates the design and manufacturing but the trade-off is simplified board layout and better signal integrity.

2. Document notation

2.1 Conformance levels

Several key words are used to differentiate between different levels of requirements and options, as follows:

2.1.1 expected. A key word used to describe the behavior of the hardware or software in the design models *assumed* by this specification. Other hardware and software design models may also be implemented.

2.1.2 may. A key word that indicates flexibility of choice with *no implied preference*.

2.1.3 shall. A key word indicating a mandatory requirement. Designers are *required* to implement all such mandatory requirements to ensure interoperability.

2.1.4 should. A key word indicating flexibility of choice with a strongly preferred alternative. Equivalent to the phrase *is recommended*.

2.2 Technical glossary

Many bus and interconnect-related technical terms are used in this document. These terms are described below:

2.2.1 backplane. A board that holds the connectors into which SCI modules can be plugged. In ring-based SCI systems, the backplane may contain wiring that connects the output link of one module to the input link of the next. Usually the backplane provides power connections, power status information and physical position information to the module.

2.2.2 board. The physical component that is inserted into one of the backplane slots. Note that a board may contain multiple nodes.

2.2.3 byte. Eight bits of data, used as a synonym for octet.

2.2.4 differential voltage signal. The voltage difference between the true and complementary signals from a driver with two single ended outputs whose signals always complement each other. Differential signals are also referred to as balanced signals.

2.2.5 driver. An electrical circuit whose purpose is to signal a binary state for transmitting information. Also called a generator in international standards.

2.2.6 flag. A signal used to delimit packets in parallel-signal-transmission implementations.

2.2.7 ground potential difference voltage. The voltage that results from current flow through the finite resistance and inductance between the receiver and driver circuit ground voltages.

2.2.8 idle symbol. A symbol that is not inside a packet, and is therefore not protected by a CRC. Idle symbols serve to keep links running and synchronized when no other data are being transmitted. The idle symbol also contains flow-control information.

2.2.9 jitter. Refers to the time-uncertainty of a transitioning edge recurring in a repetitive signal. This uncertainty is only with respect to other edges in that signal. Jitter is commonly measured using random bit patterns and accumulating an eye pattern to show the worst case difference in transitions.

2.2.10 LVDS. Refers to the low-voltage differential signal specifications contained in this document.

2.2.11 offset voltage. The driver offset voltage is the average dc voltage generated by the differential driver. $V_{os} = (V_{oa} + V_{ob}) / 2$

2.2.12 packet. A collection of symbols that contains addressing information and is protected by a CRC. A subaction consists of two packets, a send packet and an echo packet.

2.2.13 physical interface. The circuitry that interfaces a module's nodes to the input link, output link and miscellaneous signals.

2.2.14 receiver common mode voltage. The combination of three components: 1) the driver-receiver ground potential difference (V_{gpd}); 2) the longitudinally coupled peak noise voltage measured between the receiver circuit ground and the signal transmission media with the driver end shorted to ground (V_{noise}); 3) the driver offset voltage.

2.2.15 receiver differential noise margin high. The tolerable signal voltage variation from any source that still results in the receiver producing a logic high output state when the driver is stimulated by a logic high input. Differential noise margin high is calculated by subtracting the receiver's minimum differential high input voltage from the driver's minimum high differential output voltage. $V_{odh}(min) - V_{idh}(min)$.

2.2.16 receiver differential noise margin low. Tolerable voltage variation to guarantee that the receiver produces a logic low output when the driver is stimulated by a logic low input. $V_{idl}(max) – V_{odl}(max)$.

2.2.17 SCI. An abbreviation for Scalable Coherent Interface.

2.2.18 Scalable Coherent Interface. Refers to the Scalable Coherent Interface standard, IEEE Std 1596-1992.

2.2.19 signal line. An electrical or optical information-carrying facility, such as a differential pair of wires or an optical fiber, with associated driver and receiver, carrying binary true/false logic values.

2.2.20 skew. The difference in time that is unintentionally introduced between changing signal levels (incident edges) that occur on parallel signal lines. This difference results in an uncertain position with respect to time among parallel signals.

2.2.21 symbol. Refers to data within an SCI packet. A 16-bit unit of data accompanied by flag information. The flag information may be explicitly present as a 17th bit, or implied by the context. Symbols are transmitted one after another to form SCI packets or idles. The particular physical layer used to transmit these symbols is not visible to the logical layer.

2.2.22 sync packet. A special packet that is used heavily during initialization and occasionally during normal operation for the purpose of checking and adjusting receiver circuit timing.

3. Electrical specifications

3.1 Description and configuration

An LVDS interface, shown in figure 3–1, has a low voltage swing (400 mV single-ended maximum), is connected point-to-point, and achieves a very high data rate (500 Mbits per second per signal pair) and reduced power dissipation. Power is low because signal swings are small: minimum 2.5 mA are sent through a 100 Ω termination resistor. This sharply reduced power dissipation enables an important advance: integrating the line termination resistors, interface drivers and receivers, and the processing logic in the same integrated circuit.

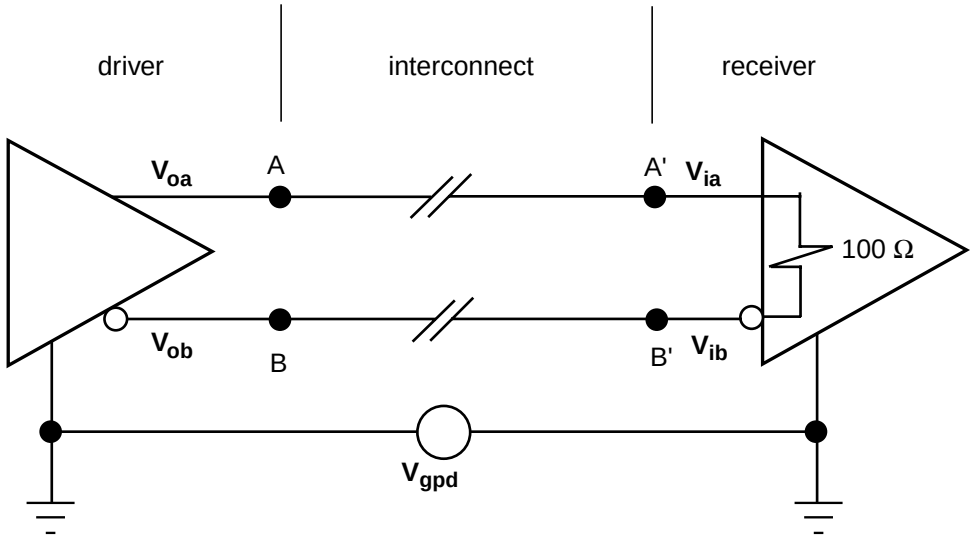

Figure 3–1: LVDS interface

Switching speed is high because the driver load is an uncomplicated point-to-point 100 Ω transmission line environment. Switching speed is also high because interface devices are all on the same piece of semiconductor material, reducing the skew due to process, temperature and supply variations between signal pairs. Connected in serial or parallel pairs, the LVDS interface forms a link used to transfer packets between integrated circuits, such as SCI nodes. For example, figure 3–2 shows circuit boards with LVDS links connected in a ring. The ring is implemented on a PCB similar in mechanical function to a multidrop bus backplane. The difference is that fewer printed circuit board (PCB) layers are needed to make the point to point connections. The PCB is simplified by eliminating the multidrop bus lines, as there is no need to route around interlayer vias used to make mechanical and electrical connections.

LVDS is independent of the physical layer transmission media. As long as the media deliver the signals to the receiver with adequate noise margin and within the skew tolerance range, the interface will be reliable. This is a great advantage when using cables to carry LVDS signals. Since all connections are point-to-point, physical links between nodes are independent of other node connections in the same system. This allows for freedom in developing a useful interconnect that fits the needs of the application.

The data path can be serial or parallel with 1, 4, 8, 16, 32, 64, or 128 bits, depending on the needs of the (see annex A—Signal encoding—for all widths except 1 and 16, which are defined in IEEE Std 1596-1992, clause 6, Physical layer).

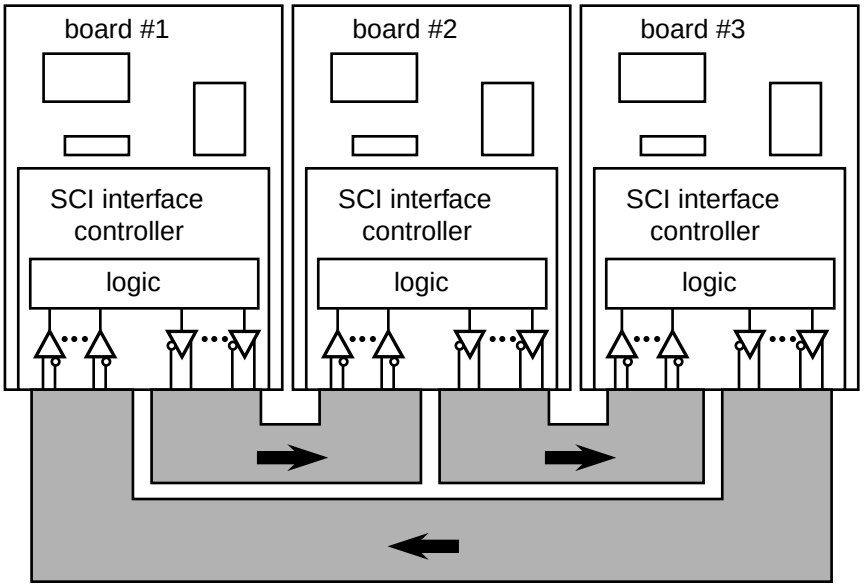

Figure 3–2: Links in SCI application (ring connection)

Electrical specifications and skew specifications are optimized for 2 V to 5 V supply voltages. The full range of semiconductor process technologies can be used to implement LVDS. It is intended that the specification be interoperable for all these technologies. The rapid trend toward reduced power supply voltage was considered in providing for signals that can be compatible with future system requirements.

The physical environment of point-to-point connections between circuit boards is further divided into two categories. The first (a general purpose link) is for circuit boards that need to operate with tolerance for V_{gpd} (table 3–1). This tolerance (approximately ± 1 V for a 2.5 V powered system) is for a general purpose system. The second (a reduced-range link) is for boards mounted on a PCB or similar environment that will guarantee less than 50 mV V_{gpd} (table 3–2). In this environment, the differential signal is reduced by reducing the driver current. This reduces the power at both driver and receiver. This is a special consideration for subsystem implementations such as IEEE Std 1596.4-1994 RamLink.

The backplane environment implies short interconnects with controlled V_{gpd}. The use of cables implies that all the skew and signal quality requirements will be met by the cables and the system designer will account for the worst case V_{gpd} and provide appropriate safeguards. The scope of the electrical specification is the differential interface of drivers and receivers. The transmission media specification, whether cables or printed circuits, is beyond the scope of this standard.

3.2 Electrical specifications

The specification for driver and receiver parameters is given in table 3–1 and table 3–2. Descriptions of these parameters are contained in the following subclauses. These specifications shall be satisfied over the product's stated power supply voltage and temperature operating range.

Table 3–1: General purpose link

Driver dc specifications:

Symbol	Parameter	Conditions	Min	Max	Units
V_{oh}	Output voltage high, V_{oa} or V_{ob}	R_{load} = 100 Ω ±1% Refer to figure 3–5		1475	mV
V_{ol}	Output voltage low, V_{oa} or V_{ob}	R_{load} = 100 Ω ±1%	925		mV
$\|V_{od}\|$	Output differential voltage	R_{load} = 100 Ω ±1%	250	400	mV
V_{os}	Output offset voltage	R_{load}=100 Ω ±1% Refer to figure 3–7	1125	1275	mV
R_o	Output impedance, single ended	V_{cm}=1.0 V and 1.4 V	40	140	Ω
ΔR_o	R_o mismatch between A & B	V_{cm}=1.0 V and 1.4 V		10	%
$\|\Delta V_{od}\|$	Change in $\|V_{od}\|$ between ‘0’ and ‘1’	R_{load} = 100 Ω ±1%		25	mV
ΔV_{os}	Change in V_{os} between ‘0’ and ‘1’	R_{load} = 100 Ω ±1%		25	mV
I_{sa}, I_{sb}	Output current	Driver shorted to ground		40	mA
I_{sab}	Output current	Drivers shorted together		12	mA
$\|I_{xa}\|$,$\|I_{xb}\|$	Power-off output leakage	V_{cc}=0 V		10	mA

Receiver dc specifications: All voltages are given with respect to receiver circuit ground voltage

Symbol	Parameter	Conditions	Min	Max	Units
V_i	Input voltage range, V_{ia} or V_{ib}	$\|V_{gpd}\|$ < 925 mV	0	2400	mV
V_{idth}	Input differential threshold	$\|V_{gpd}\|$ < 925 mV	–100	+100	mV
V_{hyst}	Input differential hysteresis	V_{idthh}–V_{idthl}	25		mV
R_{in}	Receiver differential input impedance	—	90	110	Ω

Driver ac specifications:

Symbol	Parameter	Conditions	Min	Max	Units
Clock	Clock signal duty cycle	250 MHz	45	55	%
t_{fall}	V_{od} fall time, 20% to 80%	Z_{load} = 100 Ω ±1%	300	500	ps
t_{rise}	V_{od} rise time, 20% to 80%	Z_{load} = 100 Ω ±1%	300	500	ps
t_{skew1}	$\|tp_{HLA} - tp_{LHB}\|$ or $\|tp_{HLB} - tp_{LHA}\|$, Differential skew	Any differential pair on package*		50	ps
t_{skew2}	$\| tp_{diff}[m] - tp_{diff}[n]\|$ Channel-to-channel skew	Any 2 signals on package†		100	ps

Receiver ac specifications:
Shall be maintained for (100 mV < Vid < 400 mV) throughout the receiver common-mode operating range.

Symbol	Parameter	Conditions	Min	Max	Units
t_{skew}	Skew tolerable at receiver input to meet setup and hold time requirements	Any 2 package inputs.		600	ps

*Skew measurements are made at the 50% point of the transition.
†Skew measurements made at 0 V differential (the crossing of single-ended signals).

Table 3–2: Reduced range link

Driver dc specifications: (the ac specifications of table 3–1 apply without change for driver and receiver)

Symbol	Parameter	Conditions	Min	Max	Units
V_{oh}	Output voltage high, V_{oa} or V_{ob}	R_{load} = 100 Ω ±1% Refer to figure 3–5		1375	mV
V_{ol}	Output voltage low, V_{oa} or V_{ob}	R_{load} = 100 Ω ±1%	1025		mV
$\|V_{od}\|$	Output differential voltage	R_{load} = 100 Ω ±1%	150	250	mV
V_{os}	Output offset voltage	R_{load} = 100 Ω ±1% Refer to figure 3–7	1150	1250	mV
R_o	Output impedance, single ended	V_{cm}=1.0 V and 1.4 V	40	140	Ω
ΔR_o	R_o mismatch between A & B	V_{cm}=1.0 V and 1.4 V		10	%
$\|\Delta V_{od}\|$	Change in $\|V_{od}\|$ between ‘0’ and ‘1’	R_{load} = 100 Ω ±1%		25	mV
ΔV_{os}	Change in V_{os} between ‘0’ and ‘1’	R_{load} = 100 Ω ±1%		25	mV
I_{sa}, I_{sb}	Output current	Driver shorted to ground		40	mA
I_{sab}	Output current	Drivers shorted together		12	mA
$\|I_{xa}\|$,$\|I_{xb}\|$	Power-off output leakage	V_{cc}=0 V		10	mA

Receiver dc specifications: all voltages are given with respect to receiver circuit ground voltage

Symbol	Parameter	Conditions	Min	Max	Units
V_i	Input voltage range, V_{ia} or V_{ib}	$\|V_{gpd}\|$ < 50 mV	825	1575	mV
V_{idth}	Input differential threshold	$\|V_{gpd}\|$ < 50 mV	–100	+100	mV
V_{hyst}	Input differential hysteresis	V_{idthh}–V_{idthl}	25		mV
R_{in}	Receiver differential input impedance	—	80	120	Ω

3.2.1 Driver output levels

The driver output, when properly terminated, results in a small-swing differential voltage. The relation between the single ended outputs and the differential signal is shown in figure 3–3. The differential driver is made up from 2 single ended outputs. These outputs alternate between sourcing and sinking a constant current. The differential voltage level is determined by the load resistance. The dc load seen by the driver is the receiver input impedance in parallel with the differential termination, 100 Ω, which dominates. The case where the current source is providing 4 mA is shown in figure 3–3, where the outputs are switching the current at a 50% duty cycle.

The receiver threshold limits are shown in figure 3–3, in relation to the single-ended signals that arrive at the receiver inputs. When the magnitude of the voltage difference exceeds the receiver threshold, then the receiver is in a determined logic state. For the purpose of this standard, a differential voltage greater than or equal to V_{idth}(max) is a logic high and less than or equal to V_{idth}(min) is a logic low.

Ground shift margin is built in by confining the output to a range of V_{ol} to V_{oh}. (E.g., this allows approximately 1 V of ground shift between a driver and receiver that are powered from 2.5 V supplies.) The range of allowable dc output levels for driver output voltages V_{oa} and V_{ob} is illustrated in figure 3–4. Measurement of the voltages V_{oa}, V_{ob} and the differential output voltage V_{od} is illustrated in figure 3–5.

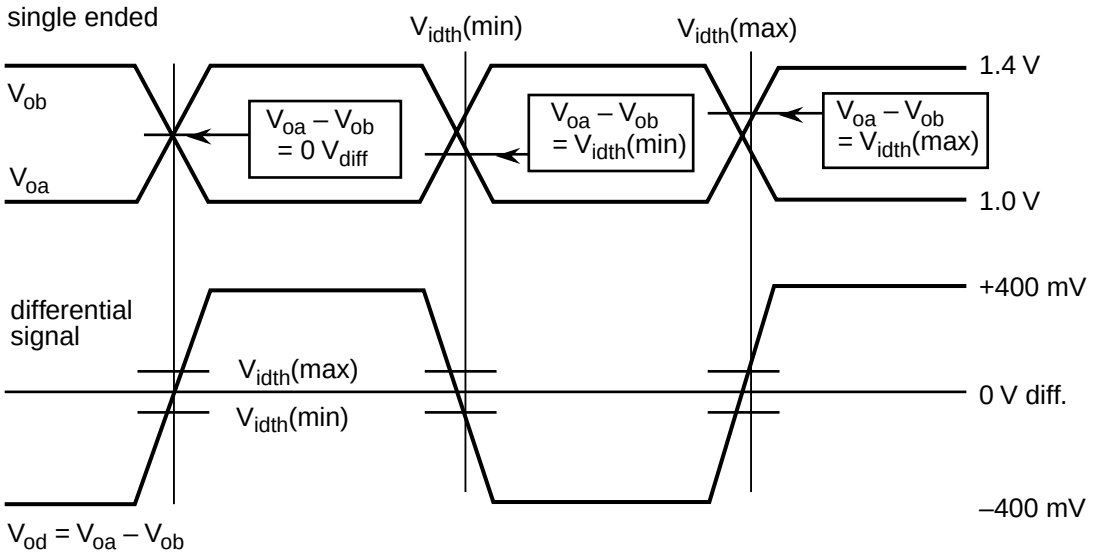

Figure 3–3: Maximum driver signal levels shown for 1.2V V_{os}

The driver output shall always be terminated in compliance with this specification. The unterminated driver output voltage shall not exceed 2.4 V. Note that the receiver may be exposed to the unterminated driver output voltage briefly when a cable from the driver is being connected to the receiver—the cable will be charged to the unterminated driver output voltage.

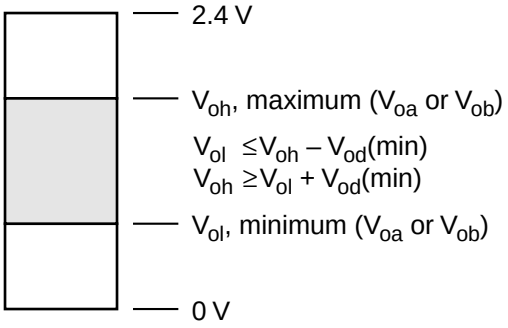

Figure 3–4: Driver signal levels

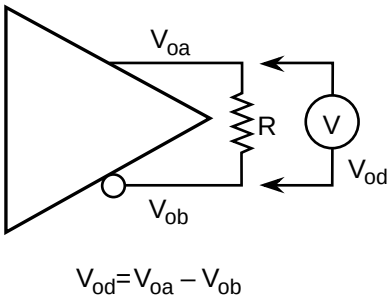

Figure 3–5: Reference circuit

The following driver dc output voltage limits refer to figure 3–4 and figure 3–5, and shall apply for a load resistance R = 100 Ω connected as shown in figure 3–5.

Ideally, the amplitude and common-mode voltage of the steady-state differential output would not change, but in practical designs, both change. The output of a driver whose differential voltage (V_{od}) and driver offset voltage (V_{os}) change when the output changes state is shown in figure 3–6. The definition for V_{od} and V_{os} is shown in figure 3–7.

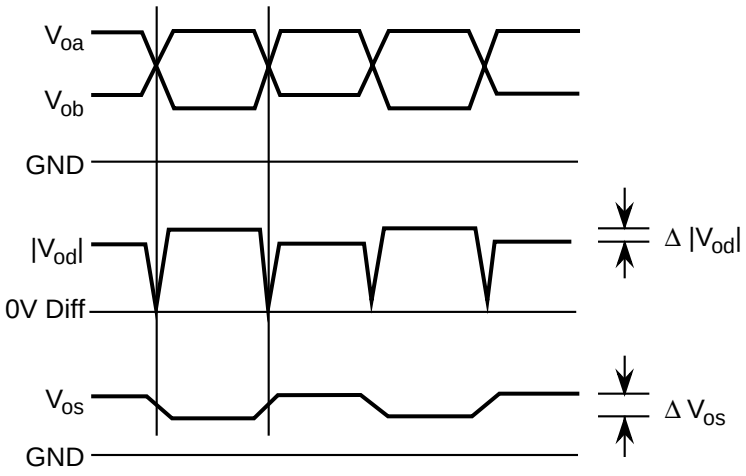

Figure 3–6: Driver signal levels

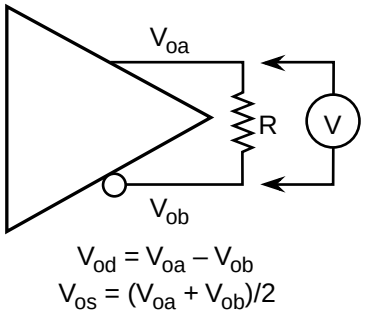

Figure 3–7: Reference circuit

The definition of ΔV_{os} and $\Delta|V_{od}|$ are explicitly stated by taking into account the varying voltage levels of the single ended outputs when in the different logic states. This can be expressed by equation 1 and equation 2.

$$\Delta V_{os} = \left| V_{os}(high) - V_{os}(low) \right| \tag{1}$$

where $V_{os}(high) = (V_{oah} + V_{obl})/2$, and $V_{os}(low) = (V_{oal} + V_{obh})/2$

$$\left|\Delta V_{od}\right| = \left| V_{od}(high) - \left|V_{od}\right|(low) \right| \tag{2}$$

where $V_{od}(high) = V_{oah} - V_{obl}$, and $V_{od}(low) = V_{obh} - V_{oal}$

The driver dc output voltage limits refer to figure 3–6 and figure 3–7, and shall apply for a load resistance R = 100 Ω connected as shown in figure 3–7.

3.2.2 Driver short-circuit specification

To ensure that the driver circuit does not damage itself or other parts of the electronics, limits on the output currents when shorted mutually and to ground are imposed.

When the driver output terminals are short-circuited to the driver circuit ground, neither current magnitude (I_{sa} or I_{sb}) shall exceed the specified value in table 3–1 or table 3–2 as appropriate. The test circuit is shown in figure 3–8.

When the driver terminals are short-circuited to each other, the current magnitude shall not exceed the specified value in table 3–1 or table 3–2 as appropriate. The test circuit is shown in figure 3–9.

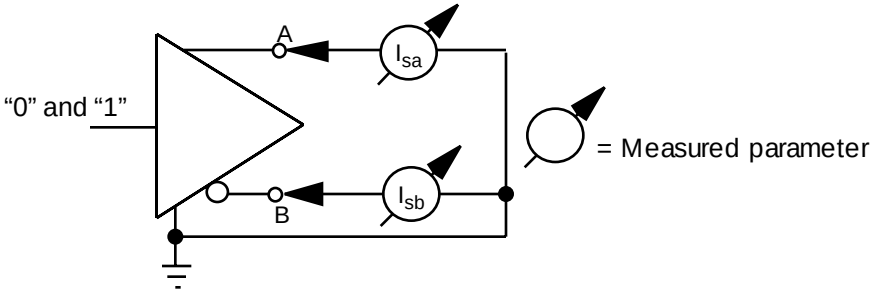

$|I_{sa}| \leq 24$ mA

$|I_{sb}| \leq 24$ mA

Figure 3–8: Short-to-ground test circuit

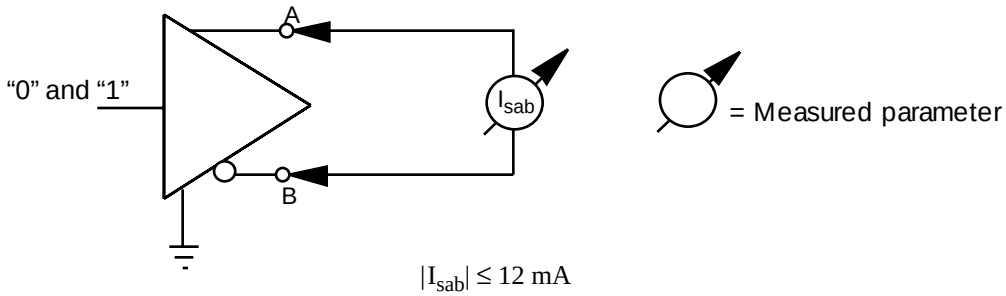

$|I_{sab}| \leq 12$ mA

Figure 3–9: Short-together test circuit

3.2.3 AC driver-output impedance

A difference between driver output impedance and signal path impedance causes reflections of incident edges arriving at the driver output from the transmission media. These waves that oppose the signal direction come from two sources: reflected signals and common-mode noise coupled onto the interconnect. To prevent common-mode noise reflected from the driver output from becoming a differential signal, the output impedance of the inverting and non-inverting outputs should be closely matched.

The upper limit of the output impedance should be as close to the transmission line impedance as possible. The lower limit on output impedance should not be much less than the impedance of the interconnect. An output impedance significantly less than the interconnect impedance will generate negative differential reflections. The upper limit can be above the interconnect impedance; however, large reflections will cause ringing and noise problems on the lines. The amplitude of the reflected signal is the product of the incident wave amplitude and the reflection coefficient (ρ), as specified by equation 3.

$$V_{reflected} = \rho \times V_{incident} \quad (3)$$

The reflection coefficient is determined by the transmission line impedance and the driver output impedance, as specified in equation 4.

$$\rho = \frac{Z_d - Z_t}{Z_d + Z_t} = \frac{2(Z_d)}{Z_d + Z_t} - 1 \tag{4}$$

Where Z_t is transmission line impedance and Z_d is driver impedance. It follows that reflections less than 10% of the incident wave will result when equation 5 is satisfied.

$$0.1 > \left| \frac{2Z_t(Z_{outa} - Z_{outb})}{(Z_{outa} + Z_t)(Z_{outb} + Zt)} \right| \tag{5}$$

Where Z_{outa} and Z_{outb} are the respective output impedance of the complementary differential drivers.

Figure 3–10 illustrates a means of measuring the reflected voltage difference due to mismatched driver output impedance. The driver would be tested when driving differential high and then driving differential low. This is a means to test the dynamic output impedance, which may differ from the static output impedance. The dynamic impedance is important in the very high frequency operation for which the driver is designed to work. Since it is difficult to test the dynamic output impedance in a production environment, this parameter can be tested, verified, and guaranteed for a design.

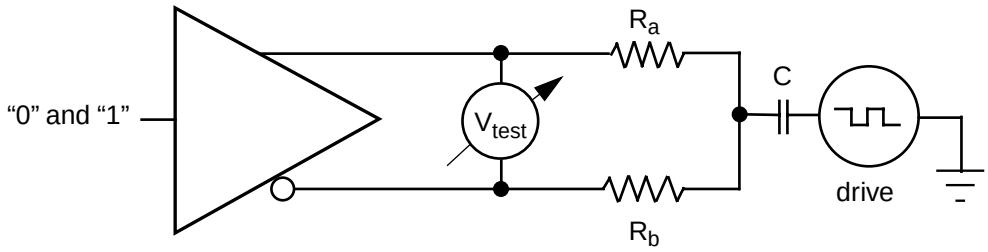

$R_a = R_b = 50\ \Omega \pm 0.1\%$
$C = 0.033\mu F \pm 20\%$
drive = 500 mV peak to peak, 10 MHz
$|V_{test}| \leq 50$ mV peak to peak
(V_{test} is the differential voltage)

Figure 3–10: Driver dynamic output impedance

3.2.4 DC driver-output impedance

Figure 3–11 illustrates a means of measuring the static driver output impedances, which have min/max as well as mismatch constraints. For each of the two possible output signal values (0 and 1), the output voltages shall be measured with driver-load common-mode (V_{cm}) voltages of 1.0V and 1.4V, yielding measured voltages of V_{oalo}, V_{oblo}, V_{oahi}, and V_{obhi}.

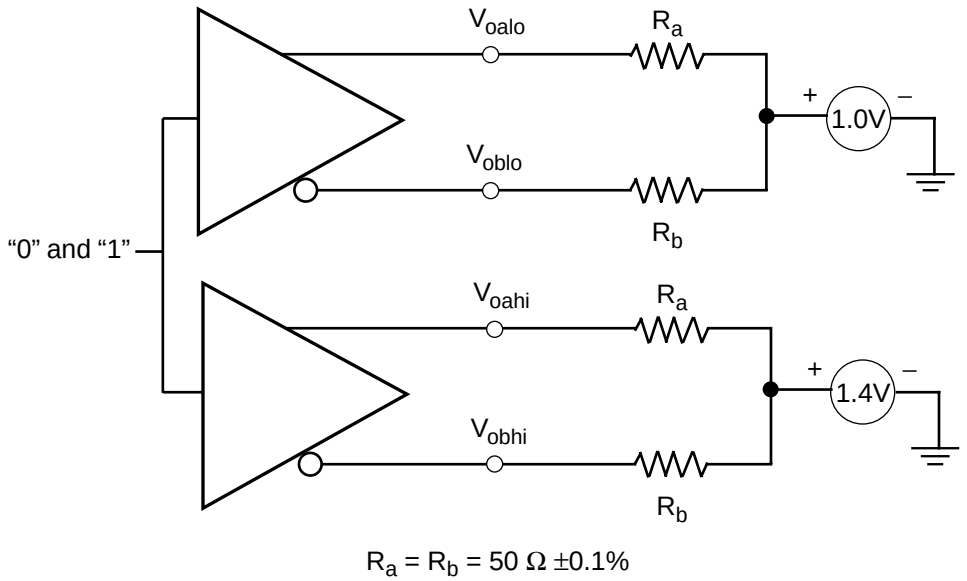

Figure 3–11: Driver static output impedance

The values of these output voltages are based on the absolute and relative impedances of the drivers. The difference in a single output voltage (for 1.0 V and 1.4 V output-load voltages) is affected by the value of the driver's output impedance on this signal. The difference in the differential output voltages (for 1.0 V and 1.4 V output-load voltages) is affected by the matching of the driver's *a* and *b* output impedances. These values shall be within the constraints specified in table 3–3.

Table 3–3: Static loaded-output voltage constraints

parameter	Corresponds to	units	minimum	maximum
V_{oahi}-V_{oalo}	R_o, terminal a	millivolts	178 (based on 40 ohms)	295 (based on 140 ohms)
V_{obhi}-V_{oblo}	R_o, terminal b			
$\lvert(V_{oahi}-V_{obhi}) - (V_{oalo}-V_{oblo})\rvert$	10% matching of reflection coefficients	millivolts	0	20

3.2.5 Driver power-off leakage current

The driver output leakage currents (I_{xa} and I_{xb}) are measured under power-off conditions, V_{cc}=0 V, as shown in figure 3–12. With the voltage on the driver output terminals between 0 V and 2.4 V, with respect to driver common, these currents shall not exceed the value specified in table 3–1 or table 3–2.

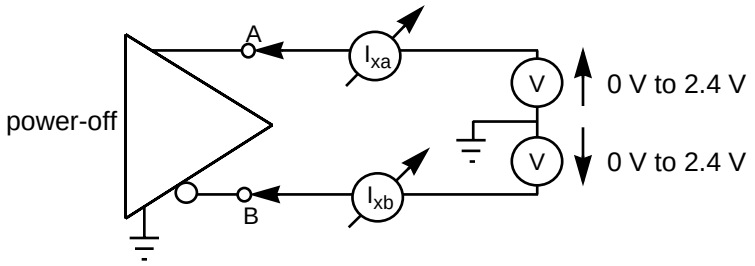

Figure 3–12: Driver power-off leakage current test circuit

3.2.6 Receiver input levels

The receiver input signal is measured differentially, figure 3–13. The receiver output state is determined by a differential input signal greater than $+V_{idth}$ or less than $-V_{idth}$, within the permitted V_i range. The table 3–1 and table 3–2 receiver specifications are the same except for the common mode operating range and the on-chip termination tolerance. The receiver common mode voltage input range, over which it must meet all other specifications, is reduced in table 3–2 because that is intended for operation in a well-controlled environment. The termination is allowed to have greater variance because it is intended to operate in a more controlled environment with less common-mode noise. For simplicity, the remaining receiver specification discussion here will apply directly to the general purpose specification. The analogous explanation will apply to the table 3–2 specification by substituting the appropriate numerical limits.

The ability to accept voltages outside a V_i range is desirable because it increases noise immunity to ground potential difference and interconnect-coupled noise. The upper limit to the differential swing is given to ensure that receiver skew specifications are maintained for this range of input signals throughout the receiver common mode range. The range of allowable dc input levels for receiver input voltages, V_{ia} and V_{ib}, is illustrated in figure 3–14. Measurement of the voltages V_{ia}, V_{ib}, and the differential input voltage V_{id} is illustrated in figure 3–15.

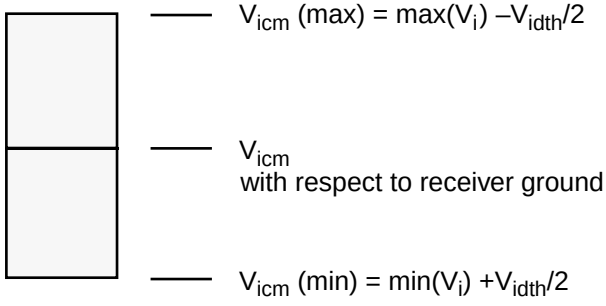

Figure 3–14: Receiver signal common mode levels, table 3–1

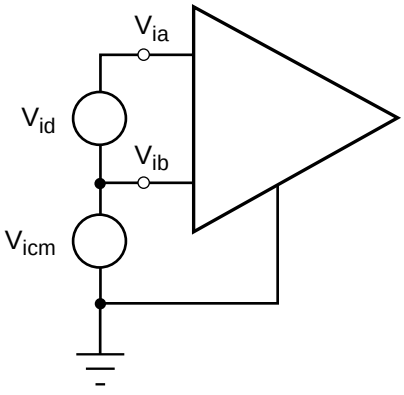

Figure 3–15: Reference circuit

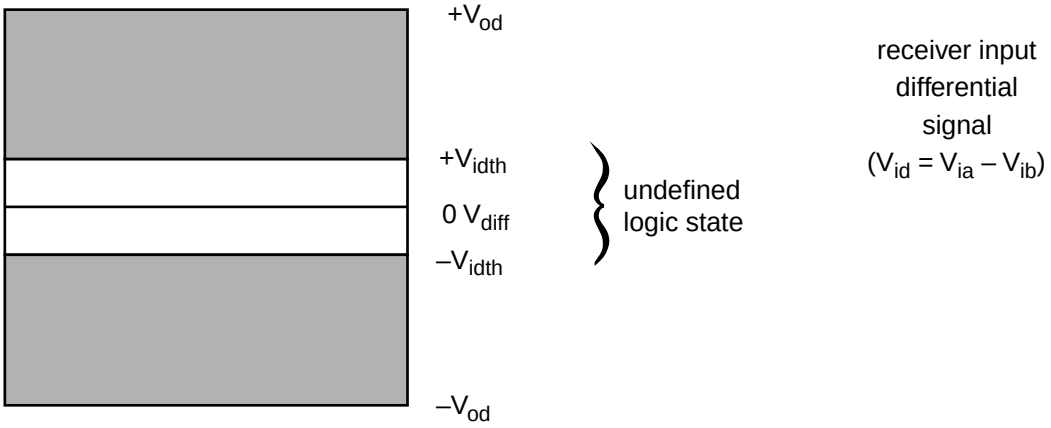

Figure 3–13: Receiver signal levels, for table 3–1

The signal common-mode level for table 3–1 is shown in figure 3–14. The receiver common-mode input voltage, V_{icm}, will be an alternating voltage depending on three superimposed conditions: the driver output condition, voltages induced on the interconnect, and reflections caused by common-mode termination imperfections.

This voltage can be expressed by accounting for the varying levels during both logic states. Equation 6 expresses the relationship of the four input voltages resulting from the three conditions previously stated.

$$V_{icm} = \frac{V_{ia} + V_{ib}}{2} \tag{6}$$

3.2.7 Receiver input impedance

A differential termination resistor (connected across the receiver inputs) should be integrated onto the receiver die. Integration is feasible because the power dissipated in the termination is less than 2 mW per receiver, and a single value can match the point to point interconnect characteristic impedance. Since an alternate termination scheme would be more useful in certain applications, there are modifications allowed to this requirement. For example, a common mode termination to 1.2 V would be desirable when the ground potential difference between receiver and driver can not be guaranteed to meet the limits specified in table 3–1. This allows the ground potential difference to be divided between the driver and the receiver. This common mode termination can still use integrated termination resistance values of 50 Ω from each input to a common externally accessible pin. This does require an additional pin as shown in figure 3–16; however, all parallel channels can bus this common-mode termination together. Figure 3–16 shows both alternatives and defines the differential input impedance as the impedance measured across the receiver inputs.

Note that the receiver input capacitance should be designed to be as low as possible. Details of integrating the termination impedance are left to the circuit designer’s discretion, but the termination should not limit the high frequency, 250 MHz, operation of the receiver.

The unpowered receiver impedance is not specified, since active circuits (which require power) are expected to be used to implement the on-chip termination resistance. However, when unpowered, the magnitude of the receiver’s leakage current (the sum of the two input currents) shall not exceed 1.0 mA.

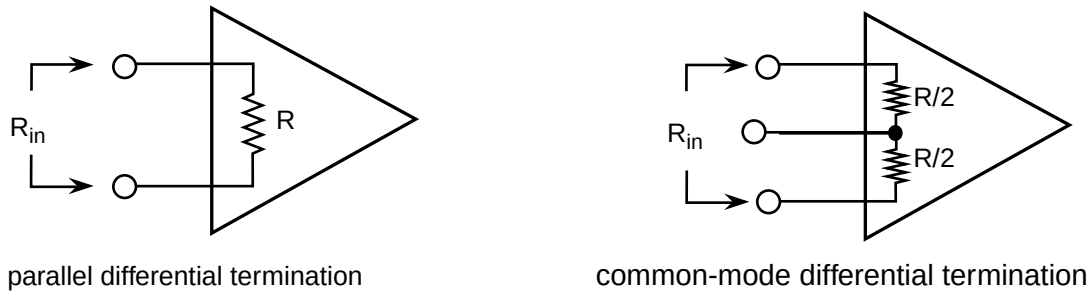

Figure 3–16: Receiver input impedance

3.2.8 Receiver threshold hysteresis

The threshold hysteresis is important in receiver design to eliminate the possibility of oscillating receiver output when the differential input is undefined. The undefined input can occur when the receiver inputs are unconnected, when the connected driver is powered down, or when transitioning between defined values. The 25 mV minimum hysteresis means that an input signal must change by more than this value to change the receiver output state. A known output condition for an open or shorted receiver input (failsafe) is implementation dependent and beyond the scope of this standard.

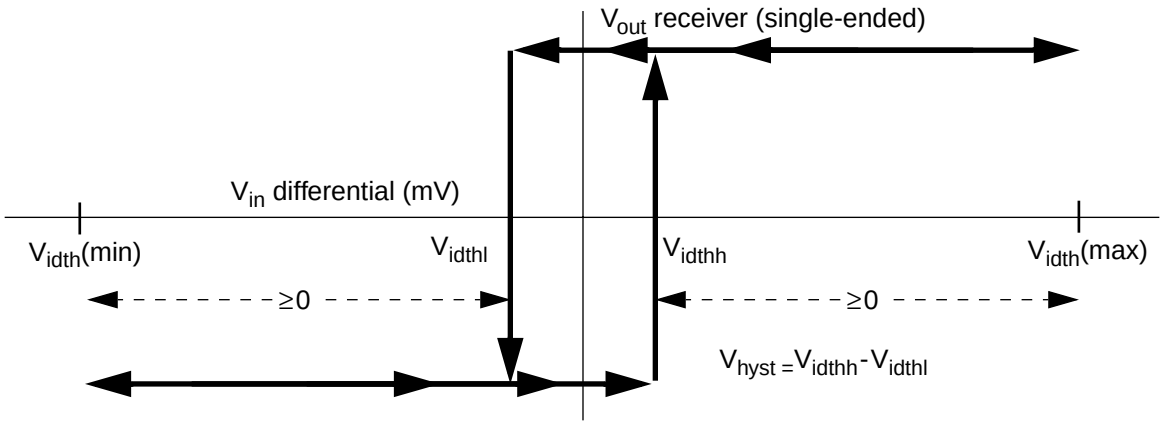

Figure 3–17: Receiver hysteresis

3.3 AC specifications

A basic goal for this specification is to preserve the high transfer rate of the ECL physical layer (1 bit /2 ns for each differential signal pair) while drastically reducing power dissipation. The skew and transition time limits specified in this section correspond to those of the base SCI standard.

3.3.1 Driver transition times and undershoot

The 500 ps maximum transition time specified in this section is equivalent to a 0.3 V/ns slew rate for a 250 mV differential signal. This is the minimum guaranteed slew rate for the signals specified in this standard. The transition time is most critical through the receiver threshold region because the high speed comparator design needs unambiguous inputs to function efficiently. The 0.3 V/ns slew rate in the threshold region shall guarantee reliable receiver switching.

The fast transitions contain high frequency components that directly affect the electromagnetic radiation created by the signal. Normally these high frequencies would create EMC problems. The differential signal advantage is produced by single-ended signals simultaneously rising and falling. These edges will generate equal and opposite electromagnetic fields that cancel each other and reduce generated fields and radiation. Therefore, it is important to have both single ended channels switching at the same time and at the same slew rate; i.e. no skewed single-ended transitions.

Undershoot, overshoot, and fast rise times can generate noise, crosstalk, and electromagnetic interference. Although the driver specifications in figure 3–18 and figure 3–19 reduce these problems, careful transmission-line design is important to maintain signal integrity.

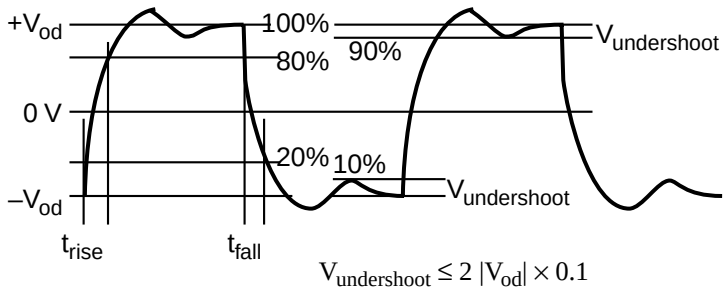

Figure 3–18: Driver waveform

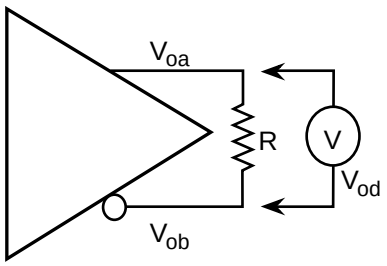

Figure 3–19: Reference circuit

3.3.2 Receiver common-mode rejection

The receiver specified for table 3–1, the general purpose specification, is intended to operate over a common-mode voltage range that will allow for about ±1 V ground-potential difference between the driver and receiver power supplies. The common-mode input voltage, V_{icm}, is the average of V_{ia} and V_{ib} measured with respect to the receiver ground potential. The equation expressing this value is given in 3.2.6. The receiver specified in table 3–1 and table 3–2 must maintain the sensitivity and skew specifications throughout this common-mode voltage range (see figure 3–14 for V_{icm}(max) and V_{icm}(min) definitions).

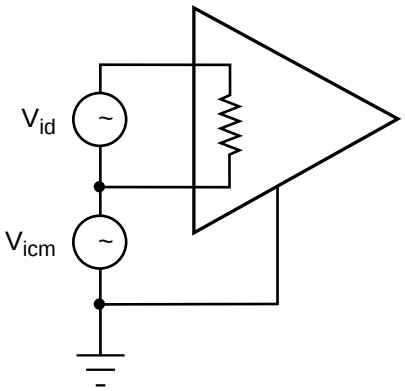

Figure 3–20: V_{icm} reference circuit

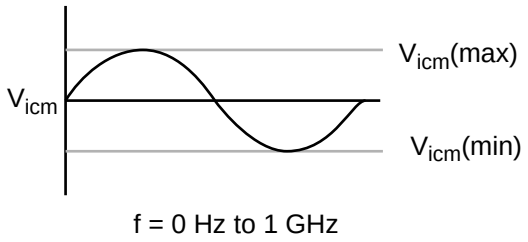

Figure 3–21: V_{icm} input waveform

3.4 Skew specifications

System interconnects comprised of parallel signals must account for skew. This standard includes skew specifications because limiting skew is important for correct data transfers in a system. For example, the SCI physical system can be thought of as a data pipeline. Bytes of data can be lined up in the physical layer, a virtual FIFO. Since more than one transmission can be contained in the data path at any time, it is not the total delay time but the timing skew that is important to reliable SCI data transfer. The skew is extremely important because it directly affects the sample window (setup and hold time) available to the receiver logic. These skew specifications are based on a total allowable skew tolerance that will still provide an adequate sample window for receiver capture logic. Since the bit width is nominally 2 ns, reliable data transfer assumes that a 600 ps skew will be the maximum allowed.

The physical layer skew is defined for the purpose of this document as the difference in time that is unintentionally introduced between changing signal levels (incident edges) that occur on parallel signal lines. This difference results in an uncertain position with respect to time between parallel signals. Jitter is not specified in this standard because the skew specification accounts for jitter as well. Jitter is the error of the time of transitions occurring in a serial transmission line. The range of parallel skew includes uncertainties that would be called jitter in a serial specification.

A parallel clock signal always accompanies SCI data signals. This clock is specified to be 250 MHz and to have a duty cycle that is greater than 45% and less than 55%. Since it would be impractical to force the circuit designer to vary the clock duty cycle to the min and max limits, testing would be done with the clock as generated. It would be tested under the same conditions as the data signals, since it is implied that they are on the same die. The requirement is for the clock to comply with the 45% to 55% duty cycle first, and then test all skew parameters to that clock. Since “any two signals” is specified for skew measurements, the relation of all signals to clock is implied because “any two” includes all pairs for skew measurement.

From the perspective of the receiver, the maximum clock-to-data signal skew is more important than the maximum data-to-data signal skew. However, from a manufacturing perspective, minimal-skew design techniques are unlikely to treat the clock signal as special, therefore skew measurements are not based upon it. Note that sophisticated receivers can measure min/max data-signal skews and dynamically adjust the clock signal delay, to reduce the effective clock-to-data skew to a little more than half of the “any-two” skew specification.

However, testing for the range of skew between the first and last signal may be inconvenient, so a more constraining specification could be used to simplify testing: require all data signals to be within ±300 ps of the clock.

The physical transmission mechanisms do not differentiate between the clock and other signals, so any given transmission link might happen to result in the clock arriving as the latest or the earliest signal. Many of the receiver technologies appropriate for very high frequency operation do not use the incoming clock to sample the incoming data, which must be resampled and shifted to the local clock domain. Such receivers merely monitor changes in clock phase as the source clock drifts relative to the receiver’s own clock, and thus do not care whether the clock transition is centered relative to the data transitions or at one extreme or another. Additionally, at these baud rates, receivers will often need to incorporate automatic deskewing circuits to function in real world environments. For example, SCI protocols provide sync packets for dynamic skew compensation (see B.1). The circuit designer will need to use dynamic skew compensation if the cables and connectors used in his systems cannot assure meeting the 600 ps maximum skew specification.

The single-ended pairs that make up the differential signal are shown in figure 3–22. This picture shows the single ac receiver specification called t_{skew}. Here, t_{skew} is pictured as only one receiver differential input to another receiver differential input. It is the total allowable skew among parallel channels. The receiver must be able to correctly sample the logic state on all parallel channels when there is this much difference between the transitioning signals as they arrive at the inputs to the high speed comparator. The t_{skew} includes

that resulting from driver, interconnect, packaging, induced noise, etc. Differential signals $V_{id}[m]$ and $V_{id}[n]$ would be measured for the difference in delay at the receiver input, using the test circuit shown in figure 3–15.

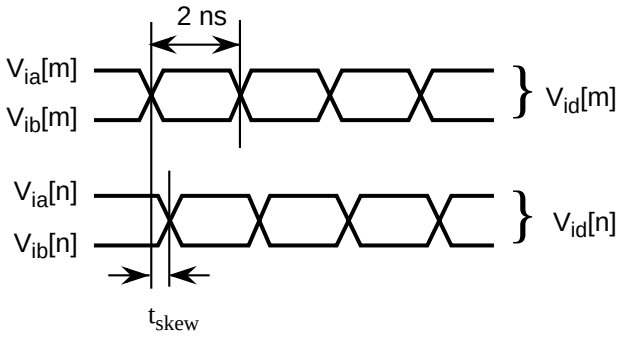

Note: $V_{id}[m]$ and $V_{id}[n]$ are any two differential signals, including clock and flag, that are present at receiver inputs.

Figure 3–22: t_{skew} diagram for receiver inputs

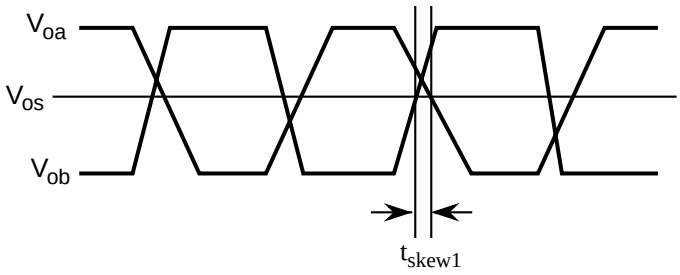

Any two complementary single-ended signals

Figure 3–23: Skew1 diagram

Since LVDS receivers are expected to be integrated on VLSI parts, it is difficult to measure skew for difference in delay through a comparator. For this reason, the receiver skew is specified as a maximum tolerable timing difference (as observed on the receiver's inputs) for which the data will be sampled correctly. This specification implies that the internally sampled data values are accessible during testing.

When generating differential signals, two skews are important. Skew1 is the skew between the high to low and low to high transitions of complementary single ended channels, figure 3–23. This skew can be the result of different propagation delays between complementary drivers, or different slew rates of the driver outputs. It is always measured at the V_{os} as defined in figure 3–7 for the single ended signals. This skew creates EMR as discussed in 3.3.1.

$$t_{skew1} = \left|tp_{HLA} - tp_{LHB}\right| \quad or \quad \left|tp_{HLB} - tp_{LHA}\right| \tag{7}$$

Where $tp_{HLA/B}$ and $tp_{LHA/B}$ are the propagation delays on outputs A and B for high to low and low to high.

Skew2 is the skew between any differential signals as they can be measured at the driver output. It results from circuit mismatches for complementary drivers and layout or packaging differences. As shown in figure 3–24, it is measured between parallel channels. Skew2 is always measured between any 2 parallel signals, at the 0 V differential point. If $tp_{diff}[i]$ is the differential delay of V_{od} through the LVDS driver *i*, and assuming that the (probably inaccessible) inputs to the driver are simultaneous, then

$$t_{skew2} = \left| tp_{diff}[m] - tp_{diff}[n] \right| \tag{8}$$

where *m* is any one of the parallel channels and *n* is any other channel.

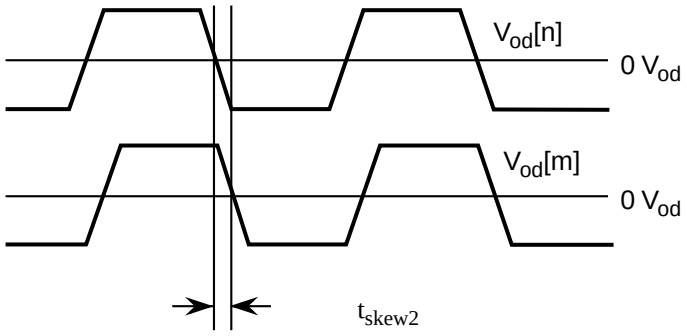

Figure 3–24: Skew2 diagram, measured between any parallel channels

The diagram in figure 3–25 shows a representative breakdown of a typical SCI signal path. The backplane could also represent a cable segment of the signal path. An estimate of how the skew budget could be allocated for each of the signal path elements up to the package input is given in table 3–4. The chip inside the package also has to tolerate the additional package-to-chip wiring skews.

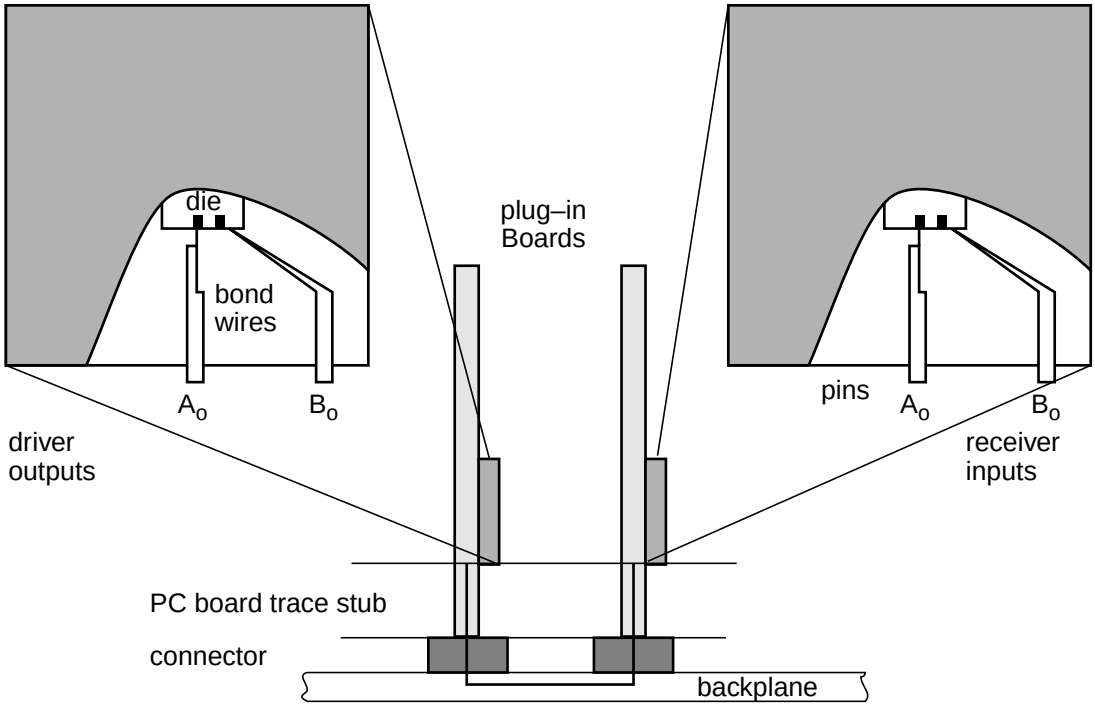

Figure 3–25: Point-to-point connection segmented for skew allocation

Table 3–4: Representative skew breakdown

signal path	from	to	skew max
package	internal signals	package's driver pins	100 ps
PCB	driver pins	input of first connector	50 ps
connector	input of first connector	output of first connector	25 ps
media	output of first connector	input of second connector	350 ps
connector	input of second connector	output of second connector	25 ps
PCB	output of second connector	receiver pins	50 ps
			total = 600 ps

Annexes

Annex A
Bibliography

(informative)

This document has been developed with point-to-point interconnects, such as the following standard, in mind:

[B1] ANSI/IEEE Std 1596-1992, Scalable Coherent Interface (SCI) (or IEC/ISO DIS 13961).[1]

[1]ANSI/IEEE publications are available from the Institute of Electrical and Electronics Engineers, Service Center, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA.

Annex B
SCI signal encoding

(normative)

B.1 SCI symbol encoding

The SCI encoding specifies how packet types, packet lengths, and idle symbols are uniquely identified. Although the logical encoding is specified in terms of 16-bit symbols, the physical encoding layers support several data-path widths. Distinct physical encoding layers can be supported without changing the logical protocols, if they define how conversions between the physical and logical encodings are performed.

For 16-bit-wide links, one 16-bit SCI symbol is transmitted in each data-transfer period. In addition, a clock signal is needed to define symbol boundaries (the data should be stable when sampled), and a flag signal is used to locate the starting and ending symbols of packets.

The zero-to-one transition of the flag signal is used to mark the beginning of each packet, and the one-to-zero transition of the flag signal specifies the approaching end of each packet. The flag signal returns to zero for the final 4 symbols of send packets and for the final symbol of an echo packet as illustrated in figure B.1. A zero always accompanies the CRC of any packet, so the zero-to-one transition can be used to identify the start of the next packet (even when there is no idle symbol between them).

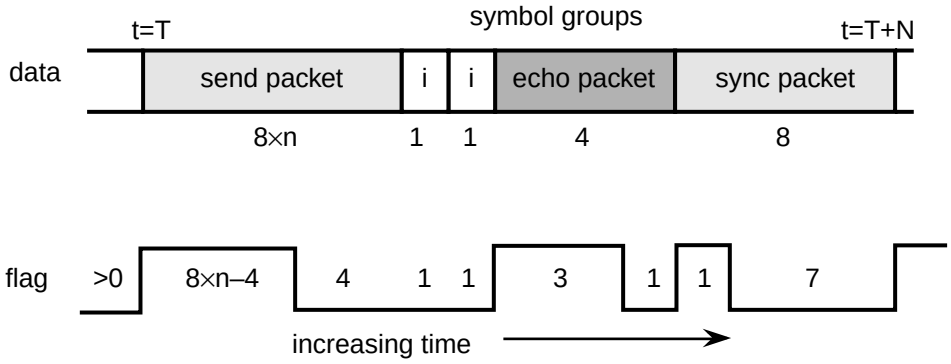

Figure B.1—Flag framing convention

The sync packets are used for initial synchronization of the physical receivers and for dynamic compensation of skew. Sync packets are always a pattern of one-bits the full width of the link, followed by zero-bits the full width of the link. Note that for P18 and wider encodings, the clock may have either a zero-to-one or a one-to-zero transition at the point in the sync packet where all other signals have a one-to-zero transition. (Packets may start at either clock transition in these wider encodings.) To simplify the design of the SCI interface hardware, the idle symbols are always 16 bits wide or the full width of the data path depending on which is larger.

It should be noted that the 1- and 16-bit encodings, which are specified in IEEE 1596-1992 and therefore not included in this annex, can also use this LVDS differential signaling standard.

B.2 Narrow parallel encoding

The 16-bit SCI symbols need not be sent in a single physical-clock-signal transition; multiple data-transfer cycles can be used. Encoding for use on physical links that are narrower than the 16-bit SCI logical symbol width use the polarity of the clock signal or extra transitions on the flag signal to mark the beginning of the logical symbol.

B.2.1 Parallel 8/10 (P10) encoding

On a byte-wide (8-data-bit) interface, half of an SCI symbol is sent in each of two data-transfer intervals. The clock signal changes before each subsymbol; the low and high clock-signal values identify the first and last subsymbol respectively. The flag line transitions occur at most once per symbol, as illustrated in figure B.2.

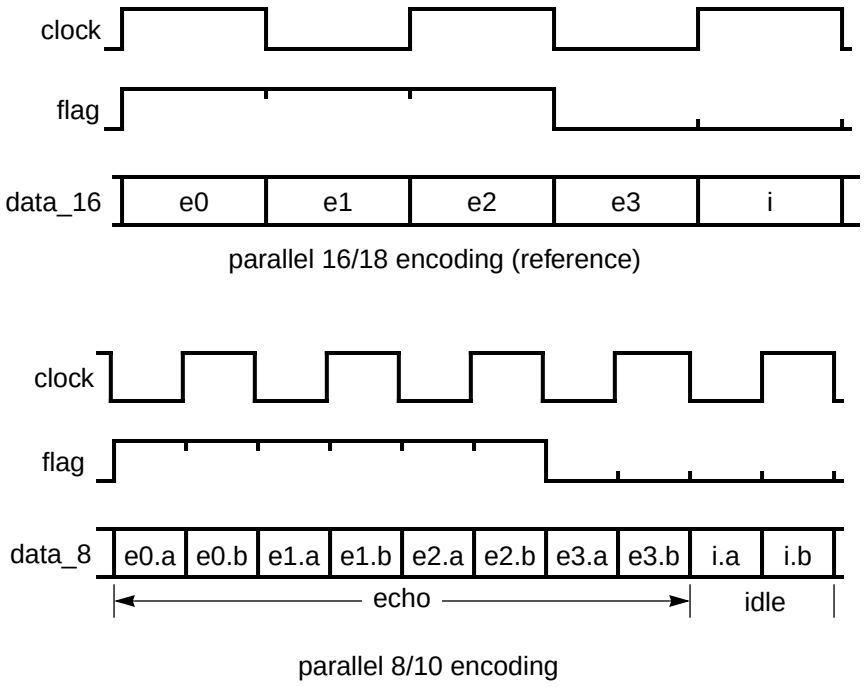

Figure B.2—Parallel 8/10 (P10) encoding

In this illustration: *e0.a* and *e0.b* bytes are the most- and least-significant bytes of the *e0* (echo-packet) symbol; *e1.a* and *e1.b* bytes are halves of the *e1* symbol; *e2.a* and *e2.b* are halves of the *e2* symbol; *e3.a* and *e3.b* are halves of the *e3* symbol; *i.a* and *i.b* bytes are the most- and least-significant bytes of the *i* (idle) symbol.

The sync packet is 16 clock-periods long (8 symbols), to simplify conversions between P18 and P10 encodings, as illustrated in figure B.3.

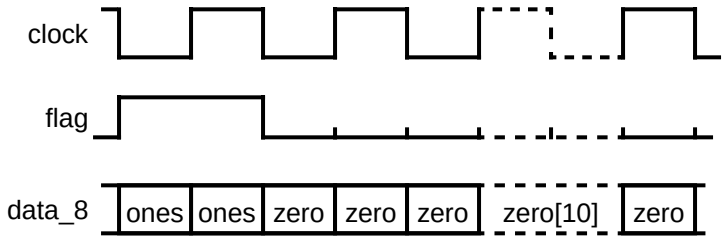

Figure B.3—Sync packet: parallel 8/10 (P10) encoding

B.2.2 Parallel 4/5 (P5) encoding

The 16-bit symbol may also be sent 4 bits at a time, using 4 subsymbols to form a logical SCI symbol. The clock and flag signals are combined, since the overhead of the extra signal is significant and 10-bit cables may then be used for two 5-bit links, one in each direction. The physical clock signal has a high-to-low transition at the start of each symbol; the flag value is derived by sampling the clock in the middle of each symbol period, as illustrated in figure B.4.

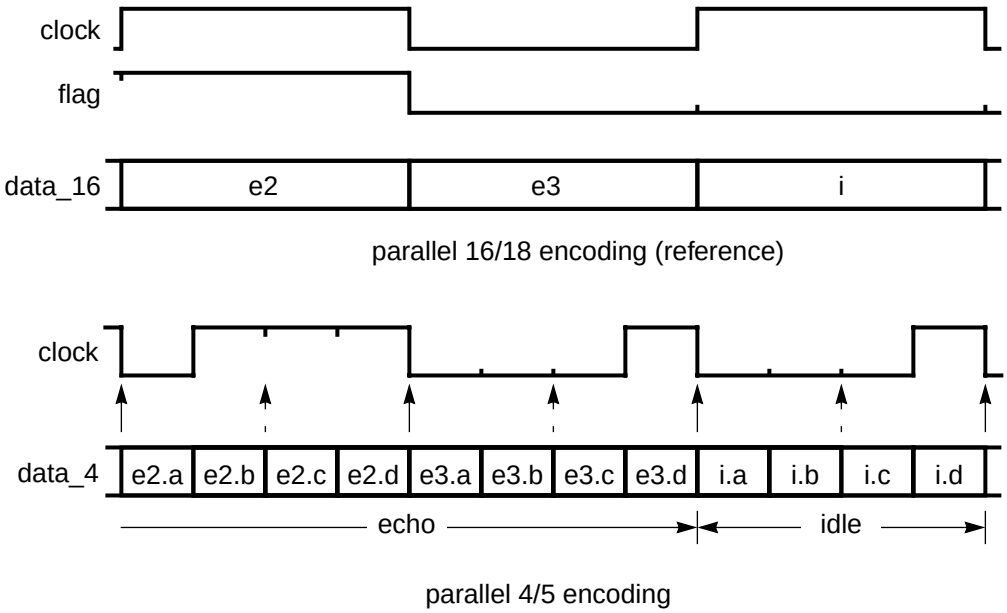

Figure B.4—Parallel 4/5 (P5) encoding

A high logical flag has a low-to-high clock-signal transition after the first subsymbol of each symbol. A low logical flag has a low-to-high clock-signal transition after the third subsymbol of each symbol.

The *e2.a*, *e2.b*, *e2.c*, and *e2.d* subsymbols are the most- through least-significant quarters of the *e2* (echo-packet) symbol; the *e3.a*, *e3.b*, *e3.c*, and *e3.d* subsymbols are the most- through least-significant quarters of the *e3* (echo-packet) symbol; the *i.a* and *i.b* subsymbols are the two most-significant quarters of the following idle symbol.

The sync packet is 16 clock-periods long (8 symbols), to simplify conversions between P18 and P5 encodings, as illustrated in figure B.5.

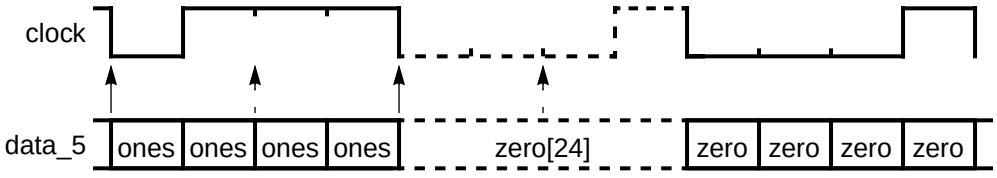

Figure B.5—Sync packet: 4/5 (P5) encoding

B.3 Wide parallel encoding

Encoding for links that are multiples of the 16-bit SCI logical symbol include the *clock* signal, provide an additional *sync* signal (to mark the beginning of a sync packet), and have one *flag* signal for each 16-bit data symbol. Transitions of the physical clock signal mark the boundaries of data-transfer intervals. The idle symbols are always the width of the link.

A sync packet consists of one data-transfer interval with ones on the *sync* signal, all *flag* signals and all data signals, followed by seven data-transfer intervals of zeroes on all these signals.

B.3.1 Parallel 32/36 (P36) encoding

On a 4-byte (32-bit) interface, two 16-bit SCI symbols are sent for each physical-clock-signal transition. There is a total of 36 signal pairs: clock, sync, and 2(flag + 16 data), as illustrated in figure B.6.

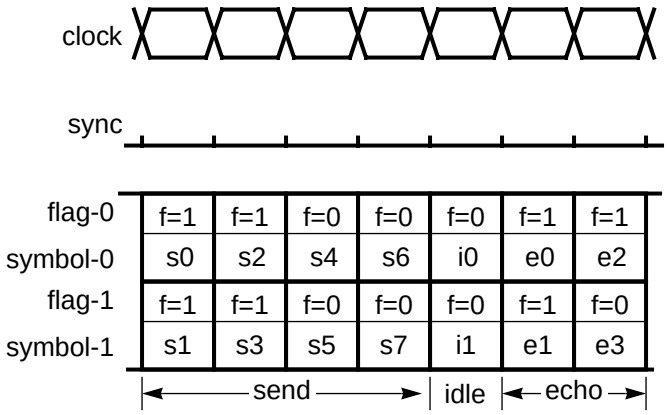

Figure B.6—Parallel 32/36 (P36) encoding

A sync packet, which remains 8 data periods in length, is illustrated in figure B.7.

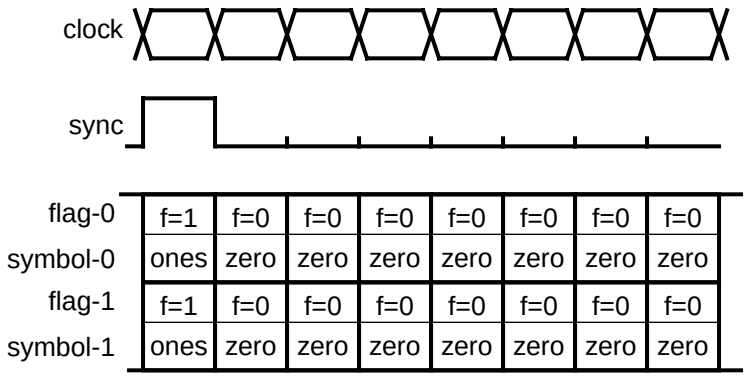

Figure B.7—Sync packet: 32/36 (P36) encoding

B.3.2 Parallel 64/70 (P70) encoding

On an 8-byte (64-bit) interface, four 16-bit SCI symbols are sent for each physical-clock-signal transition. There is a total of 70 signal pairs: clock, sync, and four sets of (flag + 16 data) signals, as illustrated in figure B.8.

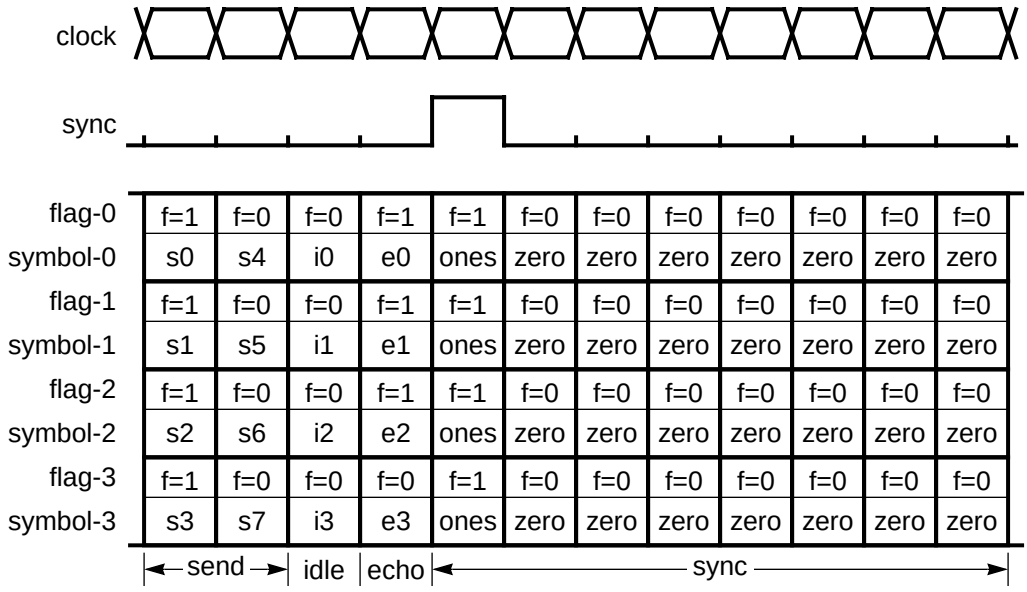

Figure B.8—Parallel 64/70 (P70) encoding

B.3.3 Parallel 128/138 (P138) encoding

On a 16-byte (128-bit) interface, eight 16-bit SCI symbols are sent for each physical-clock-signal transition. There is a total of 138 signal pairs: clock, sync, and eight sets of (flag + 16 data) signals, as illustrated in figure B.9.

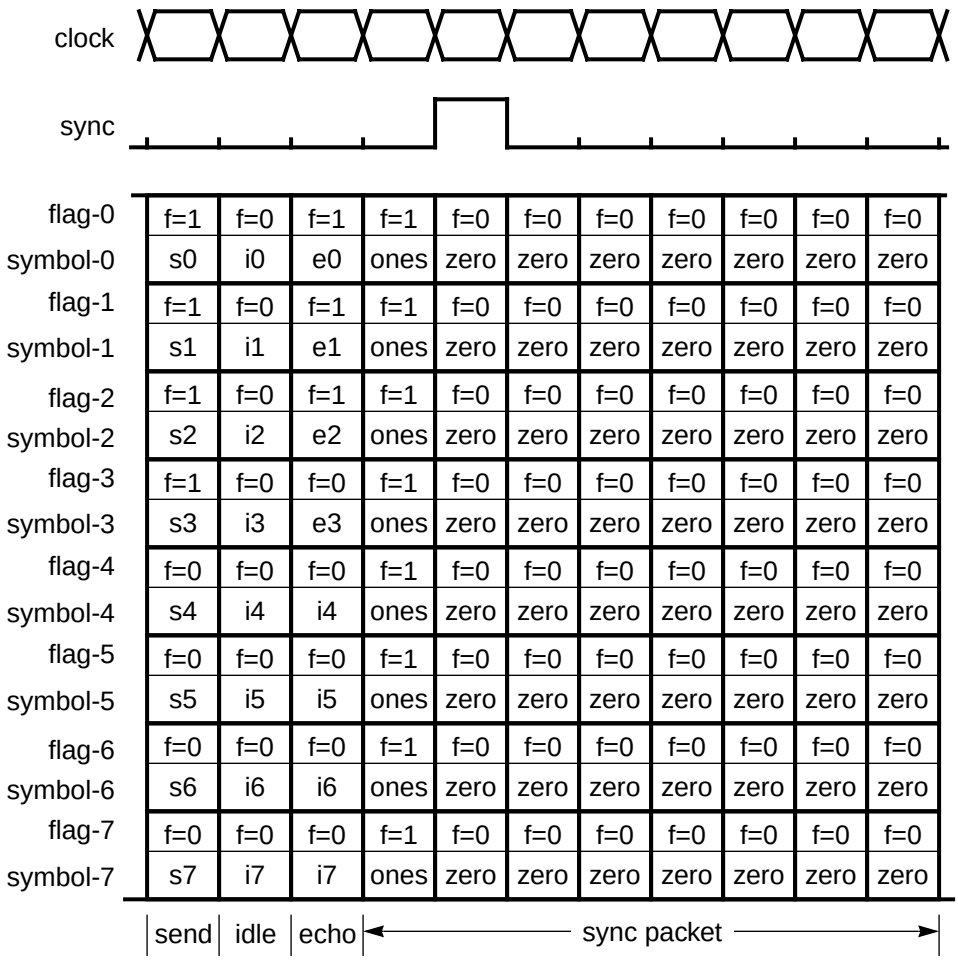

Figure B.9—Parallel 128/138 (P138) encoding

Annex C
Driver and receiver models

(informative)

C.1 Driver model

A driver design model is illustrated in figure C.1. This simplified model is only intended to give an idea of how to implement a near-constant-current differential driver. It is not necessarily a manufacturable or cost-effective design.

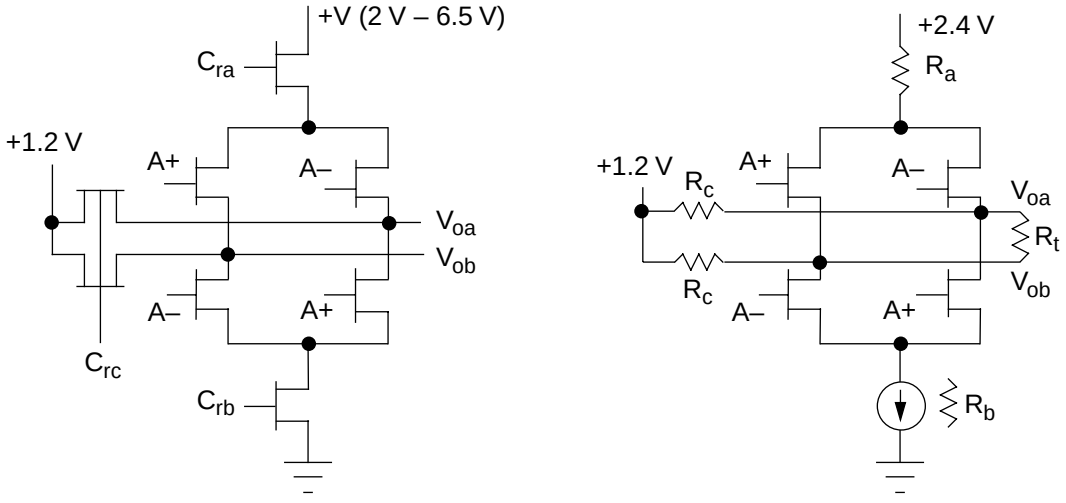

Figure C.1—Driver design model

C_{rb} is intended to operate as a current sink; however, at any particular operating condition it is equivalent to a resistance R_b. The voltage on C_{ra} is adjusted so that the attached transistor behaves like a voltage follower, providing the appearance of a resistance R_a connected to 2.4 V. The voltage on C_{rc} is adjusted so that these transistors behave like termination resistors of value R_c. R_t is the receiver termination.

Values are selected and/or adjusted so that R_a+R_{oa} and R_b+R_{ob} are closely matched, where the switching transistors have impedance of R_{oa} and R_{ob}. When so matched, the driver's output characteristics are illustrated in figure C.2.

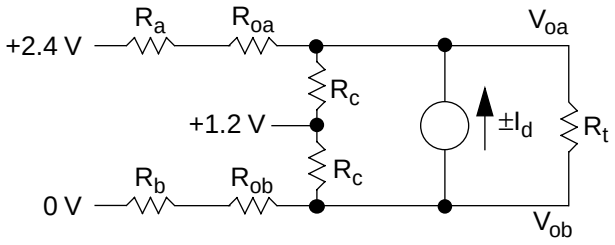

Figure C.2—Driver impedance characteristics

Although constrained by the R_o (driver's output impedance) specification, the value of R_c may be significantly greater than the transmission line impedance of 50 Ω, to minimize driver power dissipation.

C.2 Receiver model

The receiver design model includes a terminating resistance (90 Ω – 110 Ω), as illustrated in figure C.3.

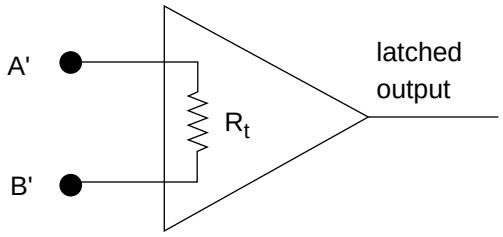

Figure C.3—Receiver design model

The design of the receiver can be simplified by recognizing that the amplifier impedance need only be large compared to the receiver's termination impedance (100 Ω). Also, it may be easier to take advantage of the fact that SCI signals are always clocked; a combined amplifier/latch (i.e. sense amp) may be easier to build than an amplifier followed by an independent latch.

C.3 Signal transmission model

The signal-transmission model, including the driver, transmission line, and termination resistance within the receiver, is illustrated in figure C.4.

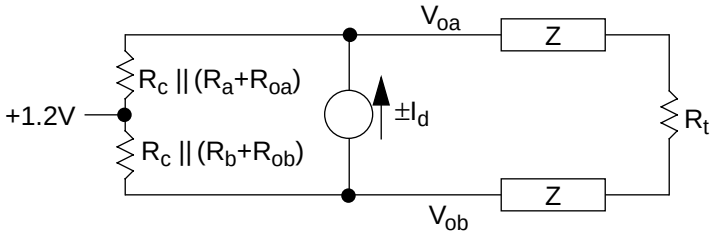

Figure C.4—Signal transmission model