Loopback Architecture for Wafer-Level At-Speed Testing of Embedded HyperTransport™ Processor Links

Alvin L. S. Loke, Bruce A. Doyle, Michael M. Oshima*, Wade L. Williams**, Robert G. Lewis**, Charles L. Wang*, Audie Hanpachern*, Karen M. Tucker, Prashanth Gurunath*, Gladney C. Asada*, Chad O. Lackey, Tin Tin Wee, and Emerson S. Fang*

Advanced Micro Devices, Inc., 2950 E Harmony Rd, Fort Collins, CO 80528-3419 USA
*Advanced Micro Devices, Inc., 1 AMD Place, Sunnyvale, CA 94085-3905 USA
**Advanced Micro Devices, Inc., 7171 Southwest Pkwy, Austin, TX 78735-8953 USA
alvin.loke@ieee.org, bruce.doyle@ieee.org

***Abstract*—We present transceiver serial loopback that enables cost-effective wafer-level at-speed testing of HyperTransport™ (HT) I/O for processor die-to-die communication. Besides facilitating known-good-die testing, this feature provides observability of multi-chip module (MCM) die-to-die links that are completely embedded without external pin visibility. We demonstrate production screening of 45-nm SOI-CMOS wafers at 6.4 Gb/s.**

I. Introduction

Point-to-point connectivity protocols, such as HyperTransport™ technology [1], have become the preferred solution for delivering high-bandwidth, low-latency communication among processor dies in multi-socket systems. This approach overcomes performance limitation in legacy front-side bus architectures in which faster data transfer is impeded by the large capacitive load and half-duplex nature of long shared buses. With increasing socket counts in AMD systems, the commensurate increase in HT ports per die heightens the need for exhaustive I/O test coverage.

This paper presents the design-for-test (DFT) transceiver loopback architecture that enables post-bump wafer-level at-speed testing of HT I/O ports and their interface with embedded NorthBridge (NB) I/O controllers. By sorting for port functionality at the wafer level, we package only known good dies in single-die substrates and, more recently and significantly, in costlier MCM substrates. Furthermore, since die-to-die links in MCM packages are embedded and do not have externally visible package pins, this DFT feature provides observability for package-level characterization and debug. We will cover system- as well as circuit-level aspects of this DFT implementation. Since the loopback link is implemented on-chip, the test uses standard probe card technology and provides repeatable and reproducible coverage without additional test hardware requirements. The loopback test is now fully integrated into AMD's production test flow. Primary use is for per-lane screening at 6.4 Gb/s on Magny-Cours—AMD's first MCM product built in 45-nm SOI-CMOS [2], featuring a total of 12 cores.

II. HyperTransport™ Links in AMD Processors

HyperTransport™ is the full-duplex point-to-point parallel link protocol used by AMD processors for die-to-die communication [1]. As seen in Fig. 1, the data transfer is source-synchronous, retiming the received NRZ data stream by a forwarded clock for good common-mode jitter rejection and low link latency. Each link direction consists of two sublinks (Sublink 0 and 1) of 10 differential lanes: 1 half-rate forwarded-clock (CLK) lane, 1 control (CTL) lane for error/interrupt handling, and 8 command/address/data (CAD) lanes. Since each sublink has its own CLK and CTL lanes, the sublinks can operate either independently or as a pair to offer a width of 16 CAD lanes. Moreover, each sublink width can be reduced to 4, 2, or 1 active CAD lanes during periods of lighter link traffic. In each HT port, an on-chip PLL derives the transmitter clock for both sublinks from a 200-MHz spread-spectrum clock (SSC) source. Worst-case channels can be as long as 30 inches of 100-Ω differential FR-4 traces with possibly two interposing connectors if the sockets are mounted on different daughter boards.

An auto-negotiated transfer rate of 0.4 to 6.4 Gb/s per lane (in 0.4-Gb/s increments) can be selected depending on real-time link bandwidth demands and power constraints. At 2.0 Gb/s and below (HT1 mode), the receiver simply retimes received data using the received forwarded clock without clock-to-data deskewing. Above 2.0 Gb/s (HT3 mode) when unit intervals are shorter, DLL-based clock-and-data recovery (CDR) is turned on to align the received clock phase to received data transitions in each data lane for better retiming margin.

On-demand operation of all HT ports in a die is directed by the common embedded NB. When data must be transmitted to another die, NB activates the appropriate HT port (into its L0 active state) to connect to the intended destination die and initiates bidirectional link training with the destination HT port. Once handshaking is established, the NB controller at each end of the link commences data transfer to the other side. When the link is no longer needed, each NB instructs its respective HT port to enter light sleep (LS1), deep sleep (LS2) or power-down (PHY-OFF) state for power saving.

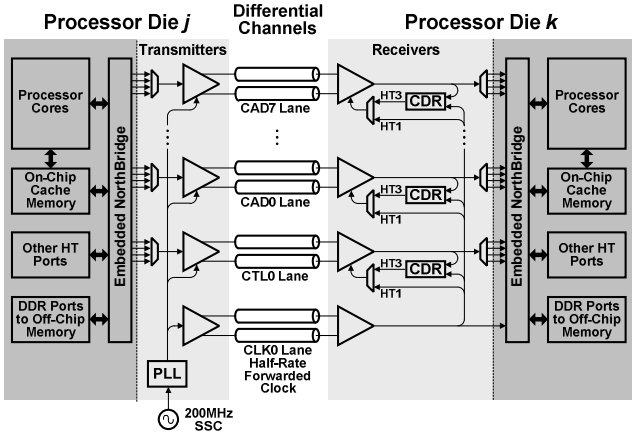

Fig. 1. HT link in AMD processors. Only one direction (from Die *j* to Die *k*) of one sublink (Sublink 0) is shown.

For multi-die processor sockets, same-package dies communicate with each other through MCM substrate traces. These HT links are embedded because they are not electrically visible outside the package in order to reduce pin count. In the Magny-Cours example, the six cores in each of two identical dies exchange data with cores in the other die through an embedded full link and embedded half link as depicted in Fig. 2. These data transfers are additionally managed by the die-to-die communication link (DDCL). Both dies are packaged in a single Socket G34 package in which the embedded full link traces span 3–4 inches on the ceramic MCM substrate (see Fig. 3).

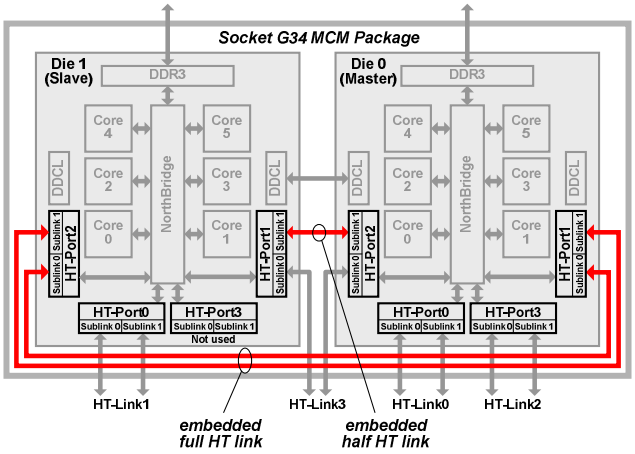

Fig. 2. Block diagram floorplan of 12-core Magny-Cours MCM processor with embedded full and embedded half HT links.

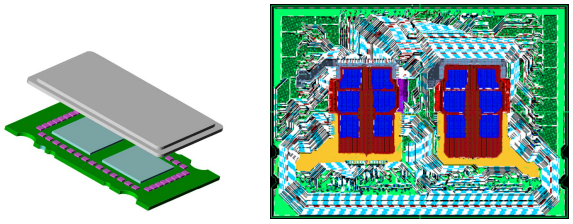

(a) (b)

Fig. 3. (a) Socket G34 package and (b) MCM substrate connectivity for twin-die Magny-Cours processor.

III. Loopback Architecture

A. System-Level Implementation

The NB I/O controller (NBICT) in Fig. 4 is the NB subunit that manages the training and normal operation of each HT port as well as data flow between HT transceivers and NB Core. NB Core provides data to HT transmitters (TXs) and distributes data from HT receivers (RXs). Prior to transmission, data can be optionally encoded through either XOR-scrambling with a lane-specific 23-bit LFSR pattern [1] or 8b/10b conversion. Encoding reduces channel intersymbol interference (ISI) and lane-to-lane crosstalk. If transmitted data are encoded, received data are correspondingly decoded. The NBICT→TX and RX→NBICT interfaces are asynchronous, so FIFOs are inserted to absorb timing uncertainties [3]. Also, each HT TX performs 4-to-1 serialization of outgoing data for bump count reduction. Received data are conversely deserialized in the HT RX.

Three on-chip loopback modes are available for testing:

1. HT Serial Loopback — HTSL (for wafer-level test),
2. HT Parallel Loopback — HTPL (for package-level test), and
3. NBICT Parallel Loopback — NBPL (for package-level test).

We focus on HTSL for wafer-level testing. Here, the TX serial outputs stream back to the corresponding RX through buffered on-chip channels. With JTAG request to initiate the test, NBICT sends common unencoded training patterns [1] to all enabled TX lanes to handshake with corresponding RXs. In HT3 mode, data transitions in the training pattern enable each CDR to align the forwarded clock phase to its corresponding data phase. Upon sending the final bit of training, NBICT sends out a test pattern provided by the tester. When post-training bits eventually arrive at NBICT, a bit-by-bit error check is performed on the received vs. transmitted bits to validate the loopback link integrity. Since the transfer is source-synchronous, CLK lane loopback must always be active to provide the forwarded clock for RX retiming. The TX PLL must also be functioning to supply the TX clock.

The two parallel loopback modes, HTPL and NBPL, provide test coverage that complements serial loopback coverage. Specifically, they test the TX output driver and RX analog front-end which are not exercised in HTSL testing. These modes can run only at the package level since they require external stimuli from another HT port or a bit error rate (BER) tester; the contact quality of wafer-level probing to wafer bumps is too unreliable for Gb/s transfers. In HTPL, data come from the RX bumps, loop to the TX FIFO, and leave the TX bumps. The FIFO is needed to synchronize the unknown phase relationship between TX PLL and RX CDR clocks. In NBPL, RX data returns to TX in NBICT in order to exercise the HT–NBICT interfaces as well as NBICT.

The general loopback configuration for all 20 lanes in a single HT port is depicted in Fig. 5. Since each sublink operates independently, we chose to perform loopback on one sublink at a time. This architecture allows pairing of transceiver lanes (CLK0/CLK1, CTL0/CTL1, CAD0/CAD8, …, CAD7/CAD15) and sharing of loopback channels to minimize wiring.

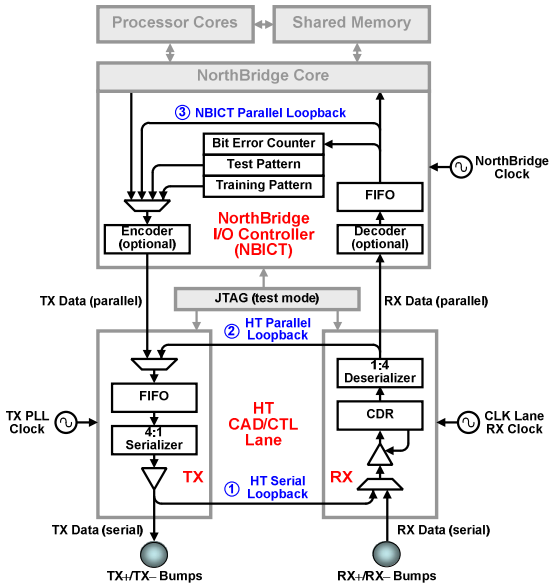

Fig. 4. HT CAD/CTL lane TX and RX interface with NBICT.

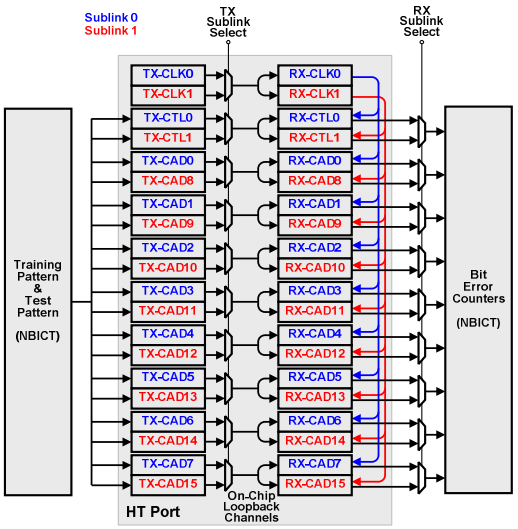

Fig. 5. General loopback configuration of entire HT port.

Furthermore, this scheme offers debug capability not possible with unshared loopback channels. By inserting the necessary multiplexers controlled by separate TX/RX sublink select signals, we can independently choose which sublink transmits and receives. So, for example, if Sublink 0 fails loopback, we can configure Sublink 0 TXs to loop back to Sublink 1 RXs or Sublink 1 TXs to loop back to Sublink 0 RXs. These additional tests can isolate identified lane failures to be TX- or RX-related.

Since each TX and RX is independently activated by NBICT, we can exercise loopback on as few as 1 CAD/CTL lane provided the CLK lane is enabled. By turning on fewer lanes, we can reduce supply bounce noise that comes primarily from TX drivers switching large currents through parasitic inductances in the probe card stackup. Fig. 6 illustrates noise profiles at the supply bumps during loopback testing with TX drivers transmitting PRBS-7 at 6.4 Gb/s. For test validation, we need to ensure that this test artifact does not introduce pessimistic test failures.

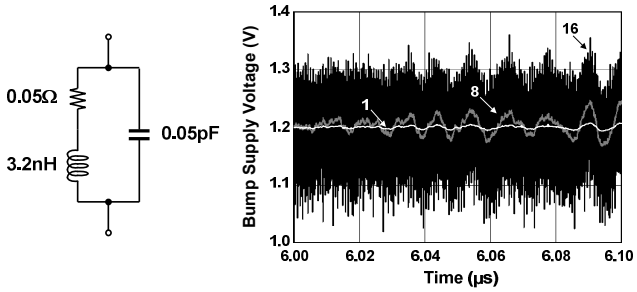

(a) (b)

Fig. 6. (a) Approximate model of probe card pin, and (b) simulated supply bump noise profile with 1, 8, or 16 TX drivers enabled.

B. Loopback Channel Implementation

The transceiver on-chip loopback channels are configured as shown in Fig. 7 for the horizontal HT port floorplan. A similar configuration exists for the vertical HT port. Unique horizontal and vertical ports were built since polysilicon gates could not be rotated by 90° in our technology for better lithography control.

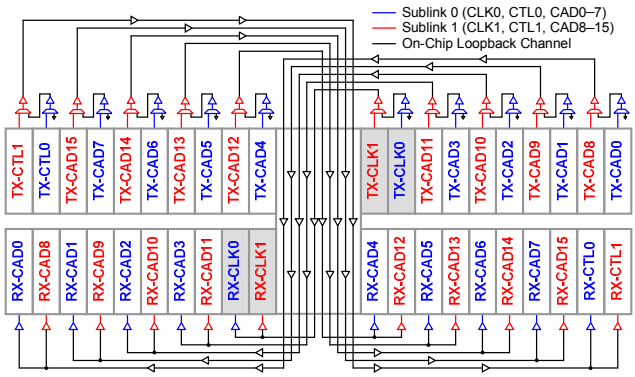

Fig. 7. Floorplan of transceiver loopback in horizontal HT port.

Loopback signaling along the shared channels is single-ended full-rail CMOS despite differential/complementary signaling throughout the TX and RX data path. This was chosen for simplicity and to halve the number of loopback routes. CMOS buffers are periodically inserted to ensure sharp data transitions and reduce ISI. These inverters were designed using Monte Carlo feedback to ensure that duty cycle distortion (DCD) from differences between rise and fall delays is acceptable. During normal HT operation, the forwarded clock is AC-coupled for DCD reduction; we need to minimize DCD in CLK loopback where AC coupling is bypassed. We also ran simulations to verify no accidental inversions in loopback buffering. Despite the unmatched wire lengths and propagation delays across the 10 loopback channels, total received skew is comfortably corrected by HT3 CDR and fall within HT1 non-CDR retiming requirements. Also, notice the daisy-chaining of multiplexers for each TX pair. The multiplexer at each TX actually belongs to the TX circuit block that is replicated across lanes and circumvents an AMD design flow restriction that allows only wires and inverters to be added outside of major circuit blocks.

C. Circuit-Level Transmitter Loopback Implementation

The TX architecture employs a hybrid voltage/current-mode 4-tap FFE, shown in Fig. 8. The loopback signal is tapped off the pre-driver cursor that switches the main output driver. Since the loopback circuitry does not capacitively load the driver output, there is no TX return loss penalty. Furthermore, there is no need for additional ESD protection (which loads the TX output even more) to safeguard the MOS gates of the loopback circuitry. This approach also offers the flexibility of operating loopback with the driver disabled to reduce the magnitude of supply bounce.

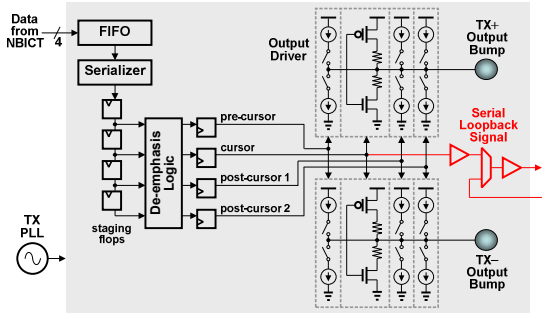

Fig. 8. TX architecture with loopback implementation.

D. Circuit-Level Receiver Loopback Implementation

The full-rate RX architecture shown in Fig. 9 consists of independent paths for HT1- and HT3-mode data recovery. In the HT1 path, the received data stream is simply amplified and delayed by the forwarded clock distribution latency prior to being sampled. In the HT3 path, the data stream is AC-coupled and equalized by a fast 1-bit speculative DFE working in conjunction with decision feedback restore (DFR) to combat DC wander from long data runs [4]. An Alexander phase detector steers a frequency-doubling interpolator-based DLL, similar to [5], for CDR function. The analog front-end preceding the data/edge samplers consists of summing amplifiers. Here, the main differential input comes from the RX bumps while the auxiliary input comes from a DAC for offset correction and DFE sampler threshold adjustment. The DAC settings are determined during power-up calibration.

In the HT1 path, the signal from the loopback channel is simply multiplexed into the signal path after the first stage of amplification. In the HT3 path, the loopback signal returns to the RX through the front-end amplifier auxiliary inputs. During loopback testing, each amplifier's main input and auxiliary input DAC are disabled. As in the TX, the loopback circuitry does not introduce additional load on the sensitive RX input bumps and consequently does not impact RX bandwidth or return loss.

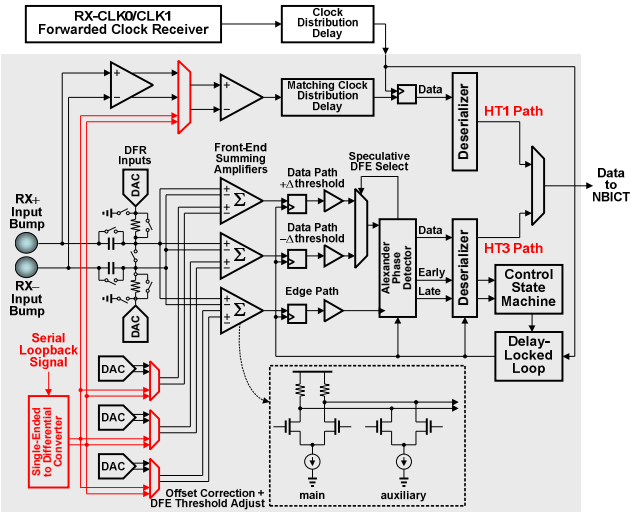

Fig. 9. RX architecture with loopback implementation.

IV. Silicon Results

The wafer-level loopback test has been enabled on an LTX Sapphire platform for screening 12-inch bumped Magny-Cours wafers. HT3 loopback was conducted at 5.2 and 6.4 Gb/s with supply voltages of 1.1 and 1.3 V. Following link training, the transceivers were exercised for a duration of 0.3–0.4 s with 10^8 cycles of alternating +K28.5 (1100000101) and −K28.5 (0011111010) symbols. A passing test returned no bit errors and BER $< 5\times10^{-10}$. Tester time prohibited us from establishing a much lower BER bound. HT1 loopback was not run since yield fallout from HT1 circuits is rare and screening for such infrequent failures was deemed an impractical use of tester time.

A sort example from an early wafer is shown in Table I. In this example, of the dies tested, seven did not pass the loopback test. Failures were essentially isolated to a single lane and occurred more frequently at higher frequency (6.4 Gb/s) and lower supply voltage (1.1 V) when operating margin is expectedly lower. The HT-Port0 all-lane failure in Die 3 suggests failure in a circuit block that is common to all lanes of both sublinks, such as the TX PLL. Revisiting Fig. 2, note that HT-Port2 of Die 7 would have been connected to an embedded link if it were packaged as a slave die. Also, since HT-Port3 is used only in the master die, Dies 1 and 5 are binned as slave dies to avoid being discarded for non-functional HT-Port3.

TABLE I
Early Example of HT Loopback Test Sort Results

Die No.	HT Loopback Fail Description
1	Port3 Sublink0 @ 6.4Gb/s – 1.1V CAD2 bit error count = 63 (saturated)
2	Port0 Sublink0 @ 6.4Gb/s – 1.1V CAD2 bit error count = 2
3	Port0 Sublink0/Sublink1 @ 5.2Gb/s, 6.4Gb/s – 1.1V, 1.3V Training failure in all CTL/CAD lanes
4	Port0 Sublink0 @ 6.4Gb/s – 1.1V, 1.3V CAD2 bit error count = 63 (saturated)
5	Port3 Sublink0 @ 6.4Gb/s – 1.1V CAD4 bit error count = 63 (saturated)
6	Port0 Sublink 0 @ 6.4Gb/s – 1.1V CAD0 bit error count = 2
7	Port2 Sublink0 @ 6.4Gb/s – 1.1V CAD4 bit error count = 63 (saturated)

V. Conclusion

We have presented a transceiver loopback implementation that enables wafer-level at-speed testing of HT I/O ports at no extra sort infrastructure cost and have successfully demonstrated test functionality at 6.4 Gb/s. This DFT feature delivers substantial cost reduction of packaging only known good dies primarily for more expensive MCM packages and provides package-level characterization and debug capability of embedded links for faster productization. The loopback test is now established for wafer-level screening of 45-nm SOI-CMOS processor products.

Acknowledgments

We are indebted to the Product Development Engineering team for establishing the loopback test in production. We thank Ari Shtulman for DFT verification support; Michael Parker and Heidi Grande for layout support; and Dicai Yang, Dean Gonzales, Tim Kasper, and Dennis Fischette for helpful information and review.

References

[1] http://www.hypertransport.org/

[2] S. Narasimha *et al.*, "High-performance 45-nm SOI technology with enhanced strain, porous low-k BEOL, and immersion lithography," *IEEE Int. Electron Devices Meeting*, Dec. 2006.

[3] W. L. Williams *et al.*, "Low latency clock domain transfer for simultaneously mesochronous, plesiochronous, and heterochronous interfaces," *IEEE Int. Symp. Asynchronous Circuits and Systems*, Mar. 2007.

[4] E. Fang *et al.*, "A 5.2Gbps HyperTransport™ integrated AC coupled receiver with DFR DC restore," *IEEE Symp. VLSI Circuits*, Jun. 2007.

[5] G.-Y. Wei *et al.*, "A variable-frequency parallel I/O interface with adaptive power-supply regulation," *IEEE J. Solid-State Circuits*, Nov. 2000.