

Multiband RF-Interconnect for Reconfigurable Network-on-Chip Communications

Jason Cong¹, Mau-Chung Frank Chang², Glenn Reinman¹, and Sai-Wang Tam²,
Computer Science Department, University of California, Los Angeles, CA 90095, USA¹
Electrical Engineering Department, University of California, Los Angeles, CA 90095, USA²,
cong@cs.ucla.edu

Categories & Subject Descriptors: C.1.4
[PROCESSOR ARCHITECTURE]: Parallel Architectures,
B.7.1 [INTEGRATED CIRCUIT]: Types and Design Styles

General Terms: Design, Experimentation, Measurement, Performance

Keywords: Chip Multiprocessor, FDMA, MORFIC, Network-on-Chip, RF-Interconnect, Reconfigurable

One of the key benefits of the scaling of CMOS is that the switching speed of the transistor improves over each technology generation. According to ITRS, f_T and f_{max} , will be 600GHz and 1 THz, respectively, in 16nm CMOS technology. With the advance in CMOS mm-wave circuits, hundreds of GHz bandwidth will be available in the near future. In addition, compared with CMOS repeaters charging and discharging the wire, EM waves travel in a guided medium at the speed of light which is about 10ps/mm on silicon substrate. The question here is: How can we utilize over hundreds of GHz of bandwidth in a future mobile system through RF-I while concurrently achieving ultra-low power operation and dynamic allocation in bandwidth to meet future Network-on-Chip needs?

One of the possibilities is to use multi-band RF-I [1][2][3][4], based on frequency-division-multiple-access algorithms (FDMA) to facilitate inter-core communications on-chip. The main advantages of RF-I include:

- Superior Signal to Noise ratio: Since all data streams modulate RF-carriers, which are at least 10GHz above the baseband, the high speed RF-interconnect does not generate and/or suffer from any baseband switching noise. This reduces possible interference to the sensitive near/sub- V_{th} operated circuit.
- High bandwidth: A multi-band RF-interconnect link has a much higher aggregate data rate than a single repeater buffer link. For example, the router of the RF-I node can first collect all the data from the nearby ULP core. Utilizing the superior bandwidth of RF-I, one RF-I node is able to handle a large amount of data generated from multiple ULP cores.
- Low Power: Compared to a repeater buffer, a multi-band RF-interconnect is able to operate at much better energy per bit in the NoC. This is especially true in ULP cores, where only several RF-nodes are enough to satisfy the

bandwidth demand between the ULP cores and the rest of the NoC. Compared to normal repeated wire networks, which consume considerable amounts of power, a few RF-I nodes only consume a very small amount of power. (see Section IV, benchmarked using pJ/bit as a metric)

- Low Overhead – High data rate/wire and low area/Gigabit and low latency due to speed-of-light data transmission (see Section IV, benchmarked using Area/(Gbit/sec) as a metric)
- Re-configurability – Efficient simultaneous communications with adaptive bandwidths via shared on-chip transmission lines
- Multicast support – Scalable means to communicate from one transmitter to a number of receivers on chip

Total compatibility and scalability with mainstream digital CMOS technology

The concept of RF-I is based on transmission of *waves*, rather than voltage signaling. When using voltage signaling in conventional RC time constant dominated interconnects, the entire length of the wire has to be charged and discharged to signify either ‘1’ or ‘0’. In the RF approach, an electro-magnetic (EM) wave is continuously sent along the wire (treated as a transmission line). Data is modulated onto that carrier wave using amplitude and/or phase changes.

A simultaneous tri-band on-chip RF-interconnect for future network-on-chip [5] is demonstrated. Two RF bands in mm-wave frequencies, 30GHz and 50GHz, are modulated using amplitude-shift keying, while the base-band utilizes a low swing capacitive coupling technique. Each RF-band and base-band carries 4Gbps and 2Gbps respectively. Three different bands, up to 10Gbps, are transmitted simultaneously across a shared 5mm on-chip differential transmission line. The energy per bit is 0.125pJ/b/mm in base-band, while RF-band is 0.09pJ/b/mm. Based on this demonstration, it is possible to improve bandwidth efficiency using N-channel multi-carrier RF-I. Those N distinct channels transmit N different data streams onto the same transmission line. The total aggregate data rate (R_{Total}) equals to $R_{Total} = R_{baseband} \times N$, where the data rate of each base-band is $R_{baseband}$ and the number of channels is N.

As an example of use of RF- reconfigurability, we recently proposed MORFIC (Mesh Overlaid with RF Inter-Connect)[2][4], a hybrid NoC design. It is composed of a traditional mesh of routers augmented with a shared pool of RF-I that can be configured as short-cuts within the mesh. In this design, we have 64 computing cores, 32 cache memory modules and 4 memory output ports – and RF-I is a bundle of transmission lines spanning the mesh, and features 16 carrier frequencies. We examined four architectures: (1) Mesh Baseline – a baseline mesh architecture without any RF-I, (2)

Mesh Wire Baseline – the baseline mesh architecture with express shortcuts between routers (conventional wire, not RF-I) that are chosen at chip design time (i.e. no adaptability to application variation), (3) Mesh Static Shortcuts – the same express shortcuts as the Mesh Wire Baseline but using RF-I instead of conventional repeated wire, and (4) Mesh Adaptive Shortcuts – the overlaid RF-I with shortcuts tailored to the particular application in execution. From the simulation results of our in-house cycle-accurate simulator [6], we demonstrated a significant performance improvement of the Mesh Adaptive Shortcuts over the Mesh Baseline, an average packet latency reduction of 20-25% [2], through the reconfigurable RF-I. We further demonstrated a 65% power reduction [4] by reducing the bandwidth of the baseline mesh by 75% - reducing the 16 Byte wide to 4 Byte wide baseline mesh. Our continued exploration of the MORFIC architecture will be instrumental in gauging future CMP interconnect design tradeoffs, and in better quantifying what benefits CMPs can expect from MORFIC in future generations of CMOS technologies down the road.

References

- [1] Jenwei Ko, Jongsun Kim, Zhiwei Xu, Qun Gu, C. Chien, M.F. Chang, "An RF/baseband FDMA-interconnect transceiver for reconfigurable multiple access chip-to-chip communication," *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, vol., no., pp.338-602 Vol. 1, 10-10 Feb. 2005
- [2] M. Frank Chang, Jason Cong, Adam Kaplan, Mishali Naik, Glenn Reinman, Eran Socher, Sai-Wang Tam, "CMP Network-on-Chip Overlaid With Multi-Band RF-Interconnect," *IEEE International Conference on High Performance Computer Architecture Sym*, pp.191-202, Feb. 2008.
- [3] M.F. Chang, et al., "RF/Wireless Interconnect for Inter- and Intra-chip Communications," *Proceedings of the IEEE*, Volume 89, Issue 4:456 – 466, April 2001
- [4] M. Frank Chang, et al., "Power Reduction of CMP Communication Networks via RF-Interconnects", *Proceedings of the 41st Annual International Symposium on Microarchitecture (MICRO)*, Lake Como, Italy, pp.376-387, November 2008.
- [5] Sai-Wang Tam, et al., "A Simultaneous Tri-band On-Chip RF-Interconnect for Future Network-on-Chip," *Accepted to be published in VLSI Symposium*, 2009.
- [6] J. Cong, et al., "MC-Sim: An Efficient Simulation Tool for MPSoC Designs," *IEEE/ACM International Conference on Computer-Aided Design*, pp.364-371, 2008.

Speaker Bio

JASON CONG received his B.S. degree in computer science from Peking University in 1985, his M.S. and Ph. D. degrees in computer science from the University of Illinois at Urbana-Champaign in 1987 and 1990, respectively. Currently, he is a Chancellor's Professor at the Computer Science Department of University of California, Los Angeles, and a co-director of the VLSI CAD Laboratory. He also served as the department chair from 2005 to 2008.