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

Jason Cong (cong@cs.ucla.edu)

Joint work with Frank Chang, Glenn Reinman and Sai-Wang Tam

UCLA

Communication Challenges

- On-Chip Issues
  - # Cores in Chip-Multiprocessor (CMP) growing
    - Increasing bandwidth demand on interconnect
  - Wires scaling poorly compared to transistors
    - Increased latency to communicate between distant points on CMP
- Off-chip limited by chip-to-chip, board-to-board, board-to-backplane communications
- Requirements on future interconnect
  - Scalable, reliable
  - Support high traffic volume with low latency
  - Constrained by
    - Power
    - Silicon Area
    - Cost (compatibility with mainstream CMOS technology)

Used vs. Available Bandwidth in Modern CMOS

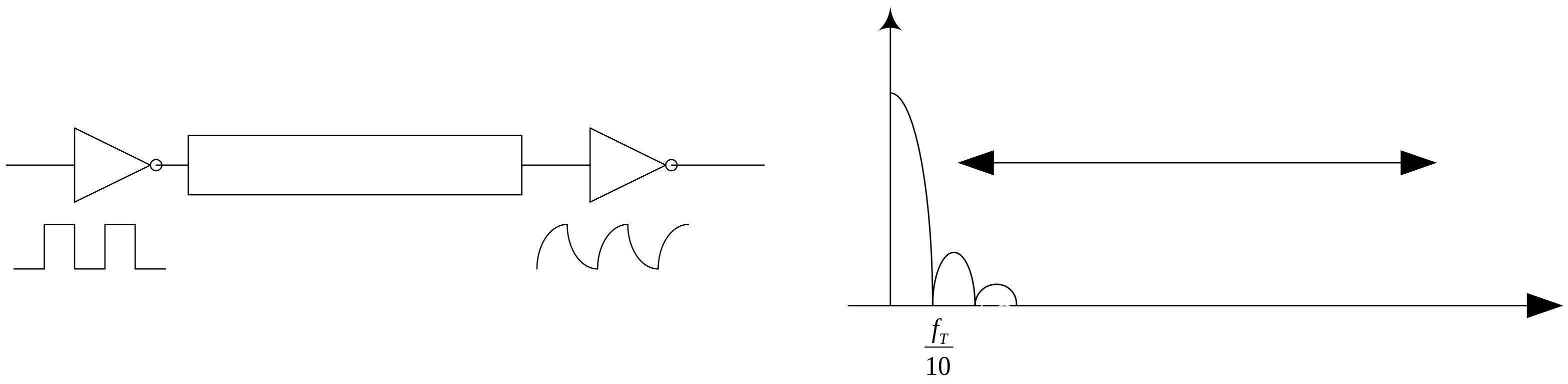

- @ 45nm CMOS Technology
  - Data Rate: 4 Gbit/s
  - f_T of 45nm CMOS can be as high as 240GHz
  - Baseband signal bandwidth only about 4GHz
  - 98.4% of available bandwidth is wasted
- Question: How to take advantage of full-bandwidth of modern CMOS?

UCLA 90nm CMOS VCO at 324GHz (ISSCC 2008*)

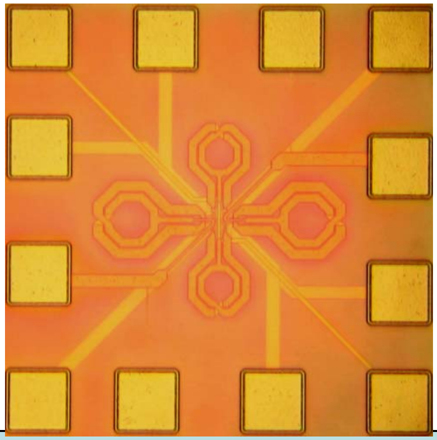

CMOS VCO designed by Frank Chang's group at UCLA, fabricated in 90nm process

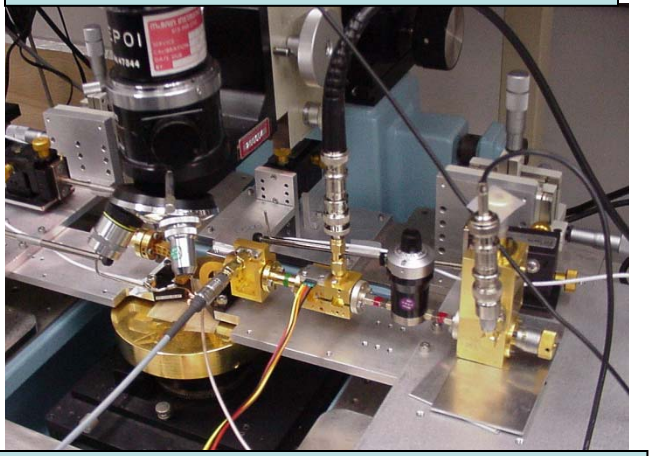

On-Wafer VCO Test Setup at JPL

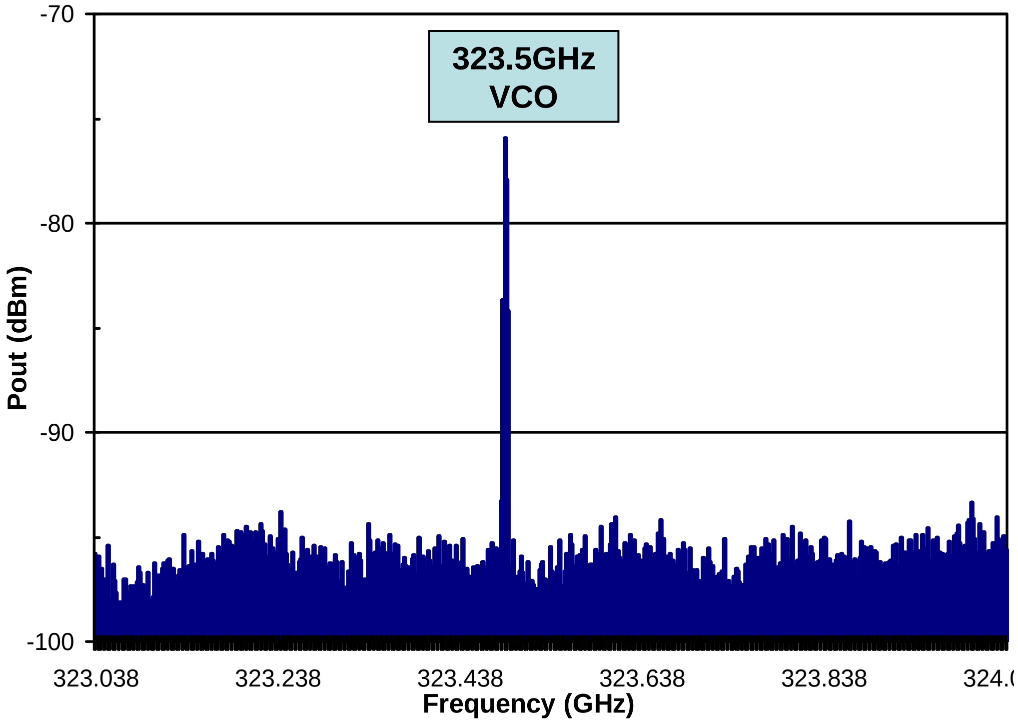

CMOS Voltage Controlled Oscillator, measured with a subharmonic mixer and driven with a 80 GHz synthesizer local oscillator. The mixing frequency is (f_{VCO} - 4*f_{LO})=f_{IF}, or f_{VCO} -4*(80 GHz)= 3.5 GHz, yielding f_{VCO}= 323.5 GHz!

*Huang, D., LaRocca T., Chang, M.-C. F., "324GHz CMOS Frequency Generator Using Linear Superposition Technique IEEE International Solid-State Circuits Conference (ISSCC), 476-477, (Feb 2008) San Francisco, CA

Multiband RF-Interconnect

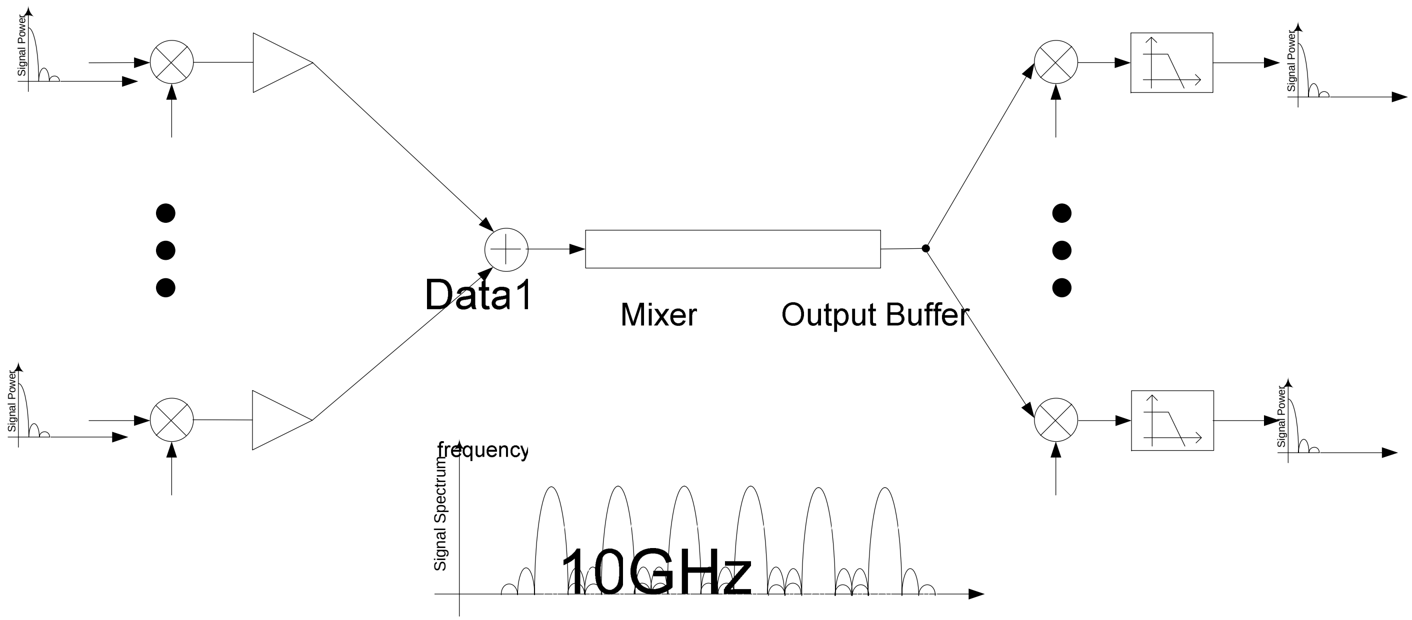

- In TX, each mixer up-converts individual baseband streams into specific frequency band (or channel)
- N different data streams (N=6 in exemplary figure above) may transmit simultaneously on the shared transmission medium to achieve higher aggregate data rates
- In RX, individual signals are down-converted by mixer, and recovered after low-pass filter

UCLA

RF-Interconnect Demonstrations

- **Off-chip (On-board) Simultaneous Dual-band Communications through RF-Interconnect (ISSCC 05)**
- **Inter-layer 3DIC RF-Interconnect (ISSCC 07)**
- **On-chip Simultaneous generation of multi-band carriers (RFIC 08)**
- **On-Chip Tri-band simultaneous communications (VLSI 2009)**

Tri-Band On-Chip RF-Interconnect (VLSI 2009*)

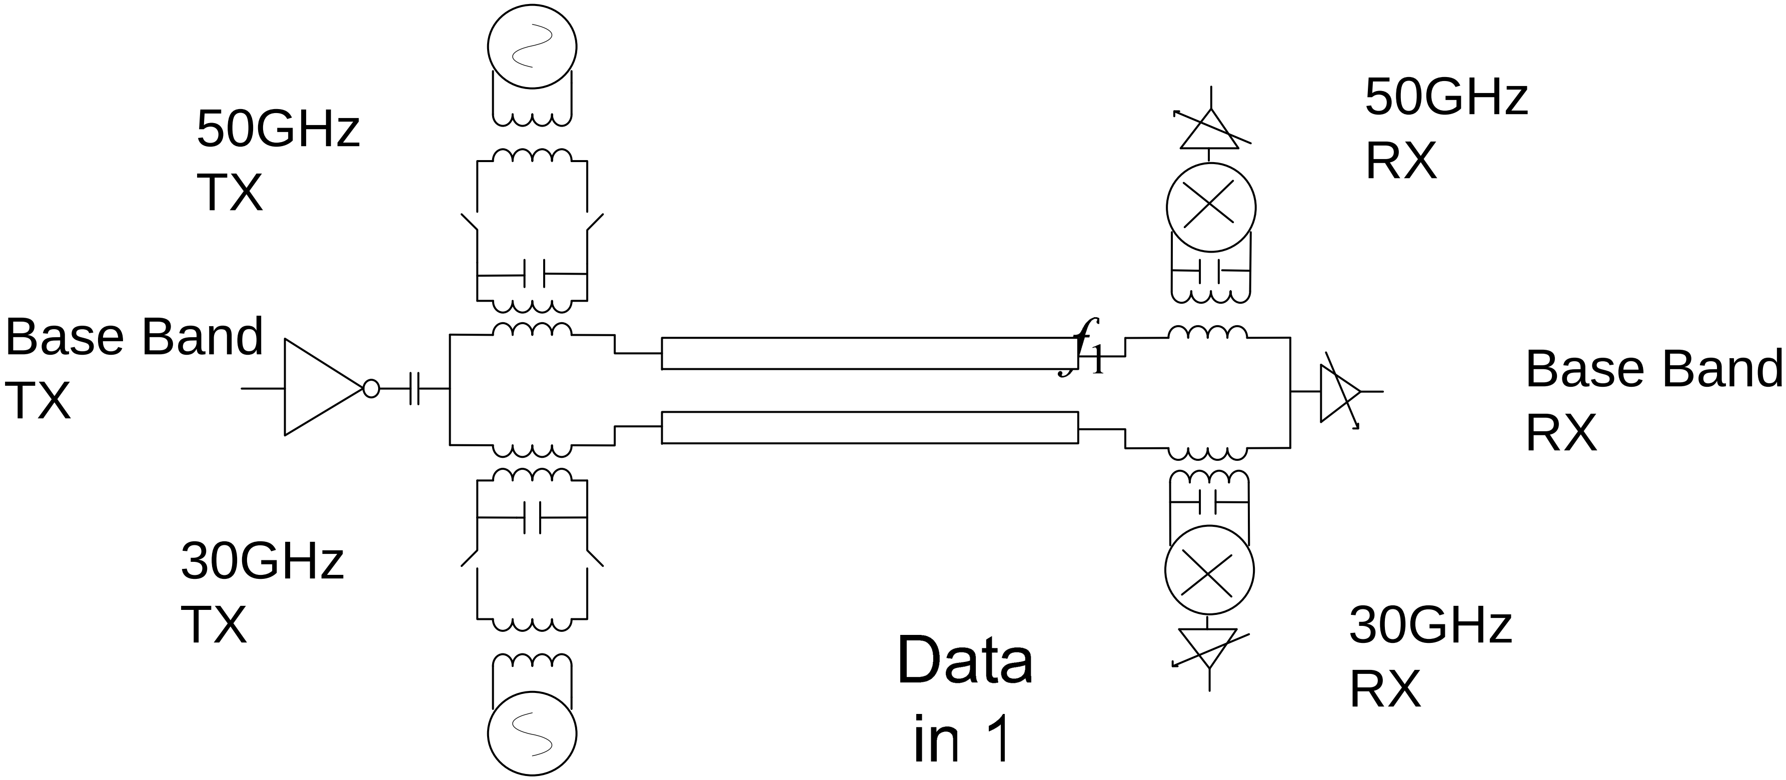

- **IBM 90nm digital CMOS process**
- **5mm differential transmission Line**
- **Total 3 Channels: 2RF + 1Baseband**
- **Differential mode for RF: 30GHz and 50GHz**
- **Common mode for baseband**
- **Total aggregate data rate is 10Gb/s1**

Data in 3

Transmission

* Sai-Wang Tam, Eran Socher, Alden Wong, M.-C.Frank Chang,
"A Simultaneous Tri-Band On-Chip RF-Interconnect for Future Network-On-Chip," IEEE VLSI Symposium 2009

Data

Tri-band On-Chip RF-I Test Results

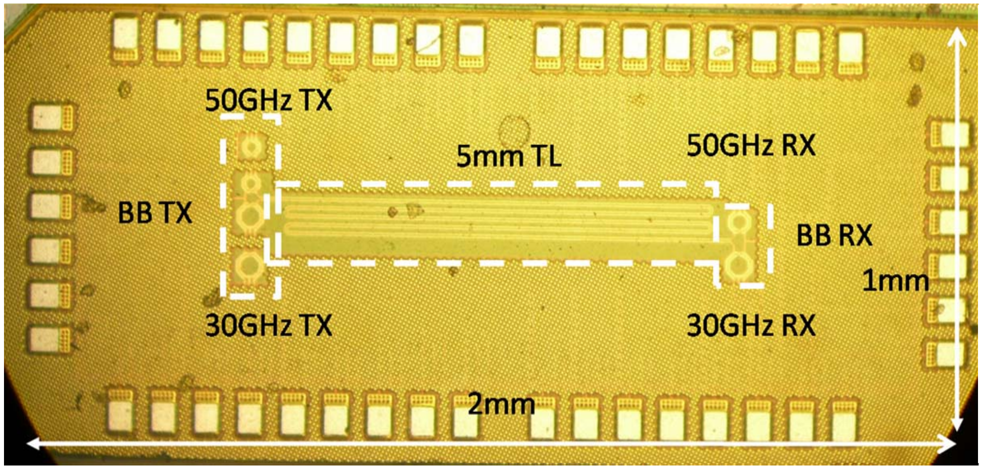

Process	IBM 90nm CMOS Digital Process
Total 3 Channels	30GHz, 50GHz, Base Band
Data Rate in each channel	RF Band: 4Gbps Base Band: 2Gbps
Total Data Rate	10Gbps
Bit Error Rate	Across all Bands <10E-9
Latency	6 ps/mm
Enegry Per Bit (RF)	0.09*pJ/bit/mm
Enegry Per Bit (BB)	0.125pJ/bit/mm

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

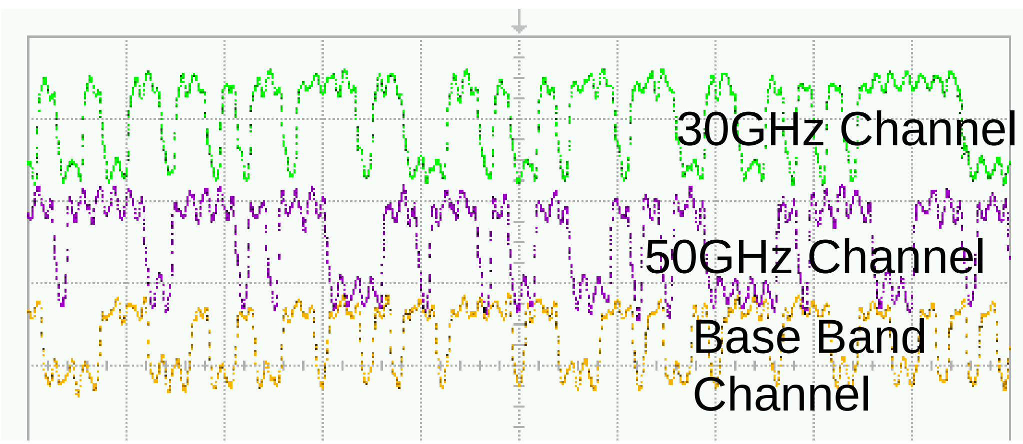

Data Output waveform

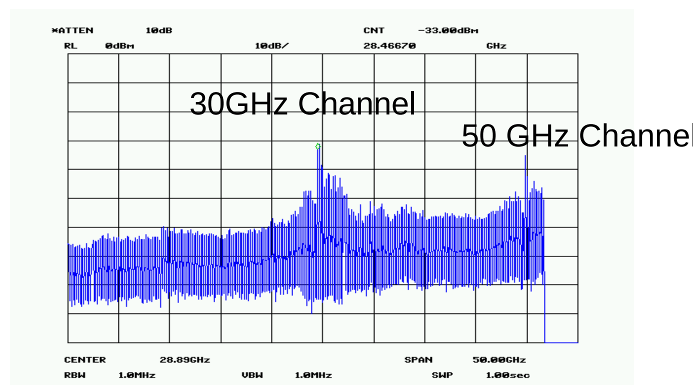

Output Spectrum of the RF-Bands, 30GHz and 50GHz

Multi-band ASK RF-I Scaling

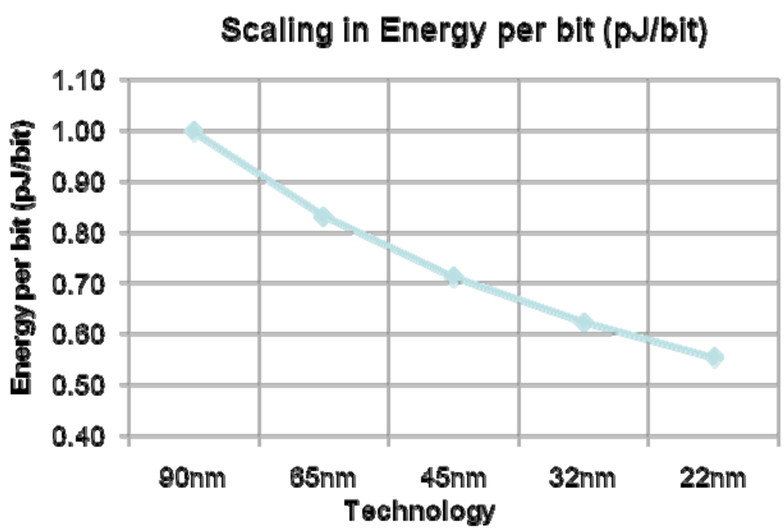

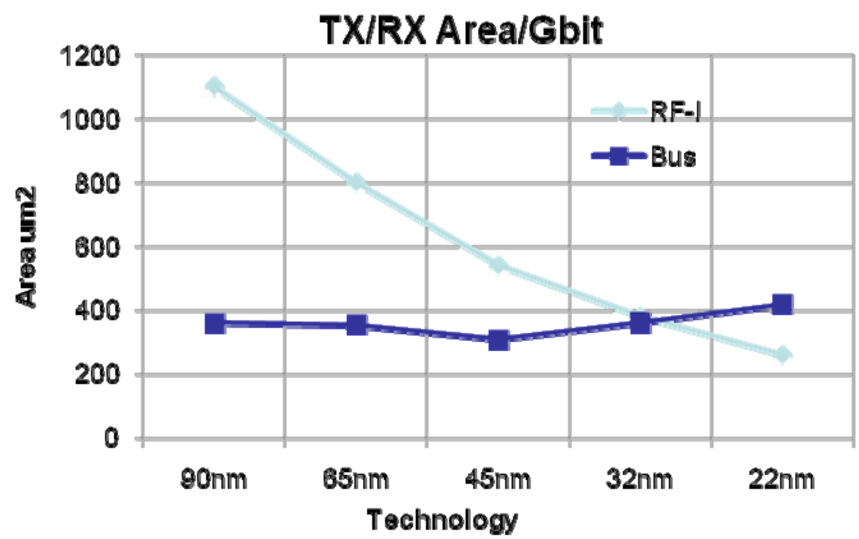

Technology	# of Carriers	data rate per carrier (Gb/s)	Total Data rate per wire (Gb/s)	Power (mW)	Energy per bit(pJ)	Area (TX+RX) mm^2	Area/Gbit (μm^2/Gbit)
90nm	3RF + 1 BB	5	20	20	1.00	0.022	1100
65nm	4RF + 1 BB	6	30	25	0.83	0.024	800
45nm	5RF + 1 BB	7	42	30	0.71	0.023	540
32nm	6RF + 1 BB	8	56	35	0.63	0.021	380
22nm	7RF + 1 BB	9	72	40	0.56	0.019	260

UCLA

Comparison between Repeated Bus and Multi-band RF-I @ 32nm

	RF-I	Repeated Bus
# of wire	13	448
Data rate per carrier (Gbit/s)	8	NA
# of carrier	7	NA
Data rate per carrier (Gbit/s)	56	1
Aggregate Data Rate	728	768
Bus Physical Width	160	160
Transceiver Area (mm^2)	0.27	0.022
Power (mW)	455	6144
Energy per bit (pJ/bit)	0.63	8

Interconnect length = 2cm

Assumptions:

1. 32nm node; 30x repeater, FO4=8ps, Rwire = 306Ω/mm Cwire = 315fF/mm, wire pitch=0.2um, Bus length = 2cm, f_bus = 1GHz, Bus Width 96Byte
2. Repeaters Area = $0.022mm^2$
3. Bus physical width = 160um
4. In that width we can fit 13 transmission line, each with 7 carriers with carrying 8Gbps

UCLA

Architectural Considerations for RF-I

- Opportunities (both on and off chip)
  - High bandwidth communication
    - Data distribution across many-core topologies
    - Vital in keeping many-core designs active
  - Low latency communication
    - Enables users to apply parallel computing to a broader applications through faster synchronization and communication
    - Faster cache coherence protocols
  - Reconfigurability
    - Adapt NoC topology/bandwidth to the needs of the individual application
  - Power efficient communication
- Challenges
  - Frequency arbitration and Tx/Rx tuning
  - Application-specific modeling

Simple RF-I Topology

- Four NoC Components
- Tunable Tx/Rx's
  - Arbitrary topologies
  - Arbitrary bandwidths

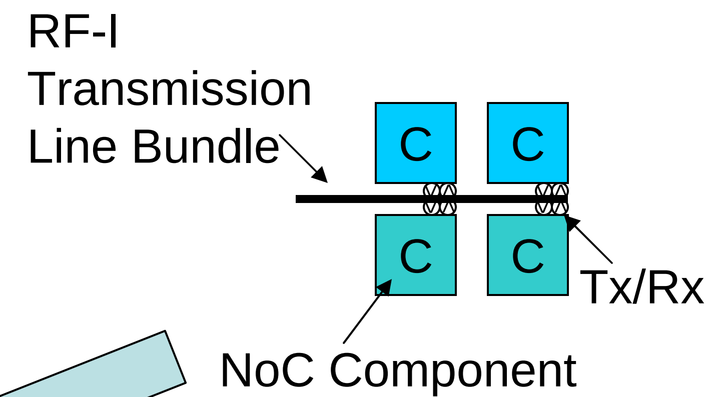

One physical topology can be configured to many virtual topologies

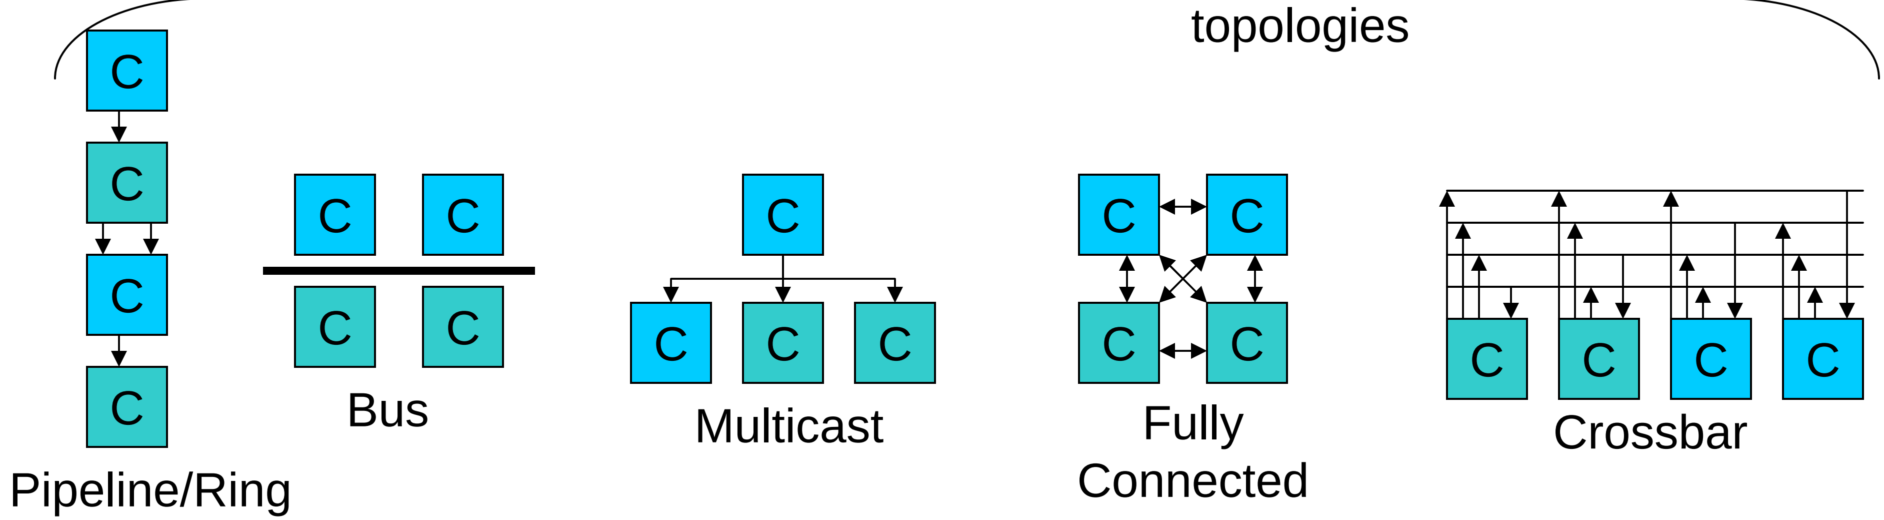

Mesh Overlaid with RF-I [HPCA'08]

- 10x10 mesh of pipelined routers
  - NoC runs at 2GHz
  - XY routing
- 64 4GHz 3-wide processor cores
  - Labeled aqua
  - 8KB L1 Data Cache
  - 8KB L1 Instruction Cache
- 32 L2 Cache Banks
  - Labeled pink
  - 256KB each
  - Organized as shared NUCA cache
- 4 Main Memory Interfaces
  - Labeled green
- RF-I transmission line bundle
  - Black thick line spanning mesh

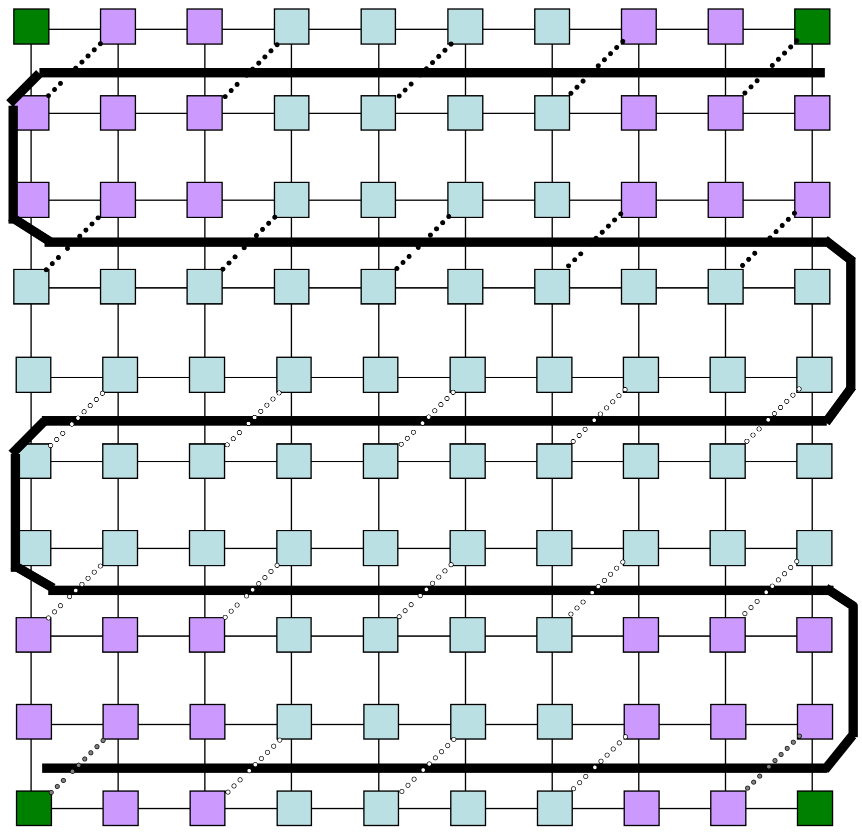

RF-I Logical Organization

- Logically:
  - RF-I behaves as set of N express channels
  - Each channel assigned to src, dest router pair (*s*,*d*)

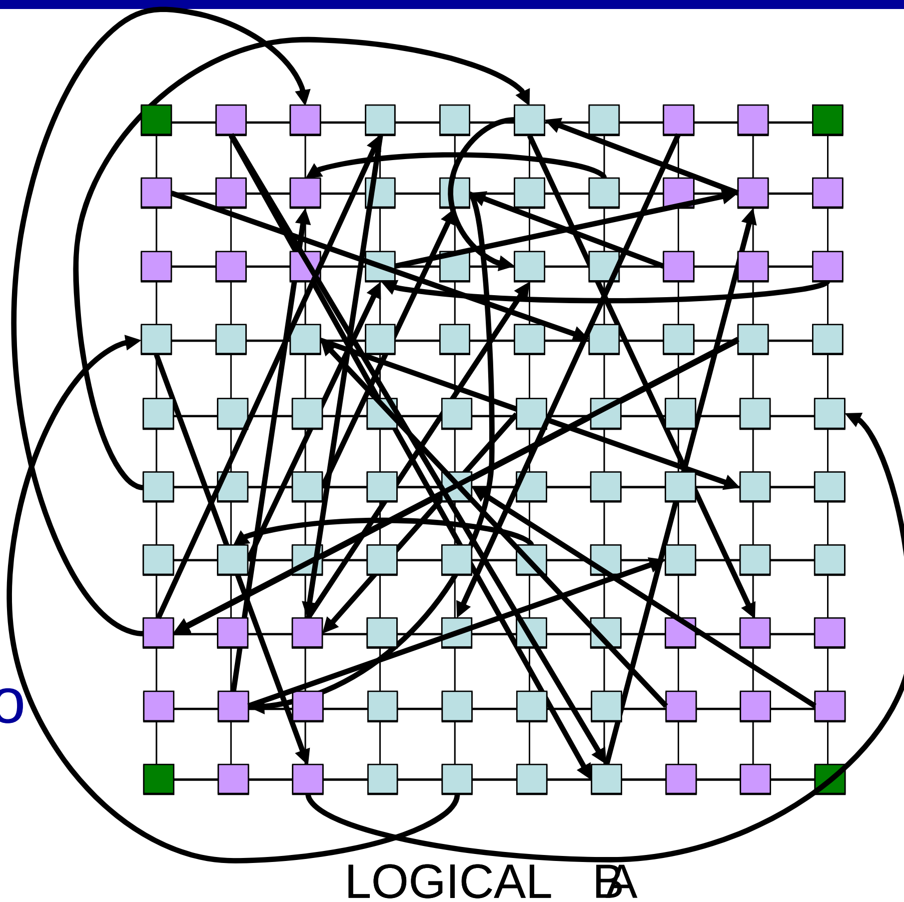

- Reconfigured by:
  - remapping shortcuts to match needs of different applications

Power Savings [MICRO'08]

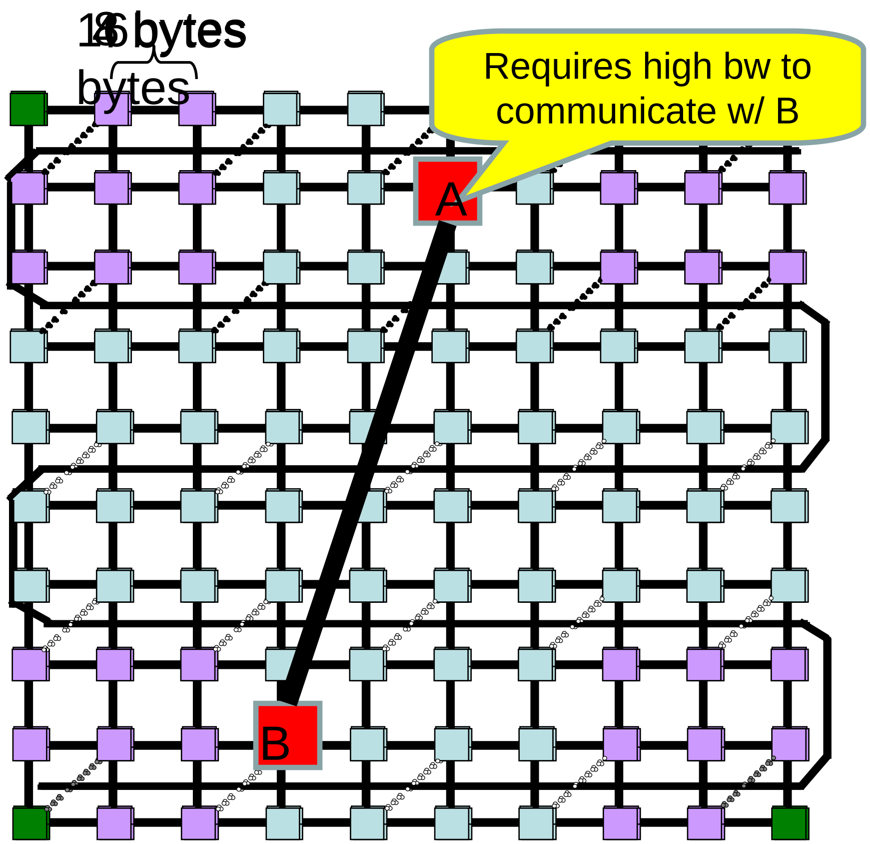

- We can thin the baseline mesh links
  - From **16B…**
  - …to **8B**
  - …to **4B**

- **RF-I makes up the difference in performance while saving overall power!**
  - **RF-I provides bandwidth where most necessary**
  - **Baseline RC wires supply the rest**

RF-I Enabled Multicast

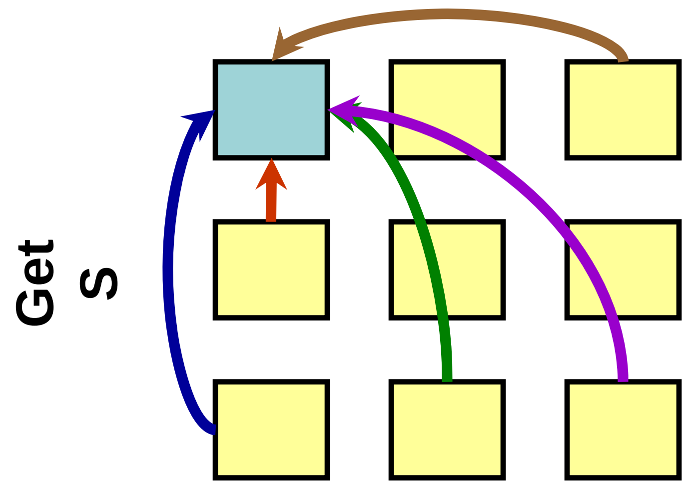

Request Scenario

Conventional NoC

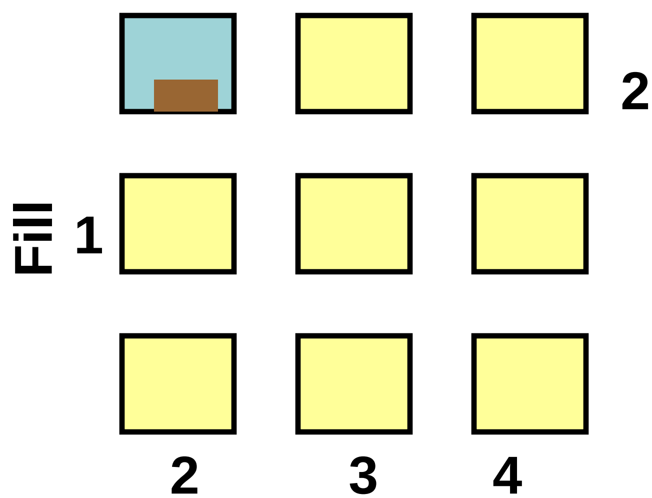

RF-I enabled NoC

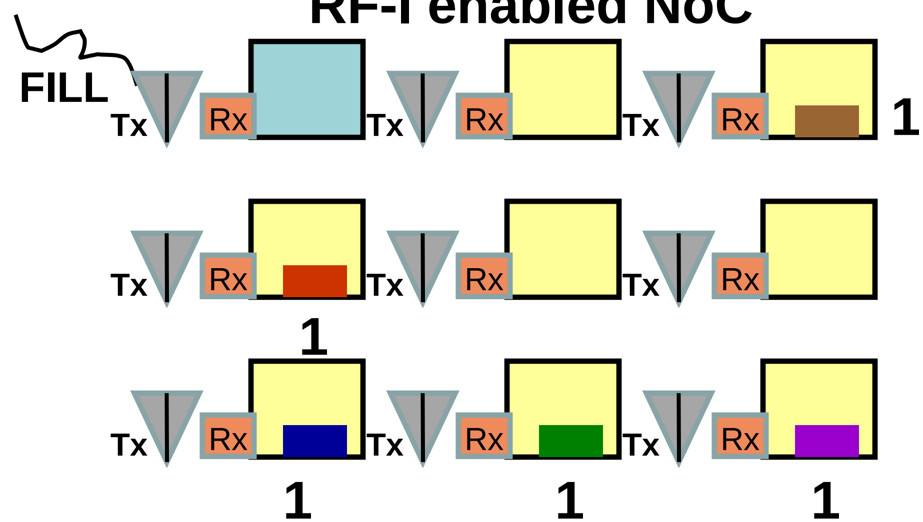

Unified Analysis

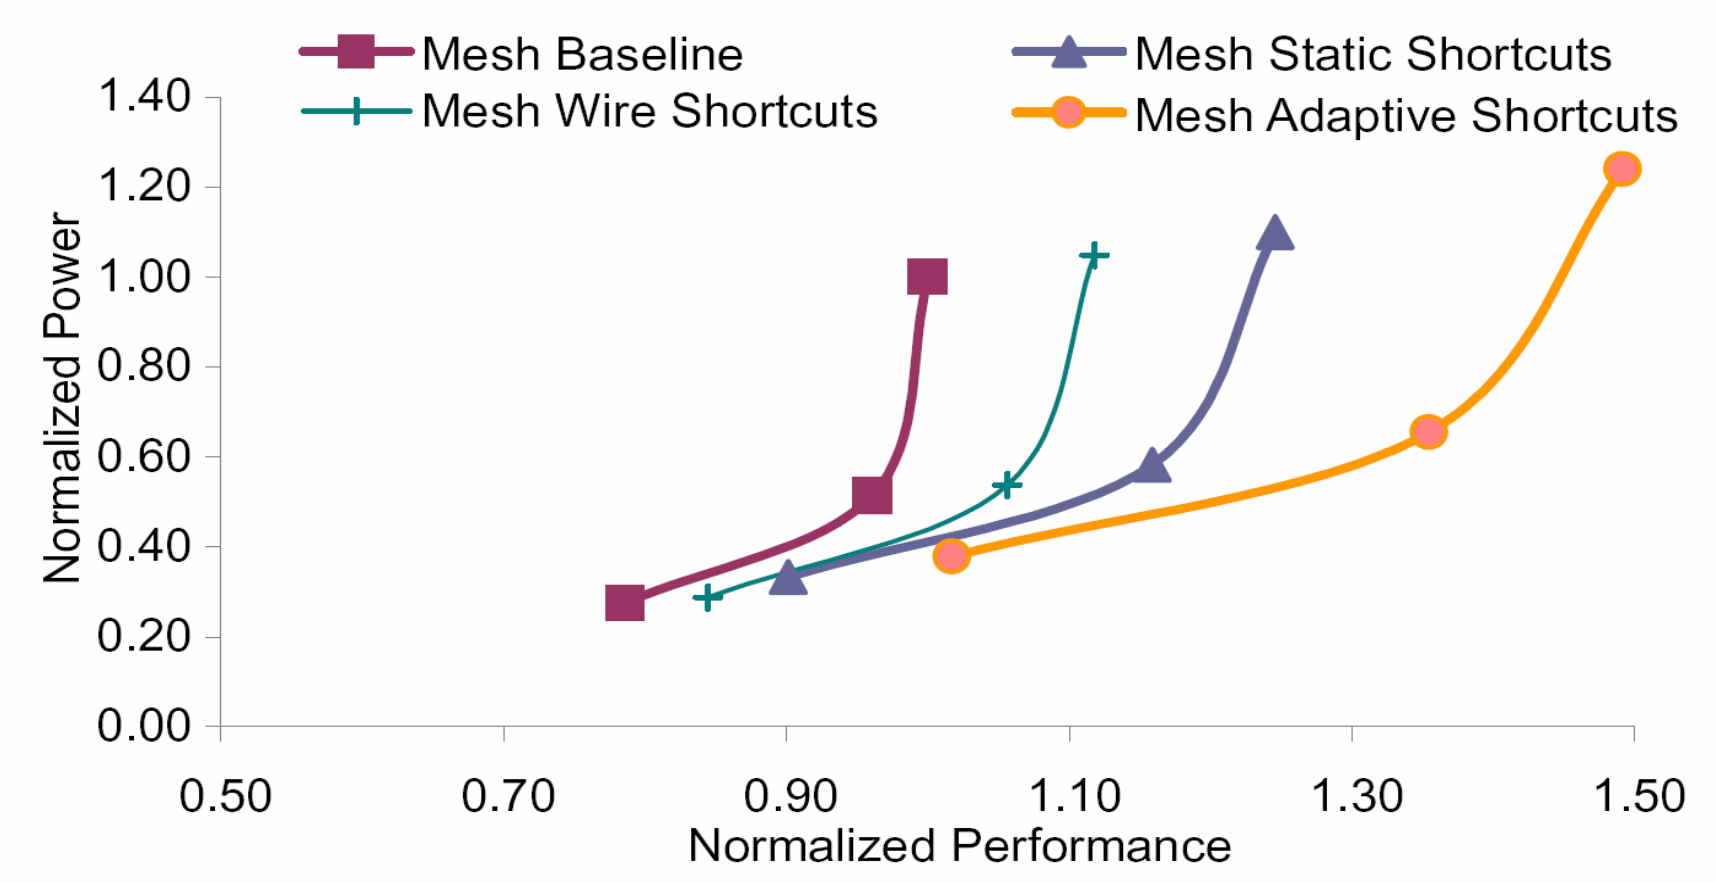

- **Adaptive RF-I enabled NoC**
  - Cost Effective in terms of both power and performance

UCLA

Acknowledgements

- DARPA and GSRC for financial
- TAPO/IBM for their foundry service