A 60GHz On-Chip RF-Interconnect with λ/4 Coupler for 5Gbps Bi-Directional Communication and Multi-Drop Arbitration

Hao Wu[1], Lan Nan[1], Sai-Wang Tam[1], Hsieh-Hung Hsieh[2], Chewpu Jou[2], Glenn Reinman[1], Jason Cong[1], and Mau-Chung Frank Chang[1]
[1] University of California, Los Angeles, CA 90095, USA
[2] Taiwan Semiconductor Manufacturing Company, Hsinchu, Taiwan, ROC

***Abstract*—A 5Gbps bi-directional *RF-Interconnect* (RF-I) with multi-drop and arbitration capabilities is designed and realized in 65nm CMOS. The baseband data are modulated in RF-I by using a 60GHz carrier in ASK format. An on-chip differential transmission line (TL) is used as the communication channel, which minimizes the latency (9ps/mm) only under the speed-of-light limitation. We insert λ/4 directional couplers for implementing multi-drops without signal reflection. We also use MOS switches along the signal path to reconfigure/arbitrate communication priority for multi-drops. This design consists of four TX/RX drops along a 5.5mm TL ring, supports destructive reading with fixed priority, and can reconfigure any drop as the transmitter. The tested data rate of the RF-I is 5Gbps with lower than 10^{-12} BER. The average power consumptions for the link are 1.33pJ/b and 0.24pJ/b/mm.**

I. Introduction

On-chip interconnects, especially those used in future chip multi-processors and network-on-chip (NoC), have been projected to be the limiting factor in terms of bandwidth, power and latency. RF modulated and transmission-line-based interconnects (*RF-I*) have been demonstrated as superior in latency, scalability, re-configurability, bandwidth and power efficiency [1]-[3]. According to [2], RF-I global links combined with local RC wires provide the inter-core network with either 1.7X performance gain under the same power or 5X power savings under the same performance. Meanwhile, the large number of peer-to-peer on-chip interconnects in NoC becomes another major issue, which inevitably leads to large power and area consumption for connecting cores. Therefore, the next generation of NoC architecture demands the benefit of having on-chip high-speed interconnect with multi-drop and arbitration capabilities as data buses [2][4]. Hence an RF-I solution can be extremely beneficial by itself and can be further enhanced by adding the multi-drop and arbitration capabilities. The arbitration capability is especially critical so that all computing cores can share a common communication channel efficiently with fairness and without collision. Previously published multi-drop works are either too power-hungry (optical link involves power-consuming O/E and E/O conversions) [4], or has long latency and lacks arbitration capability because of reflections introduced by multiple drops at the baseband and unpredictable receiving priorities [5].

In this work, we propose a novel multi-drop RF-I with arbitration capability based on λ/4 directional couplers. This architecture offers superior scalability and re-configurability, while retaining RF-I's benefits of high data rate, low latency and low energy per bit. In this paper, section II will present the proposed RF-I architecture, on-chip λ/4 directional coupler and TX/RX circuit blocks; and section III presents measurement results; section IV summarizes the overall work.

II. RF-Interconnect with Multi-Drop Arbitration Architecture

Taking into account the aforementioned considerations, we designed and implemented a high-speed RF-I with 60GHz carrier and multi-drop arbitration capability in 65nm 1P6M GP CMOS with a 3.4μm thick top metal. The carrier frequency of 60GHz is carefully chosen as lower frequency results in larger coupler size incompatible with future NoC's core size and larger carrier to bandwidth ratio that potentially causes dispersion and requires power hungry equalization techniques; whereas higher frequency causes higher TL loss and higher VCO power consumption which implies degradation of link's power efficiency.

A. RF-I System with Multi-Drop Arbitration Capability

In this design, the RF-I with multi-drop and arbitration consists of four drops, each of which can perform as transmitter (multi-caster) or receiver. Fig. 1 shows the system block diagram.

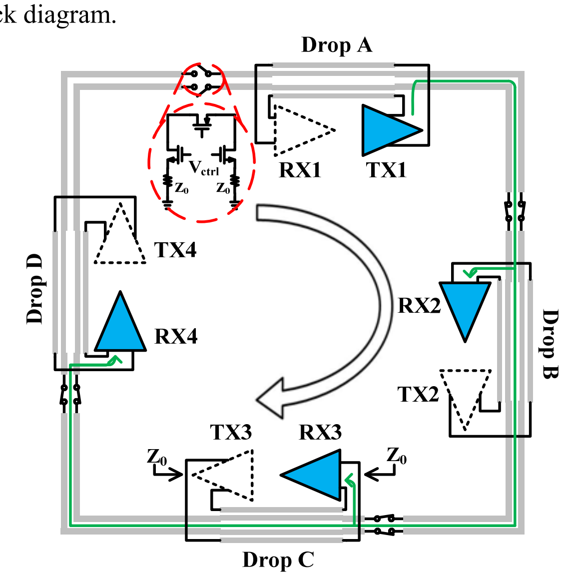

Fig. 1. 4-drop RF-I with Arbitration Capability (Drop A multicasts)

At each drop, a λ/4 directional coupler and a high-speed (5Gbps) ASK transceiver are employed. The λ/4 directional

978-1-4673-1556-2/12/$31.00@2012 IEEE

coupler is known for its directivity, isolation and matching properties at specific ports. The isolation and matching suppress reflection issues in multi-drop system, whereas directivity sets pre-defined priority. Digitally controlled switches are also implemented between adjacent drops for both impedance matching of arbitrary multi-cast and destructive reading arbitration. The switch has 1.5dB loss and 40dB isolation when it is on and off. With transmission-gate-based switches, signal on the main channel can be terminated to Z_0 (100Ω differentially) or passed to subsequent drops [2][4]. In this method the RF-I is implemented with a scalable daisy-chain arbitration scheme. Fig. 1 shows the scenario that drop A multi-casts to drops B, C, and D. Signal flow is highlighted in green. Inactive components (drop A' RX and other drops' TX) are powered off to save power.

With the switches, the system can adapt to the following fashions based on the global link condition in the NoC: 1) an arbitrary drop can be reconfigured as transmitter with a fixed set of priorities for the receivers by simply opening the switch before the transmitting node and closing the other ones; 2) drops with lower priorities can be destructed reading by opening the switch before them; 3) the physical link can be divided to two separated sections as A transmits to B whereas C communicates with D (the number of available sections increases when more nodes are attached to this RF-I bus). The two links will not interfere with each other since the isolation of switch is high enough. This architecture brings superior flexibility and scalability, which potentially saves considerable bandwidth, latency and power on the NoC interconnects [2]. Although 4 drops are implemented in this design due to silicon area constraints, more drops can be added as long as TX's output power is large enough for the last drop to detect the data. For real world applications, the number could be 8 or 16, whichever is more convenient for implementing an NoC.

B. Channel and On-Chip Directional Coupler

The main communication channel is made of differential transmission lines with a total length of 5.5mm between farthest nodes under multicast scenario.

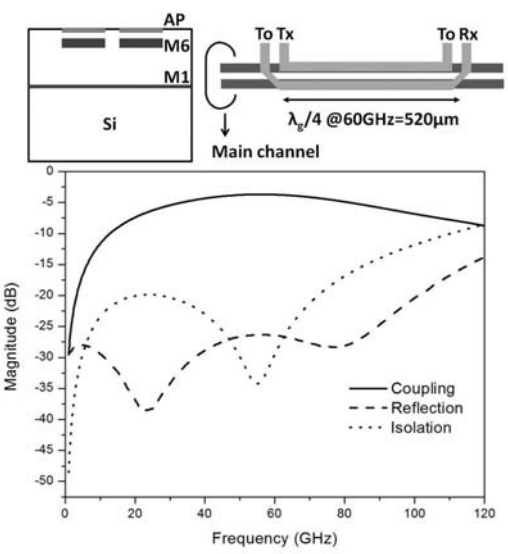

Fig. 2. On-chip directional coupler and its simulated performance

One challenge of high-speed transmission-line-based multi-drop interconnect is the signal reflections from various drops along the channel. In our design for example, the loss on the TL is 1.2dB/mm, and the distance between two adjacent drops is 1.5mm. The round trip path and switch loss is only -6.6dB which means reflection at each drop has to be low enough to achieve SNR of 18dB required by ASK [3]. On-chip directional coupler is introduced to solve the problem. We take advantage of the short wavelength of mm-wave carrier, so the on-chip coupler could be realized in a small size. The coupled lines are placed on top of the main channel, costing no extra silicon area and providing tight coupling. Due to the directivity of couplers, the signals will propagate in one direction along the channel. It also results in the priority ordering of drops.

In Fig. 2, a 3.7dB coupling loss is achieved at 60GHz, and it remains almost flat within a 10GHz bandwidth. The main channel reflection and isolation between the transmitter and receiver on the same drop are both below 26dB. It results in the reflection caused ISI below 32.6dB (reflection plus round trip path loss) per drop, so it is no longer the limiting factor to achieving the required SNR. This solves the reflection issues that have plagued traditional multi-drop interconnect designs.

C. 60GHz ASK RF Transceiver

To minimize the energy consumption and design complexity, a simple ASK transceiver is proposed. In contrast to other modulation such as BPSK, ASK is a non-coherent communication scheme that does not require any synchronization. It simplifies the transceiver architecture and reduces power consumption by avoiding any phase or frequency locking circuits. The carrier frequency is tuned to 60GHz to enable the small coupler size.

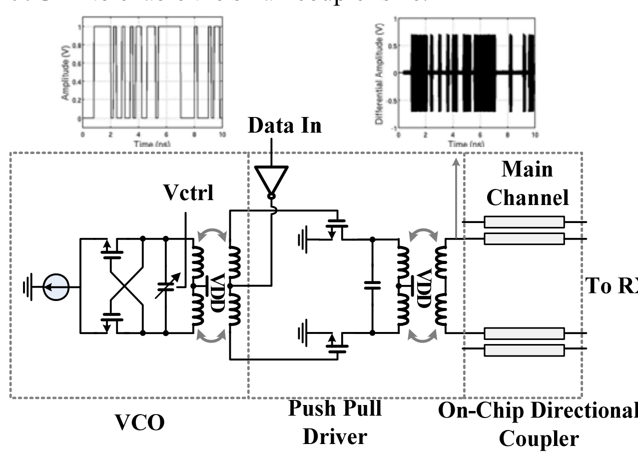

Fig. 3. 60GHz ASK transmitter

As shown in Fig. 3, the transmitter consists of a VCO and a simple single-stage, push-pull driver with embedded ASK modulator. The VCO is a free-running, cross-coupled differential pair with a resonant tank. In a non-coherent modulation scheme, such as ASK, it can tolerate output frequency drift up to several GHz. As long as the frequency is within the coupling band of the coupler and receiving band of the receiver, the signal can be demodulated.

The transmitter VCO's output is inductively coupled to the differential push-pull driver through an on-chip transformer with a 1-to-1 ratio. The input data bits are directly fed to the center tap of the transformer's secondary coil. The digital data bits modulate the on/off state of the driver by modulating the bias (common mode) of the input transistors. By combining the modulator and driver into one single stage, this design avoids an additional mixer to perform ASK modulation, and therefore achieves high efficiency and simplicity. At the end, the modulated ASK signal is inductively fed into the directional coupler.

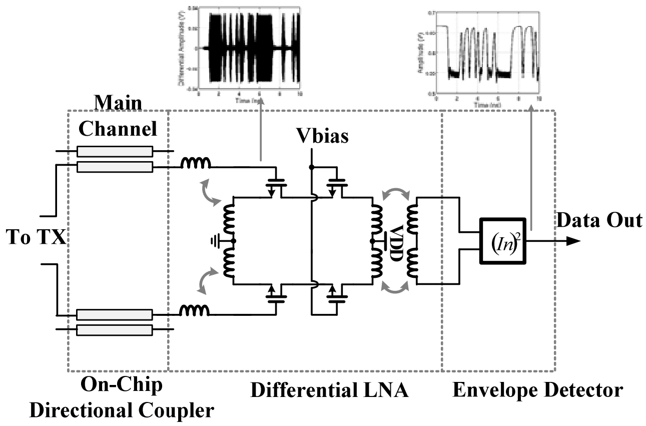

Fig. 4. 60GHz ASK receiver

The receiver (Fig. 4) consists of a single-stage low noise amplifier and an envelope detector. The output power from the transmitter is around 3dBm, and the path loss from the transmitter to the farthest receiver (worst case) is 27dB. Therefore, the received power at the input of the LNA is about -24dBm. The differential common source LNA with inductive degeneration has an 18dB gain at 60GHz, 9GHz bandwidth and 6dB noise figure. The cascode configuration offers better stability. From link budget analysis, the farthest distance this architecture can support is more than 12mm with 8 drops, assuming adjacent drops are 1.5mm apart. After the LNA, the amplified signal is then fed into a differential mutual mixer to detect the envelope. From simulation, this detector is able to recover a 60GHz ASK signal's envelope up to 6Gbps, and the input sensitivity is about $10mV_{pk\text{-}pk}$. Because the mixer's devices work at sub-threshold, its power consumption is very low. Entire RX consumes power of 5mW.

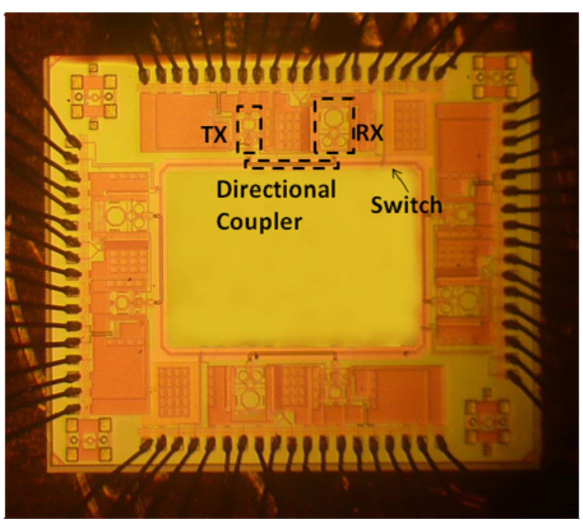

Fig. 5. 5Gbps RF-I with multi-drop and arbitration die-photo

III. Measurement Results

The RF-I with 60GHz carrier and multi-drop arbitration capabilities is fabricated in the TSMC 65nm 1P6M process with the die-photo in Fig. 5. The transmitter and receiver occupy core areas of $0.0048mm^2$ and $0.034mm^2$ respectively. The link can operate properly with conventional digital logic circuits placed directly under its passive structure, which gives better area utilization [3]. Besides the entire RF-I, a stand-alone transmitter test-chip is measured to characterize the frequency tuning of the VCO, and output power of the driver.

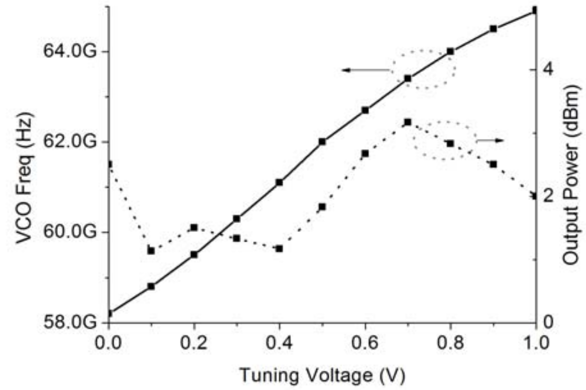

Fig. 6. Transmitter output frequency and calibrated output power

Fig. 6 shows the measured TX output frequency and output power with different tuning voltages. In the real RF-I application, this frequency does not have to change with time. For the driver, its output power is the main metric of interest because it determines the link budget. The output power is reasonably flat within the band of interest. Its ripple is within 3dB across the VCO's tuning range; therefore, the driver's bandwidth is large enough for this application. The designed driver is able to deliver more than 3dBm across the band.

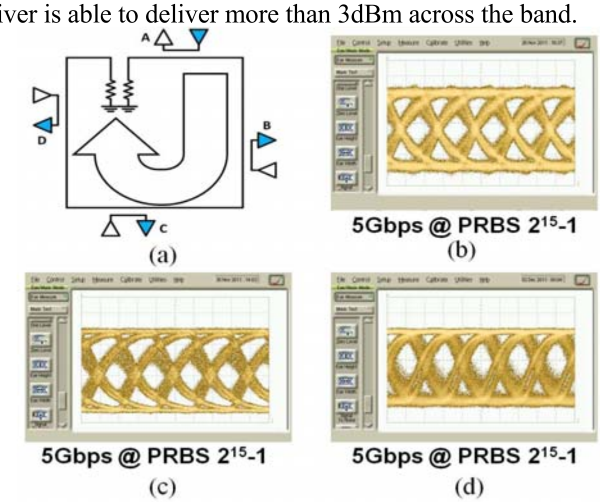

Fig. 7. Eye diagrams when drop A multi-casts to drops B, C, D

For the complete RF-I with multi-drop arbitration, eye diagrams and BER at 5Gbps for all drops in various configurations are measured to demonstrate the multi-drop and arbitration capability. Fig. 7 shows the diagrams when drop A transmits and drops B, C, D receive. This demonstrates the capability of multi-drop. Fig.10 shows that the BER is lower than 10^{-12}. Link analysis and BER also imply that the link bandwidth is limited by speed of the circuitry rather signal's

SNR. Therefore more drops can be added into this multi-drop system without degrading link bandwidth.

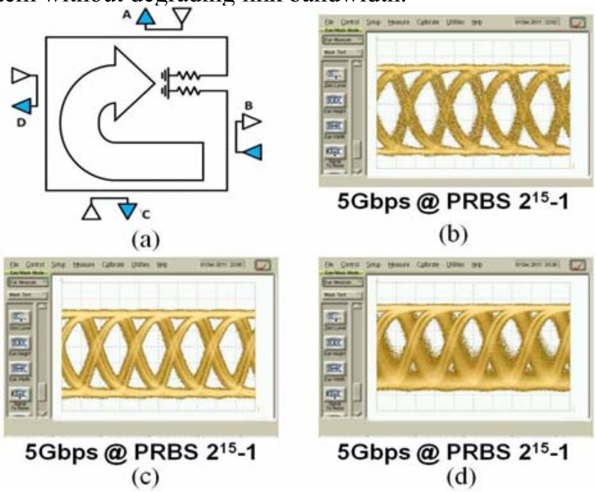

Fig. 8. Eye diagrams when drop B multi-casts to drops C, D, A

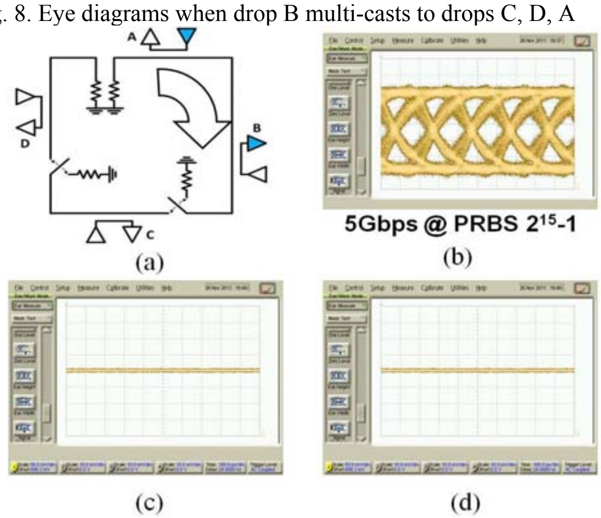

Fig. 9. Eye diagrams when drop A transmits and only B receives

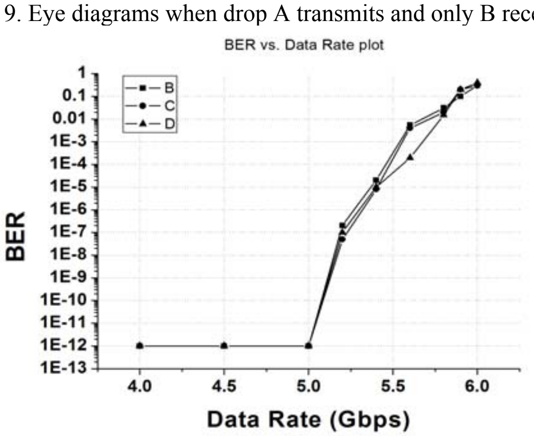

Fig. 10. Measured BER vs. data rate at different drops

Fig. 8 shows the eye diagrams when drop B transmits and drops C, D and A receive. This demonstrates the capability of an arbitrary drop multi-casts function. In fact, it is expected that all drops will perform exactly the same because of the structure symmetry. Finally, destructive reading is demonstrated in the scenario that drop A transmits to B (Fig. 9) and drops C, D are blind, as the switch before C is open. It also implies that the system can operate as multiple non-interfering links (A->B, C->D). This function is not measured due to testing equipment constraints.

IV. Conclusion

To summarize, a 5Gbps *RF-Interconnect* with 60GHz carrier and multi-drop arbitration functions has been realized. With a PRBS pattern of 2^{15}-1, the BER is lower than 10^{-12} for all four drops. This is the first CMOS on-chip interconnect demonstrated with full multi-drop and arbitration capabilities. It is highly scalable and reconfigurable from the NoC architecture point of view. Meanwhile, its power consumption is comparable to the state-of-the-art baseband-only interconnect, but with superior latency, extra multi-drop and arbitration capability (Table 1).

TABLE I
Comparison of State-of-the-art

	This work	[5]	[7]
Category	**RF**	Baseband	Baseband
Multi-drop	**Yes**	Yes	No
Arbitration Cap.	**Yes**	No	No
Data rate (Gbps)	**5**	8	5
Power (pJ/b/mm)	**0.24***	0.18	0.21
Latency (ps/mm)	**9**	30	50

*VCO power (5mW) can be shared by all (many tens) parallel RF-I links in NoC and does not burden individual link significantly.

Acknowledgment

This work is supported in part by the Center for Domain Specific Computing, funded by the U.S. National Science Foundation Expedition in Computing Award CCF-0926127. We would like to thank Dr. Zhiwei Xu, Dr. Ning-Yi Wang and Dr. Adrian Tang for their technical discussions and Yen-Hsiang Wang for the testing board preparations.

References

[1] M. F. Chang, J. Cong, A. Kaplan, M. Naik, G. Reinman, E. Socher, and S.-W. Tam, "CMP network-on-chip overlaid with multi-band RF-Interconnect," *IEEE HPCA Sym*, pp. 191-202, Feb. 2008

[2] M. F. Chang, J. Cong, A. Kaplan, C. Liu, M. Naik, J. Premkumar, G. Reinman, E. Socher, and S.-W. Tam, "Power reduction of CMP communication networks via RF-Interconnect," *IEEE/ACM Sym. Micro.,* pp.376-387, Nov. 2008

[3] S.-W. Tam, E. Socher, A. Wong, and M. F. Chang, "A simultaneous tri-band on-chip RF-Interconnect for future Network-on-Chip", *VLSI Sym.*, pp. 90-91, Jun. 2009

[4] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: Illuminating future Network-on-Chip with Nanophotonics," *IEEE ACM ISCA*, pp. 429-440, Jun. 2009

[5] H. Ito, M. Kimura, K. Miyashita, T. Ishii, K. Okada, K. Masu, "A bidirectional- and multi-drop-transmission-line interconnect for multipoint-to-multipoint on-chip communications," *IEEE J. of Solid-State Circuits,* Vol.43, No.4, pp. 1020-1029, Apr. 2008

[6] J. Postman and P. Chiang, "A survey addressing on-chip interconnect: energy and reliability considerations," *ISRN Electronics,* vol. 2012, pp. 1-9, 2012

[7] D. Schinkel, E. Mensink, E. Klumperink, E. Tuijl, and B. Nauta, "Low-power, high-speed transceivers for Network-on-Chip Communication," *IEEE Trans. VLSI Systems,* vol. 17, no. 1, pp. 12-21, Jan. 2009