Characterization of HyperTranport Interfaces on Agilent 93000 SOC Series

Todd Foulds, Karsten Hampel, Clemens Leichtle, and Andreas Olenyi
Application Consultants
Agilent Technologies

Abstract

As HyperTransport technology boosts the bandwidth for chip-to-chip interconnects to a new domain, new test methodologies are required to address the specific test challenges, and extensive characterization is a must. As Agilent's contribution to the HyperTransport consortium compatibility program, which in the first phase provides checklists and guidelines for testing HyperTransport interfaces to its members, this white paper provides the necessary background information for conducting characterization measurements on an ATE-system. In particular, it describes the main challenges associated with this task, and provides detailed information on how these test challenges are solved. Real measurement data acquired on the Agilent 93000 SOC Series complement the implementation description of the individual test methodologies. Hence, it serves as a practical guideline for test engineers involved in the characterization of HyperTransport interfaces.

1 Introduction

This paper describes the methodology for a thorough characterization of HyperTransport interfaces. It provides technical details of the implementation, together with actual measurement results for a 1.6 Gbps HyperTransport device. The main focus is on the most challenging tests such as high-speed AC tests for level and timing including setup and hold times, skew, transition times, as well as jitter measurements. The fact that HyperTransport represents a source synchronous bus, requires that measurements as well must often be done in a source synchronous fashion.

All of these measurements are parametric in nature, and are crucial for not only ensuring complete compatibility and maturity of a high-speed interface, but also for fully exploiting its design margins. This is especially valid for jitter measurements: to understand the underlying sources of jitter (e.g. information, which typically the design community is interested in), simple jitter parameters such as peak-to-peak or root-mean-square (rms) values are not sufficient. For this purpose, spectral jitter information is of tremendous help. Thus, it is very desirable to not only measure jitter histograms, but also the device' jitter spectral content in a fast and accurate manner.

For a smooth and rapid transition of a device from characterization to production and finally to high-volume manufacturing, it is desirable that all of the necessary measurements can be performed on only one ATE-system, without the need of using additional (bench-type) equipment. This clearly reduces significantly the test correlation effort. Moreover, this also allows to quickly switch back to characterization-type measurements in the case serious yield and quality problems in production occur.

This paper is organized as follows:
After giving an overview on technical aspects and the main testing challenges of HyperTransport interfaces in section 2, we outline the HyperTransport characterization methodology and main tests in section 3. Details on the implementation of individual tests complemented by exemplifying measurement results are given in the subsequent sections. Finally, we highlight the most important points for characterizing HyperTransport devices.

Note that all the presented measurement data have been obtained by solely using the Agilent 93000 SOC Series. All specifications used in the paper are based on the HyperTransport specification Rev.1.04 [1].

2 About HyperTransport™ Technology

HyperTransport technology is a high-speed, high-performance, point-to-point link designed to meet the bandwidth needs of tomorrow's computing and communications platforms. It provides an extremely fast connection that complements externally visible bus standards like the Peripheral Component Interconnect (PCI), as well as emerging technologies like InfiniBand. It is targeted at the networking, telecommunications, computer and high performance embedded applications and any application in which high-speed, low latency and scalability is necessary.
Table 1 summarizes the main features of the HyperTransport technology. A complete overview can be found in [2].

Feature/Function	HyperTransport Technology
Bus Type	Dual, unidirectional, point-to-point links
Link Width	2, 4, 8, 16, or 32 bits
Protocol	Packet-based, with all packets multiples of four bytes (32 bits). Packet types include Request, Response, and Broadcast, any of which can include commands, addresses, or data.
Bandwidth (Each Direction)	100 to 6400 Mbytes/s
Data Signaling Speeds	400 MHz to 1.6 GHz
Operating Frequencies	400, 600, 800, 1000, 1200, and 1600 Megatransfers/second
Duplex	Full
Max Packet Payload or Burst Length	64-byte packet
Power Management	ACPI-compatible
Signaling	1.2-V Low-Voltage Differential Signaling (LVDS) with a 100-ohm differential impedance
Multiprocessing Support	Yes
Environment	Inside the box
Memory model	Coherent and noncoherent

Table 1: Overview on HyperTransport technology.

2.1 HyperTransport Link

The HyperTransport link is designed to deliver a scalable and high performance interconnect between CPU, memory, and IO devices. Although it represents a *bi-directional* data path, it actually consists of two independent source synchronous clocked *uni-directional* sets of wires. Each set of wires includes CAD[n:0], CLK[m:0], and CTL, where n=1, 3, 7, 15, or 31, and m=0, 0, 0, 1, or 3, respectively. Note that commands, addresses, and data (CAD), all use the same wires for signaling, thus reducing pin requirements.

CAD signals always travel in packets (see Figure 1). These packets are multiples of four bytes (32 bits) in length.

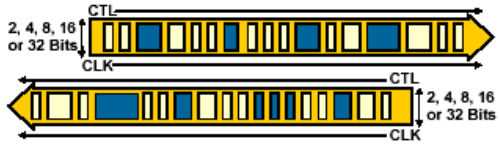

Figure 1: HyperTransport technology data paths.

HyperTransport link packets are transmitted on the high-speed CADOUT and CTLOUT wires and timed to CLKOUT, which is nominally delayed from CADOUT and CTLOUT by ½ of a bit time. In the receiver, the packets are received on high-speed CADIN and CTLIN wires and captured by sampling with CLKIN.

The high-speed data rates, which are scalable from 400 million transfers per second (Mbps), via 600 Mbps, 800 Mbps, 1.0 Gbps, 1.2 Gbps, and up to a maximum of 1.6 Gbps, are achieved by low-swing differential signaling with on-die differential termination. Commands, addresses, and data traveling on a HyperTransport link utilize a double data rate technique, i.e. transfers take place on both the rising and falling edges of the clock signal. For example, if the link clock is 800 MHz, the data rate is 1600 Mbps. An implementation of HyperTransport links with 16 CAD bits in each direction with a 1.6 Gbps data rate provides a bandwidth of 3.2 Gigabytes per second in each direction, for an aggregate peak bandwidth of 6.4 Gbytes/s, or 48 times the peak bandwidth of a 33-MHz PCI bus.

The four sideband signals RESET#, PWROK, LDTREQ#, and LDTSTOP# are single-ended LVCMOS signals supplied by the motherboard, and are used for example for link reset and power down initiation.
PWROK is a required input to each HyperTransport device to indicate that all required system power supplies are within specification and that the reference clock is within specification. RESET# is a required input to each HyperTransport device to indicate the system reset state. LDTSTOP# and LDTREQ# are used in systems requiring power management to signal requests for power related system activities. The AC and DC device output and input requirements for these signals are defined in the specification [1]. Table 2 summarizes the meaning of the different HyperTransport signals, whereas Figure 2 illustrates a 16x16 implementation of a HyperTransport link.

Signal Name	Description	Comment
CAD	**Commands, Addresses and Data:** Carries command, address, or data information.	CAD width can be different in each direction.
CTL	**Control:** Used to distinguish control packets from data packets.	
CLK	**Clock:** Forwarded clock signal.	Each byte of CAD has a separate clock signal. Data is transferred on each clock edge.
PWROK	**Power OK:** Power and clocks are stable.	Single-ended.
RESET#	**HyperTransport Technology Reset:** Resets the chain.	Single-ended.
LDTSTOP#	**HyperTransport Technology Stop:** Enables and disables links during system state transitions.	Used in systems requiring power management. Single-ended.
LDTREQ#	**HyperTransport Technology Request:** Requests re-enabling links for normal operation.	Used in systems requiring power management. Single-ended.

Table 2: Signals used in HyperTransport technology.

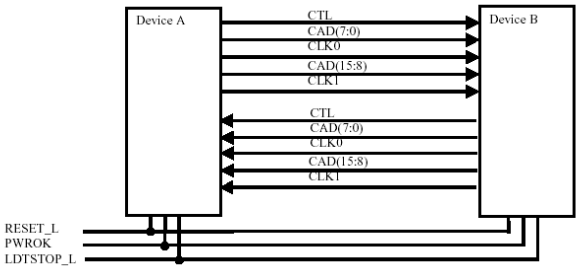

Figure 2: Example of a HyperTransport link of width 16x16 bits.

2.2 HyperTransport Physical Layer

Conceptually, the architecture of the HyperTransport interface can be mapped into five different layers, whose structure is similar to the Open System Interconnection (OSI) reference model. The physical layer defines the physical and electrical characteristics of the protocol. This layer interfaces to the physical world and includes the data, clock, and control lines. The CAD/CLK/CTL signaling used in HyperTransport technology is a type of low voltage differential signaling (LVDS). The reasons for this are the following:

1. Differential signaling provides a return current path for each signal, thus reducing the number of power and ground pins required in each package.
2. Differential signaling rejects common mode noise as generated e.g. by crosstalk and ground bounce.
3. The peak-to-peak signal voltage is actually twice the physical voltage of either trace.

In differential signaling, there is no implicit reference since the voltage potential of a differential signal is produced across two nets. Therefore, a common reference is not needed between the driver and the receiver. Instead, the true signal and its complement are subtracted, and the polarity of the result defines the logic state: a logic 1 occurs when the true signal is greater than the complement, whereas a logic 0 occurs when the true signal is less than the complement.
For HyperTransport, the signal swing ranges from 300 mV to 900 mV and swings around 600 mV (common mode voltage). The termination of each signal is done at both the source and the receiver. There are no external motherboard pull-ups or pull-downs. This is illustrated in Figure 3.

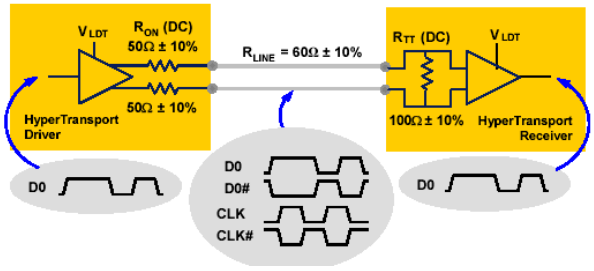

Figure 3: Enhanced low-voltage differential signaling as used by HyperTransport technology.

The on-die differential termination as well as the AC and DC device output and input voltage requirements are defined in the specification [1].

2.3 HyperTransport Test Challenges

High data rate
Since HyperTransport operates at high data rates up to 1.6 Gbps, at-speed functional tests require ultra fast ATE pin electronics, which accommodate this speed range. Moreover, especially for characterization tests, it is necessary to go even beyond the specified speed by at least 5%. The Agilent 93000 SOC Series with its NP1700-model for data rates up to 1.7 Gbps is exactly targeted at this speed challenge. For future speed requirements, the NP1700-model can be simply upgraded by software to the NP2500-model, which covers speeds up to 2.5 Gbps.

Low Voltage Differential Signaling
Since HyperTransport employs low voltage differential signaling, it is of great advantage when the ATE pin electronics supports this signaling to make sure that the test environment matches the application. For this reason, the NP-models of the Agilent 93000 SOC Series feature a native differential driver and receiver architecture, i.e. the driver is implemented as a current switch for ensuring highest symmetry of the logic signal and its complement. The receiver offers both a 100Ω cross- as well as a center-

tap termination mode, and is designed for highest bandwidth to allow highly accurate parametric measurements. Note that the differential drivers and receivers also have the full flexibility to operate as two independent single-ended pins, and that the operating mode can be easily switched from differential to single-ended and vice versa from one test to another. This is particularly desirable in cases where, for example, differential pins are also used for scan.

Source Synchronous Operation
HyperTransport utilizes a source synchronous timing scheme (i.e. clock and data origin from the same timing source) in order to maintain the proper clock-to-data timing relation for lowest clock-to-data jitter. For high data rates and devices having many switching transistors at the same time (e.g. microprocessor units), it is necessary to operate tests also in a source synchronous fashion since otherwise timing margins are significantly reduced mainly due to PLL long-term drift and jitter. In contrast to conventional ATE hardware, which only operates in a tester synchronous mode (i.e. with a fixed timing reference), the NP-models pin electronic cards also feature a built-in source synchronous operation mode to account for this HyperTransport challenge by using the device output clock as trigger for strobing the device data. The operation mode of the NP-receivers from source synchronous to tester synchronous can be easily switched from one test to the other, i.e. there are no hardware modifications (e.g. loadboard changes) necessary.

Jitter measurements
For high-speed interfaces such as HyperTransport, accurate jitter measurements are of particular importance since exceeding jitter potentially belongs to the device's dominating failing mechanisms. For accurate jitter measurements, the ATE system must not add additional jitter. The NP-models use therefore a retiming architecture for lowest jitter resulting in the best-in-class edge placement accuracy of +/- 50 psec. Together with the edge timing resolution of 1 psec and an outstanding edge placement linearity, it allows to do jitter measurements with an accuracy of typically +/- 5 psec. Since this can be done with the pin electronics itself, i.e. without using any additional instrument, jitter measurements on many pins are done in parallel in an ultra-fast fashion – e.g. 60 pins in less than 1 sec.
Additionally, in the device's characterization phase, spectral jitter information is of tremendous help in order to understand the underlying sources of jitter. For this reason, the NP-models not only allow to measure jitter histograms (which basically represents a probabilistic view), but also the device's jitter spectral content in a very fast and accurate manner.

Parametric measurements
For ensuring complete compatibility and maturity of a high-speed interface, and fully exploiting its design margins, thorough characterization of HyperTransport links require accurate parametric measurement such as AC level, rise/fall time, and data eye measurements. This poses the requirement on the ATE-system for highest bandwidth of the signal path in order not to distort the signal but examine the real device's output. Note that throughout the NP-receive path, a bandwidth of larger than 3 GHz is maintained; combined with the built-in scope–like measurement capability of the NP-receivers with a level accuracy of 10mV and a resolution of 1 mV, this allows to do parametric measurements with the pin electronics with the required accuracy.

Impedance measurements
HyperTransport uses on-die termination for minimizing signal reflections and ringing. Hence, input and output impedance tests are necessary for which the ATE-system needs full DC measurement capabilities. The NP-models therefore provide complete DC test capabilities including per-pin PMU for both HyperTransport drivers and receivers.

3 HyperTransport Characterization Methodology

Parametric tests are necessary to decide whether the device meets various rise and fall times, setup and hold times, low and high voltage thresholds, and low and high current specifications. Functional tests determine whether the internal logic of the chip behaves as intended.

The goals for characterization tests are the following:

- Ensuring complete compatibility
- Ensuring design maturity
- Determination of the operating margins of the device
- Finding worst case conditions
- Determination of testing margins
- Preparation of a smooth transition to production
- Maximizing high-volume manufacturing yield

The different tests and resulting challenges imposed on an ATE-system are summarized in

Table 3.

Characterization Test	ATE Challenges	
At-Speed Functional Test	• data rate > 1.6Gbps with at least 5% margin • easy upgrade	• future extension to higher speeds
Impedance Tests	• full DC capabilities with necessary accuracy	
Differential Swing/ Common Mode Voltage	• low swing drive/receive • level resolution • level linearity	• level accuracy • high bandwidth
Jitter Measurements	• low system jitter • high bandwidth	• timing linearity • timing resolution
Eye Mask Test	• high bandwidth • level/timing resolution	• timing/level linearity
Transition Times	• high bandwidth • timing linearity • timing resolution	• level linearity • low system jitter
Lane-to-Lane Skew Setup/Hold times	• wide timing skew range • timing linearity	• timing accuracy • timing resolution
Transition To Production/HVM	• same platform • integrated solution	• high throughput • parallel resources

Table 3: HyperTransport characterization test list and ATE challenges.

4 Impedance Characterization

The demand for higher bandwidth in today's digital systems is increasing the need of higher data transfer rates. At these high data speeds, signal integrity is critical to ensure correct data transmission. In order to address highest signal integrity, HyperTransport uses an on-chip termination on both ends with 50 Ohm line impedance on the transmitter side and a 100 Ohm cross termination on the receiver side. Table 4 shows the DC specification for the input and output impedance.

Parameter	Description	Min	Typical	Max	Units
R_{TT}	Diff. Input Termination	90	100	110	Ω
R_{ON}	Serial Output Impedance	45	50	55	Ω
ΔR_{ON}(pullup)	High Drive Magnitude Change	0		5	%
ΔR_{ON}(pulldown)	Low Drive Magnitude Change	0		5	%

Table 4: HyperTransport impedance specification parameters.

Since the driver is differential, both line impedances influence signal transitions on both legs (positive and negative); hence it is important that both terminations are symmetrical for minimum skew between both legs. ΔR_{ON}(pulldown) and ΔR_{ON}(pullup) specify the difference of both when driving a logical 0 and 1, respectively. The impedance tests are done in a steady-state: by forcing a voltage and measuring the resulting current, one can determine the impedance by making use of Ohm's law.

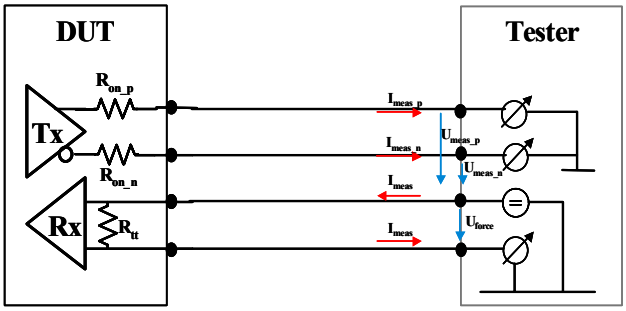

Figure 4: Tester setup for impedance tests.

Figure 4 shows a typical setup used for impedance measurements. First, the output impedance R_{on} of the HyperTransport link is measured. To do this, a V_{OL} pattern is used to drive all the outputs low for the pull-up compensation measurement. Then for the pull-down compensation measurement a V_{OH} pattern is used to drive the outputs high. For both measurements the pin PMU is then used to force 600mV on all the HyperTranport outputs and the current is measured. Figure 5 shows the results of such a measurement.

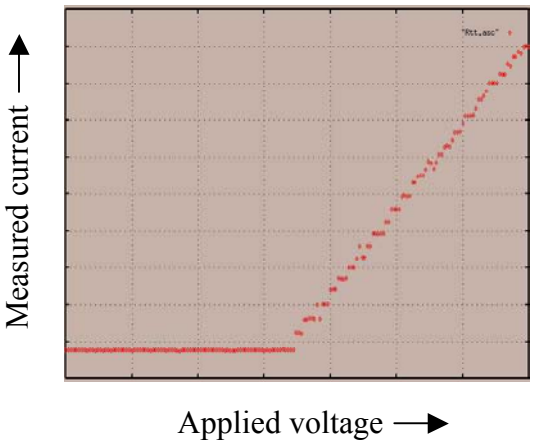

Figure 5: V-I impedance measurement results.

For the ΔR_{ON}(pullup) and ΔR_{ON}(pulldown) the measurement results for the negative leg and the positive leg are compared: both values should be the same in the range of 5%.

For the input impedance R_{tt}, the value of the impedance between true and complement signals on the HyperTransport inputs is determined by forcing a voltage and measuring the resulting current.

4.1 Conclusion

For impedance measurements, the ATE-system needs full DC measurement capabilities. If the chip uses the high-speed buffer also for DFT, such as SCAN, it is necessary that the ATE allows to switch between PMU, SCAN and high-speed path within one testflow. Figure 6 shows an example for such a flow. This allows the setup of an automated characterization testflow within one insertion.

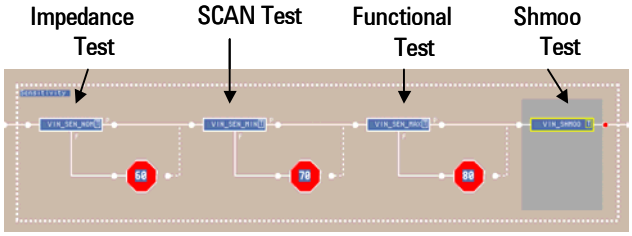

Figure 6: Testflow setup for different tests.

5 Level Characterization

For the Rx, as well as for the Tx side of the HyperTransport link, the specification lists four parameters, respectively, which needs to be analyzed for both DC, as well as for AC conditions. Whereas DC employs a static state, AC measurements are required to be performed at full speed. Since these AC measurements represent the main level test challenge, the following discussion is restricted to this case only.

5.1 Input Levels

Table 5 lists the relevant level parameters for AC input level measurements (Rx sensitivity tests).

Parameter	Description	Min	Typical	Max	Units
V_{ID}	Input differential Voltage	300	600	900	mV
ΔV_{ID}	Change in V_{ID} magnitude	-125		125	mV
V_{ICM}	Input common-mode voltage	385	600	845	mV
ΔV_{ICM}	Change in V_{ICM} magnitude	-100		100	mV

Table 5: Input AC level parameters defined by the HyperTransport specification.

Thus, the differential input voltage V_{ID} and the common-mode voltage V_{ICM} must be varied over the whole range to guarantee the device' Rx level sensitivity (see Figure 7).

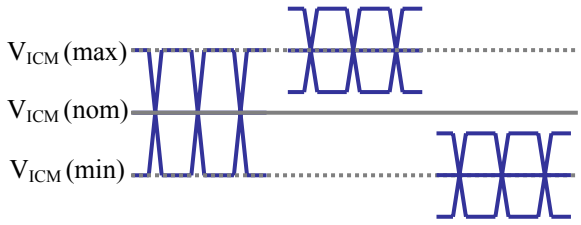

Figure 7: Input differential and common-mode voltage variation for at-speed Rx sensitivity tests.

5.2 Output Levels

Similar to input level tests, the HyperTransport specification defines also four parameters for the Tx output levels shown in Table 6.

Parameter	Description	Min	Typical	Max	Units
V_{OD}	Output differential Voltage	400	600	820	mV
ΔV_{OD}	Change in V_{OD} magnitude	-75		75	mV
V_{OCM}	Output common-mode voltage	440	600	780	mV
ΔV_{OCM}	Change in V_{OCM} magnitude	-50		50	mV

Table 6: Output AC level parameters defined by the HyperTransport specification.

For differential signaling, as used by HyperTransport, true differential comparators are required, since single-ended comparators (referenced to ground) are sensitive to common mode voltage variations. True differential comparators use the complement of the true signal as the reference; thus, the comparator result is either a logic 0 or a logic 1 without any actual voltage information. However, for determination of the differential and common-mode output voltages, it is necessary to re-extract actual voltage information from a differential signal in a scope-like manner on a per-bit basis. In this context it is important to always maintain a clean receive path with negligible signal distortion and degradation.

Figure 8 displays such at-speed scope-like measurements for 2.0 Gbps signals of ClkOut and CadOut pins, which allows fast pattern and voltage verification.

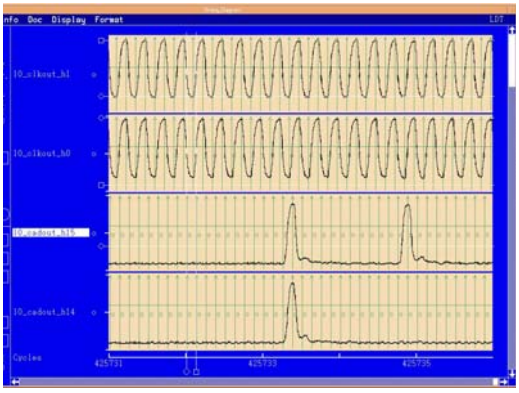

Figure 8: 2.0 Gbps ClkOut and CadOut signals measured via the 93000 scope-tool.

Figure 9 displays a small pattern portion of actual measurement data performed on a pin running at 2.0 Gbps. Note that common-mode variations (displayed in the middle) are mostly governed by signal transitions as apparent from the upper graph displaying the overlay of the positive and the negative pins of the differential pin pair.

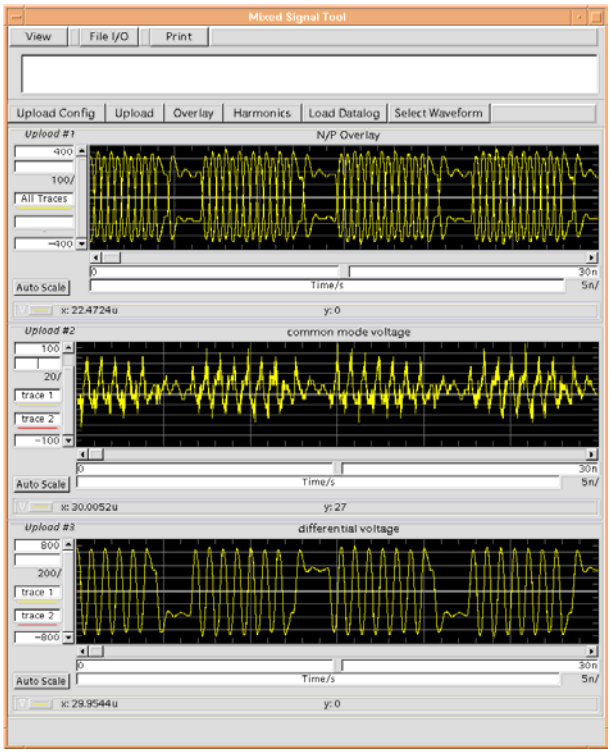

Figure 9: Overlay of the positive and negative pins of a differential pin pair, common-mode, and differential voltage, extracted from the differential signal on a per-bit basis.

5.3 Conclusion

For accurate AC level characterization measurements, the ATE-system is required to not only offer a low swing stimulus capability (while covering a wide stimulus range), but also a high voltage linearity combined with the highest accuracy and resolution. Figure 10 shows the performance of the Agilent 93000 NP-drivers measured by the NP-receivers in a direct drive-to-receive loopback configuration.

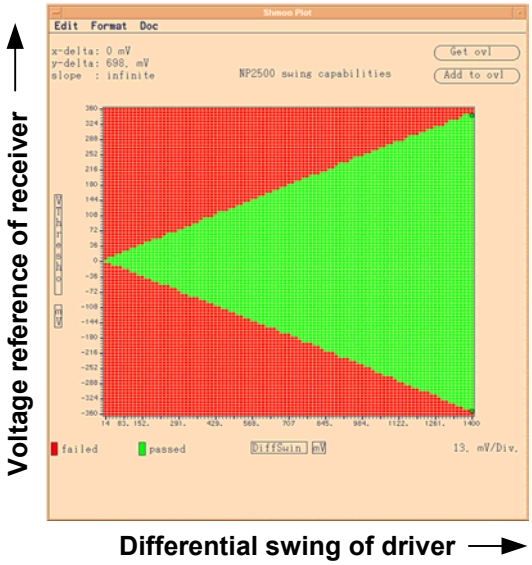

Figure 10: NP-driver swing capabilities measured by NP-receiver via direct drive-to-receive loopback.

Output voltage characterization measurements heavily rely on comparators, which have a sufficient bandwidth to accurately follow ultra-fast AC-levels. Figure 11 shows measured signals of various speeds injected by a signal generator into a 93000 NP-receiver. From this figure one can conclude that a NP-receiver is able to do parametric voltage measurements on signals even faster than 2.5 Gbps since the bandwidth clearly exceeds 3 GHz.

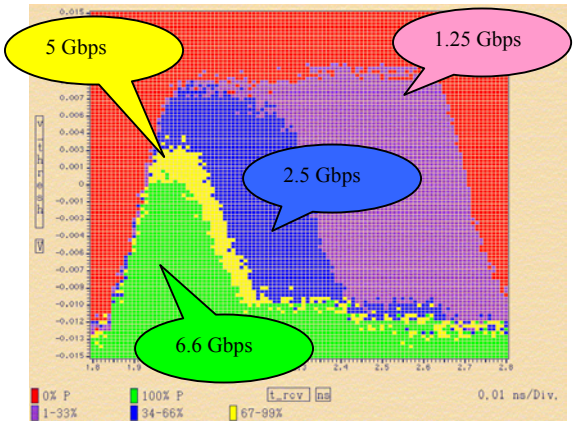

Figure 11: Ultra high-speed signals measured by NP-receiver.

6 Timing Characterization

The achievable external bandwidth of a chip is often the limiting factor of the whole digital system. To achieve highest bandwidth between chips with lowest latency, source-synchronous parallel interfaces, like HyperTransport, have been developed. Along with the data, additional timing information is sent, which is used to sample the data on the receiving chip. The source-synchronous clocking allows a more relaxed design of the clock distribution compared to centrally clocked-systems. More details will be discussed in section 7. Nevertheless, in order to guarantee the setup and hold times of the system, it is important that the skew between clock and data and between the data in one clock group is within the specification.

6.1 Receiver Input Setup & Hold Time

On the receiver side mainly two parameters are specified:

- T_{su} : receiver input setup time
- T_{hd} : receiver input hold time

The setup time T_{su} is measured from the crossing point of the data transition to the clock transition. The hold time is measured from the crossing point of the clock to the data transition. The specific values of these parameters depend on the link speed, and can be found in the HyperTransport specification.

For measuring the setup time for the device, the timing of each single data line is skewed from a relaxed value towards the clock transition while monitoring whether the device's logic is passing. For hold time measurements the data again is skewed from a loose value towards the clock transition until the device is failing. This is done over the whole pattern to guarantee the worst case, see Figure 12. Since the input voltage influences the setup and hold times, it is recommended to repeat those measurements while sweeping V_{ID} and V_{ICM} from the minimum to the maximum values.

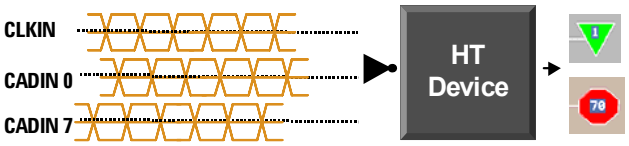

Figure 12: Setup and hold time measurements.

In addition to the setup and hold time, the specification discusses two additional receiver timing parameters, namely T_{cadvrh} and T_{cadvrs}, but since they are system oriented, it is not necessary to measure them on an ATE.

6.2 Output Skew

For HyperTransport, the transmitter clock has a phase shift of 90˚ to the data; this gives the best timing margins for the receiver in order to guarantee that the correct data is sampled.

The specification defines a parameter T_{CADV}, which at the same time also gives a measure for the group skew between the data/control lines in one clock group as illustrated in Figure 13. The values for T_{CADV} for the different link speeds can be found in the specification.

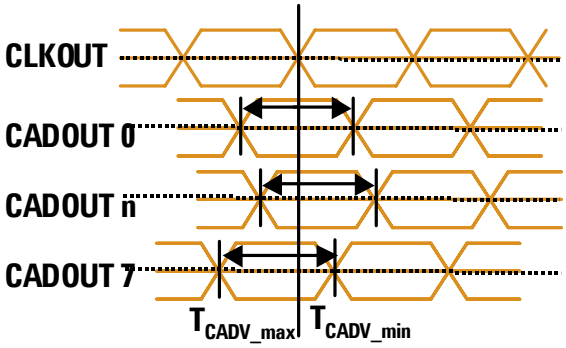

Figure 13: T_{CADV} measurements.

In order to measure T_{CADV} , the timing of the compare pattern is moved backwards and forwards with respect to the clock signal. The starting point of the search is the nominal value where the clock signal is delayed by one-

half of a bit time. The values for $T_{CADVmin}$ and $T_{CADVmax}$ are defined relatively to the clock transition.

To insure that the worst case conditions are covered, it is recommended to do the measurements over the whole pattern and under conditions where the skew is expected to be maximized.

For a fast visual feedback of the transmitter's signal performance over the whole pattern range, it is helpful to perform eye-diagram measurements for each single data line and overlay them to get the resulting minimum eye width of the transmitter. The shmoo moves the compare pattern over the outcoming data stream and at the same time the compare voltage is varied. Figure 14 shows the results of such a shmoo where the link is operated at 2 Gbps. The green area is the common eye with the different colors showing the overlap of all output pins in one clock group.

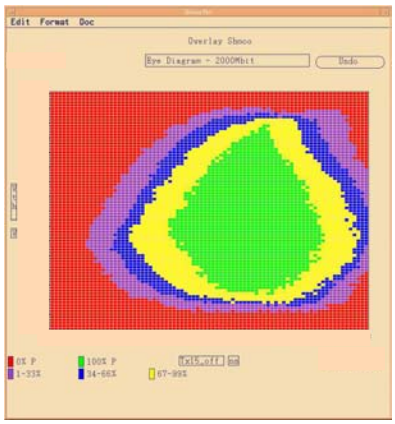

Figure 14: Overlaid eye diagram for CadOut pins running at 2 Gbps.

6.3 Single-ended Output/Input skew

The intra-skew on the transmitter and receiver, T_{ODIFF} and T_{IDIFF}, is defined by the timing difference between the true and complement signal of one differential pair.

In order to characterize these parameters, measurements need to be done in a single-ended fashion, which means that drive and compare edges can be varied individually for the true and complement signal.

6.4 Rise/Fall Time

Proper transition times are important for minimizing output jitter: when the bit-time is getting shorter than the transmitter settling time, the output no longer swings fully and the values of the previously transmitted bits affect the waveform of the current bit - a phenomenon usually referred to as data-dependent jitter or inter-symbol interference. The characterization of device jitter will be discussed in section 8 in more detail.

Rise and fall time tests measure the time interval between two voltage levels of a rising or falling edge, respectively. Thus, these measurements are done in a single-ended fashion. To do so, a compare edge is moved over a rising or falling transition while the compare voltage is swept. Figure 15 shows the result of a shmoo from which one can easily calculate the device's rise/fall time for different threshold values. Note that such measurements don't need to be done in such an interactive manner, but there are more efficient ways based on intelligent searches, which can be performed in parallel on many pins.

Figure 15: Rise time of a single CADOUT pin running at 1.6 Gbps.

6.5 Link Frequency

For characterizing the speed margins of the link, the reference clock of the device is reprogrammed to a faster frequency, and all measurements described above should be repeated. It is recommended that the link frequency is tested at least 5% above the maximum speed in order to guarantee enough process margins. However, for a thorough characterization it is of course desired to find the maximum speed of the HyperTransport link; in this case the reference clock of the device is increased until eventually setup and hold time violations occur and the

device fails. Figure 16 schematically shows the setup to test the frequency margins of the device.

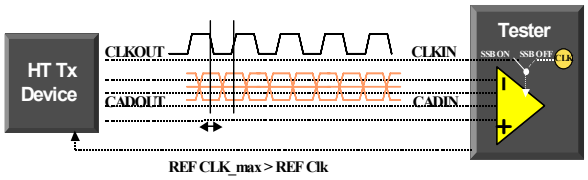

Figure 16: Characterization of the maximum frequency of a HyperTransport link.

6.6 Conclusion

In this section we have discussed various timing parameters for a HyperTransport link and how they are measured on an ATE-system. To make accurate timing measurements, the timing accuracy, resolution, and edge placement linearity of the measurement instrument is crucial: for example for being able, to measure timing variations in the pico-second range, the tester needs to have an edge resolution as high as one pico-second.

Besides the timing characteristics of the tester, also the receiver bandwidth is very critical for rise and fall time measurements. For example, if the device has a maximum edge rate of 150 ps, the following minimum tester's receiver bandwidth is required:

BW = 0.35 / Tr = 0.35 / 0.150e-9s = 2.33 GHz

A detailed explanation for the factor 0.35 can be found in [3].

7 Source-Synchronous Characterization

As mentioned in the introduction, source-synchronous timing simplifies considerably the clock distribution thus increasing the robustness of the digital system. Figure 17 shows a simplified block diagram of a source-synchronous system.

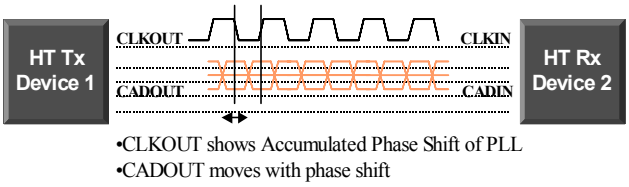

Figure 17: Source synchronous Tx and Rx setup.

On the left a HyperTransport transmitter is shown which sends a clock signal along with the data. At the receiver, the clock signal is used to latch the data for further processing. Currently, the maximum clock speed in the HyperTransport specification is 800 MHz.

For testing a source-synchronous interface, it is important that the ATE-system is able to synchronize on the transmitted clock in order to unveil the complete valid data eye (as sketched in Figure 18).

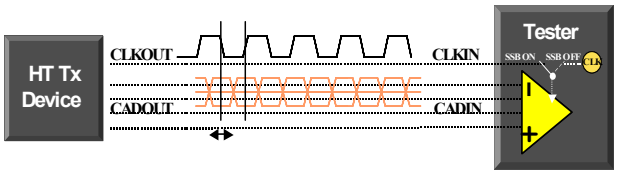

Figure 18: Source synchronous tester setup.

Simple Shmoo measurements allow for a fast visualization of the performance of the link transmitter: the left-hand side of Figure 19 shows a measurement result obtained by using the tester's master clock ("SSB OFF") for sampling the transmitter data. Here, the data-eye is clearly closed by the presence of jitter. This is in contrast to the result shown on the right-hand side, where here the transmitter clock was used to trigger the tester's compare edges ("SSB ON").

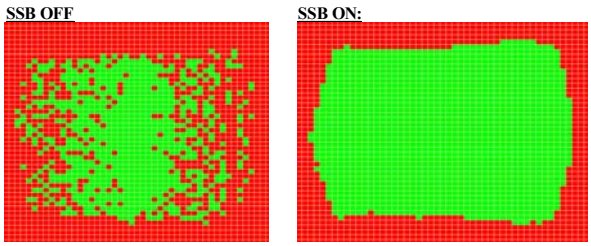

Figure 19: Difference between tester synchronous ("SSB OFF") vs. source synchronous ("SSB ON") measurement results.

However, for a thorough characterization of the timing uncertainties, it is necessary to do a jitter frequency separation on the clock and data pins as described in section 8.

7.1 Conclusion

For ATE-systems, the source-synchronous behavior of the HyperTransport link is one of the biggest challenges, since the tester needs to synchronize its strobes on the transmitter clock. Moreover, for being able to quickly check the actual closing of the data eye due to its source-

synchronous nature, it is really desirable that the source-synchronous operation mode can be switched on and off within one testflow as indicated in Figure 20.

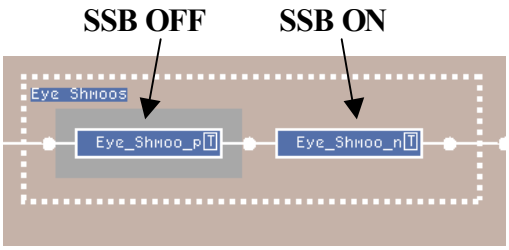

Figure 20: Eye diagram setup with source-synchronous "on" and "off".

8 Jitter Characterization

As mentioned previously, a very important parameter during characterization (and sometimes even in production) is jitter. By definition, jitter represents variations of timing events resulting from threshold crossings of a signal relative to a reference event. Jitter is mostly caused by the superimposition of additive noise onto a signal. In digital circuits, ground bounce resulting from a heavy switching activity on the circuit, as well as crosstalk (i.e coupling from other on- and off-chip channels or nearby traces) are usually the dominant mechanisms for jitter. Other sources include random fluctuations due to the inherent thermal and shot noise of active and passive system components, as well as signal reflections and ringing caused by improper termination.
In addition to these (typically) short-term uncertainties, the HyperTransport specification also discusses long-term transmitter PLL variations resulting from temperature and power supply drifts, as well as accumulated PLL phase errors.

8.1 Jitter Histogram

The most straightforward characterization of jitter is the recording of a jitter histogram. Such a *probabilistic* view on jitter gives a statistical distribution of the actual occurrence of the jittering event. From such a histogram, different measures of the jitter can be derived; these measures are not restricted to values such as the total jitter (i.e. the width of the distribution), or the root-mean-square (rms), which are the most commonly used ones, but can be extended to any other meaningful distribution parameter, yielding an even more representative jitter parameter for an actual signal.

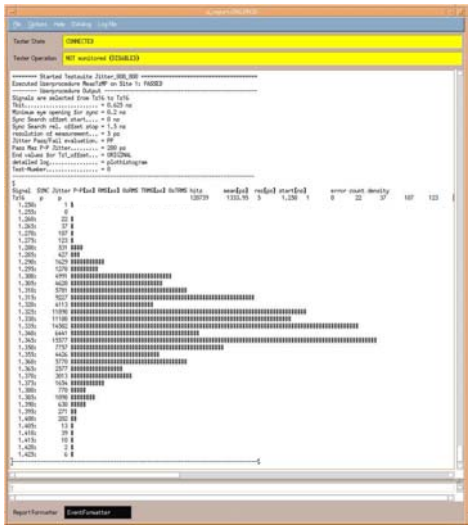

Figure 21: Jitter histogram of a ClkOut pin running at 1.6 Gbps.

Figure 21 shows actual jitter histogram data measured on a ClkOut pin running at 1.6 Gbps. Whereas on the vertical axis, the temporal occurrence of the event is shown, the horizontal axis gives a measure for the probability of the event. From this histogram we immediately can recognize that the underlying jitter sources in this case are mostly random in nature, since this curve resembles very much a Gaussian (or "normal") distribution. This measurement, which has been done on the 93000 NP-model, is based on functional tests using the following methodology visualized in Figure 22:

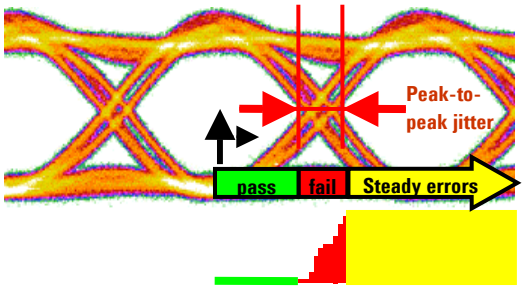

Figure 22: 93000 jitter measurement methodology.

the compare edge is moved from the center of the data eye towards the transition region until a functional test starts to fail and leads to errors. The numbers of errors (i.e. the number of failing transitions) are counted, and the edge is further moved into the transition region. This is done until one finally reaches a steady error count which means, that all transitions in the used pattern fail. Hence, the number of failing bits, when running a functional test, is used as a

measure for the probability of the occurrence of the threshold crossings. Moreover, it is quite obvious, that due to the statistical nature of jitter, the more transitions the pattern contains, the more accurate the measurement actually is. Especially in production testing, one therefore needs to find a reasonable tradeoff between achieved accuracy and affordable test time.

8.2 Jitter Spectrum

For a thorough understanding of the underlying sources of jitter, e.g. information needed by the design community (which might be also required for process engineers), in addition to jitter histograms, spectral jitter information is of tremendous help.
Figure 24 show a ClkOut pin jitter power density spectrum for two different devices: for the device in , the pin exhibits a relative small amount of random jitter only (recognized as a typical band-limited spectrum).

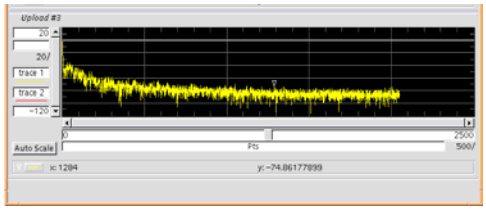

Figure 23: Power density jitter spectrum of a ClkOut pin of "good" device.

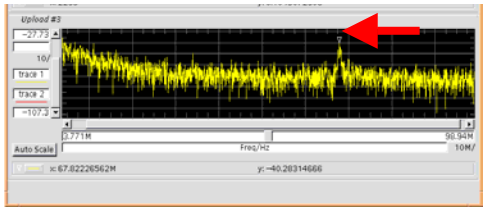

Figure 24: Power density jitter spectrum of a ClkOut pin of a "bad" device.

In contrast, the device in
Figure 24 shows, in addition to a much more pronounced amount of (broad-band) random jitter, a clearly visible frequency spike (indicated by the red arrow) on top of it. This deterministic jitter contribution gives a valuable hint on a possible design weakness or an actual manufacturing process problem. Note that these measurements help tremendously to understand why device B is actually failing, whereas device A is fully functioning. Again, the measurements have been done on the Agilent 93000 NP-models without any external instrument.

8.3 Conclusion

The measurement instrument largely determines the achievable accuracy of jitter measurements, which imposes the following requirements on the measurement system:

1. The timing resolution, i.e. the resolution with which the compare edges can be positioned, must be fine enough (i.e. must have the same or better resolution than jitter histogram variations) since this directly translates into the resolution of the jitter measurement.
2. Since any non-linearity of the compare edge placement directly translates into measurement errors, the measurement system must provide the highest possible edge placement linearity.
3. The bandwidth of the measurement instrument must be large enough to accurately follow threshold crossings.
4. The measurement system must not add significant system jitter inherent in the instrument to the actual signal, since otherwise the measurement would be significantly inaccurate.

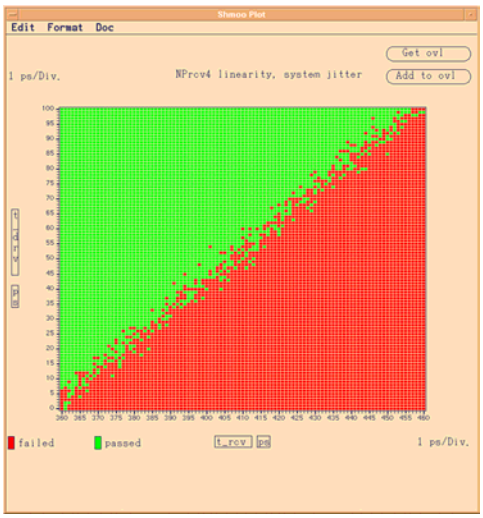

Figure 25: 93000 NP-model timing linearity, resolution, and system jitter from a direct drive-to-receive loopback measurement.

The excellent timing linearity of the 93000 NP-models together with its outstanding 1 ps timing resolution and lowest system jitter as apparent from Figure 25, demonstrates that the NP-models indeed fulfill all these requirements. This measurement was done in a direct NP-driver to NP-receiver loopback configuration via a shorted loadboard: whereas on the horizontal axis, the receiver edge is moved in 1 ps steps, on the vertical axis, the driver edge is moved with the same resolution. We can also

recognize the small system jitter of the complete drive-to-receive path (including loadboard contributions) of the NP-models as "sprinkles" around the straight line of 45 degrees.

Note that with this unique jitter measurement capability of the NP-models it is no longer necessary to use special scope-type lab instrumentation for accurate jitter measurements.

When it comes to production, an important aspect of jitter measurement is the necessary execution time. For achieving the required throughput goal, the measurement system must be able to do jitter measurements on several (or many) pins in parallel, i.e. the measurement system needs parallel resources for these measurements. With the pin electronics of the NP-models a throughput of 60 pins in less than 1 s can be achieved. This outstanding speed can be realized, since the error count is directly available in the hardware of the measurement system; otherwise data needs to be captured and subsequently uploaded to the system controller for calculating the error count by software analysis of the captured data – a task which is very time-consuming.

9 Summary

For characterization of HyperTransport links, parametric measurements such as impedance, AC level, timing, and jitter measurements play the dominant role. All of these measurements are very challenging for the measurement system, since it requires a timing architecture which is able to provide the highest accuracy and resolution, combined with an excellent edge placement linearity over a wide range while minimizing its own system jitter.
Moreover, since HyperTransport uses a source-synchronous timing scheme, the measurement system must also support such an operation mode in order to match the final application.
In addition, the determination of the link's voltage characteristics require a small-swing stimulus capability and a high receive sensitivity of the measurement instrument with an excellent voltage linearity, as well as an overall receive bandwidth larger than 2.3 GHz as discussed in section 6.4.
Finally, for a smooth transition from the characterization to the production phase, it is very desirable that all of the measurements can be done on one and the same ATE-platform without using any additional external instruments; this indeed minimizes test correlation efforts, since many parametric measurements developed during characterization might be still required during early production to ensure the product quality while developing confidence in the design and the maturity of the process.

References

[1] HyperTransport Specification version 1.04.
[2] White paper "HyperTransport Technology I/O Link – A High-Bandwidth I/O Architecture"
[3] H. Johnson and M. Graham: High-speed digital design, Prentice Hall, 1993