Challenges and Opportunities of Wireless Network-on-Chip for Manycore Architectures

Antonio Franques

(franque2@illinois.edu)

University of Illinois at Urbana-Champaign

NOCARC 2019 – Columbus, OH

Motivation

Transistor size ↓

Cores per chip ↑

Parallelism ↑

Communication ↑

Congestion ↑

Latencies (hops, queues) ↑

Motivation

Computations with broadcast and fine-grained data sharing do not scale well in shared-memory architectures

Master Thread

```
counter++;

barrier_wait(b)
```

Worker Threads

```
barrier_wait(b)

x = counter;
```

Manycore with a Network-on-Chip

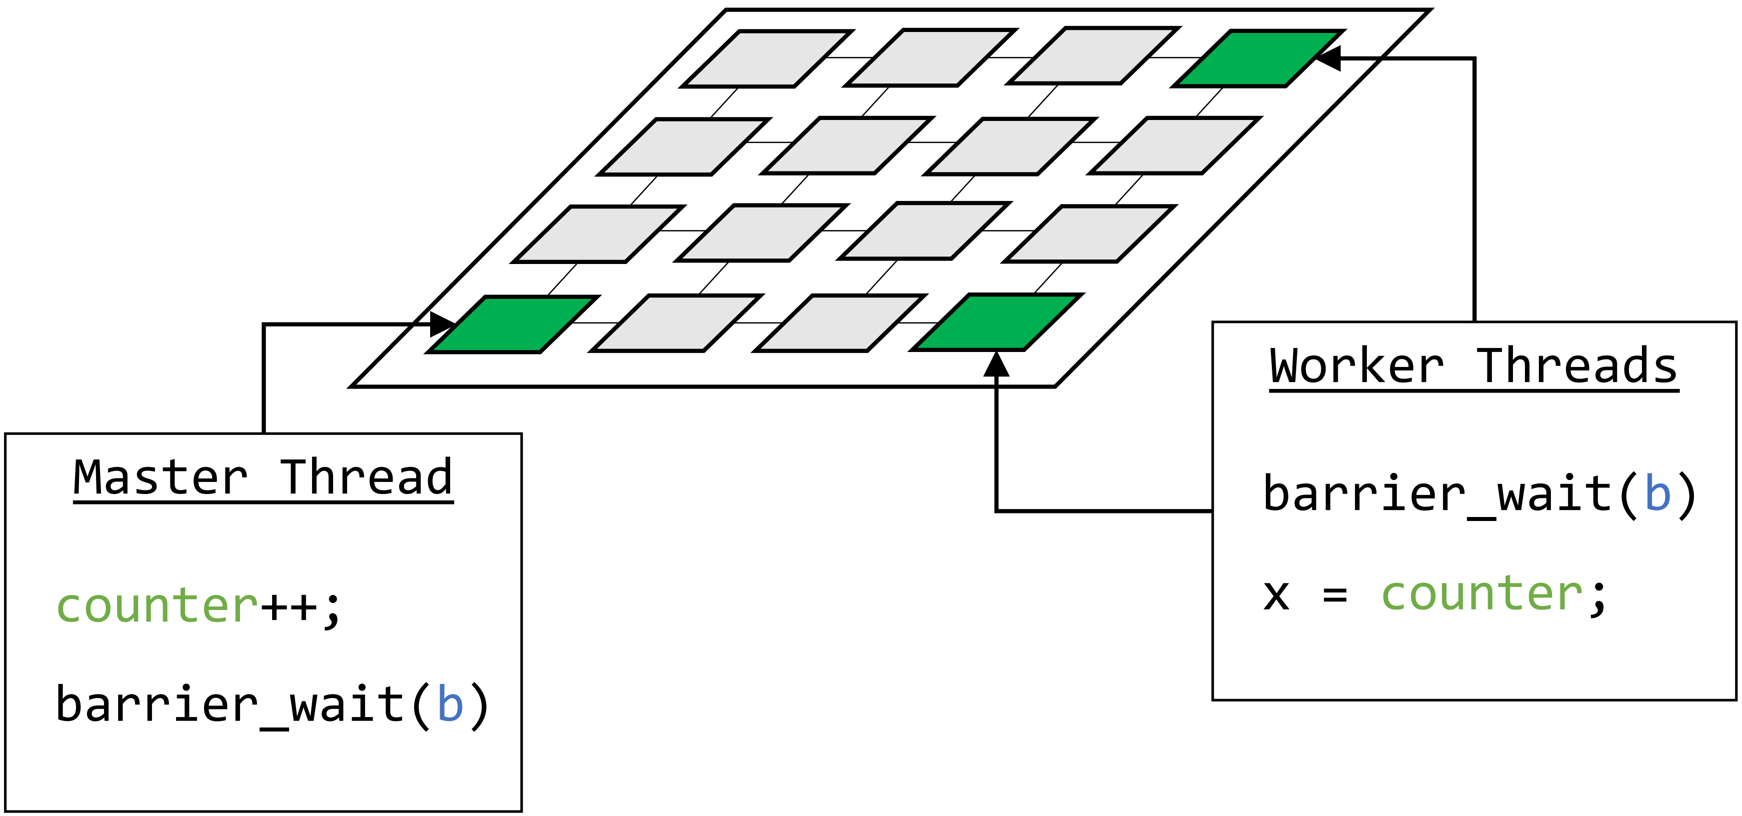

Manycore with a Network-on-Chip

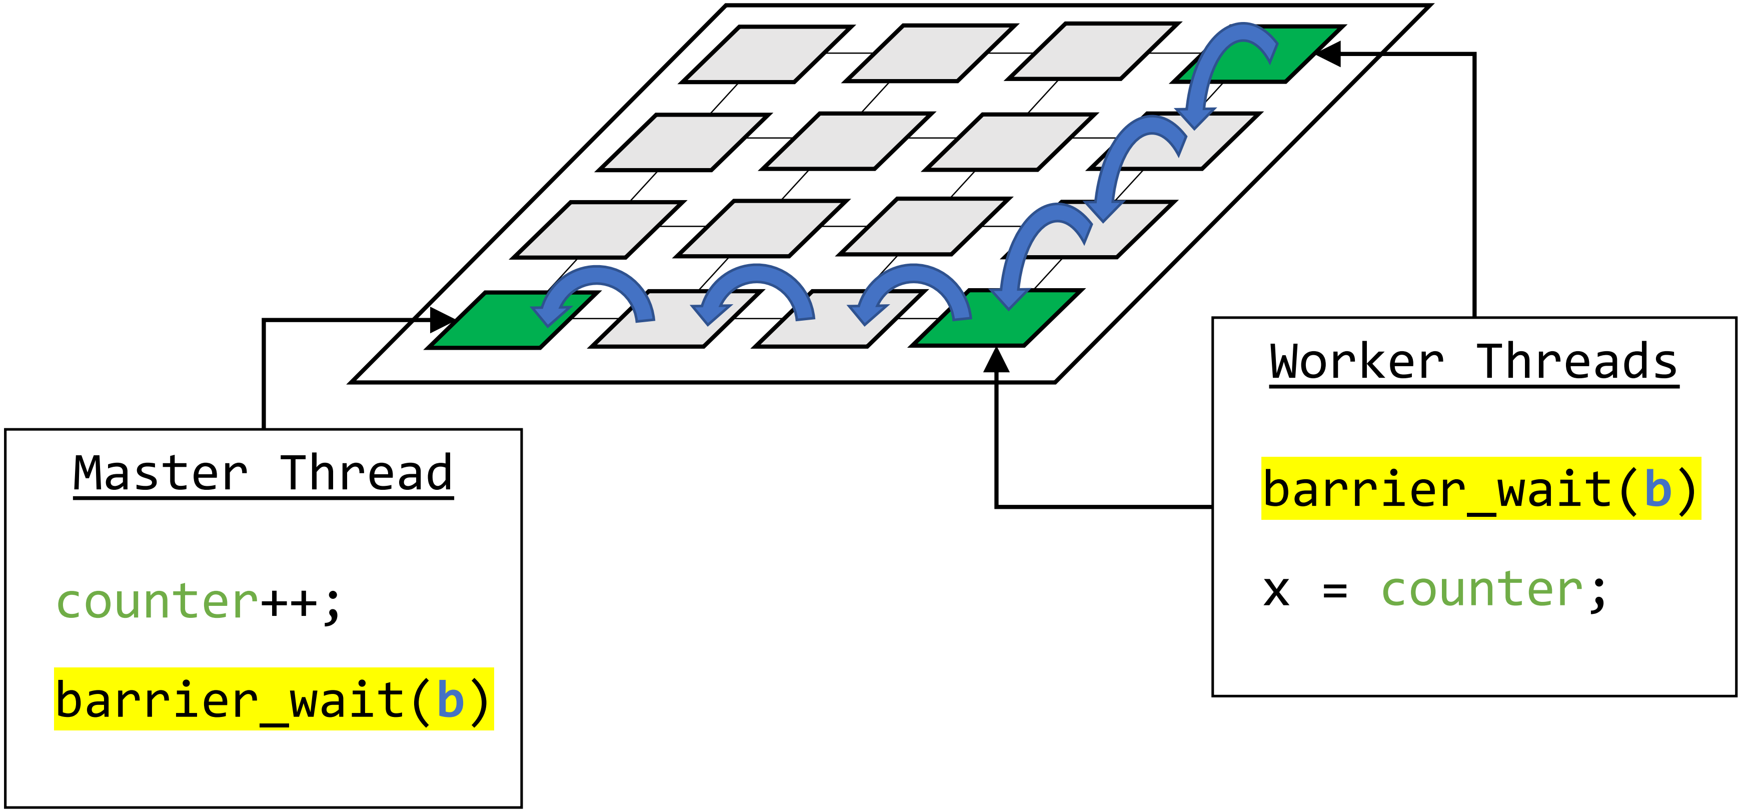

Manycore with a Network-on-Chip

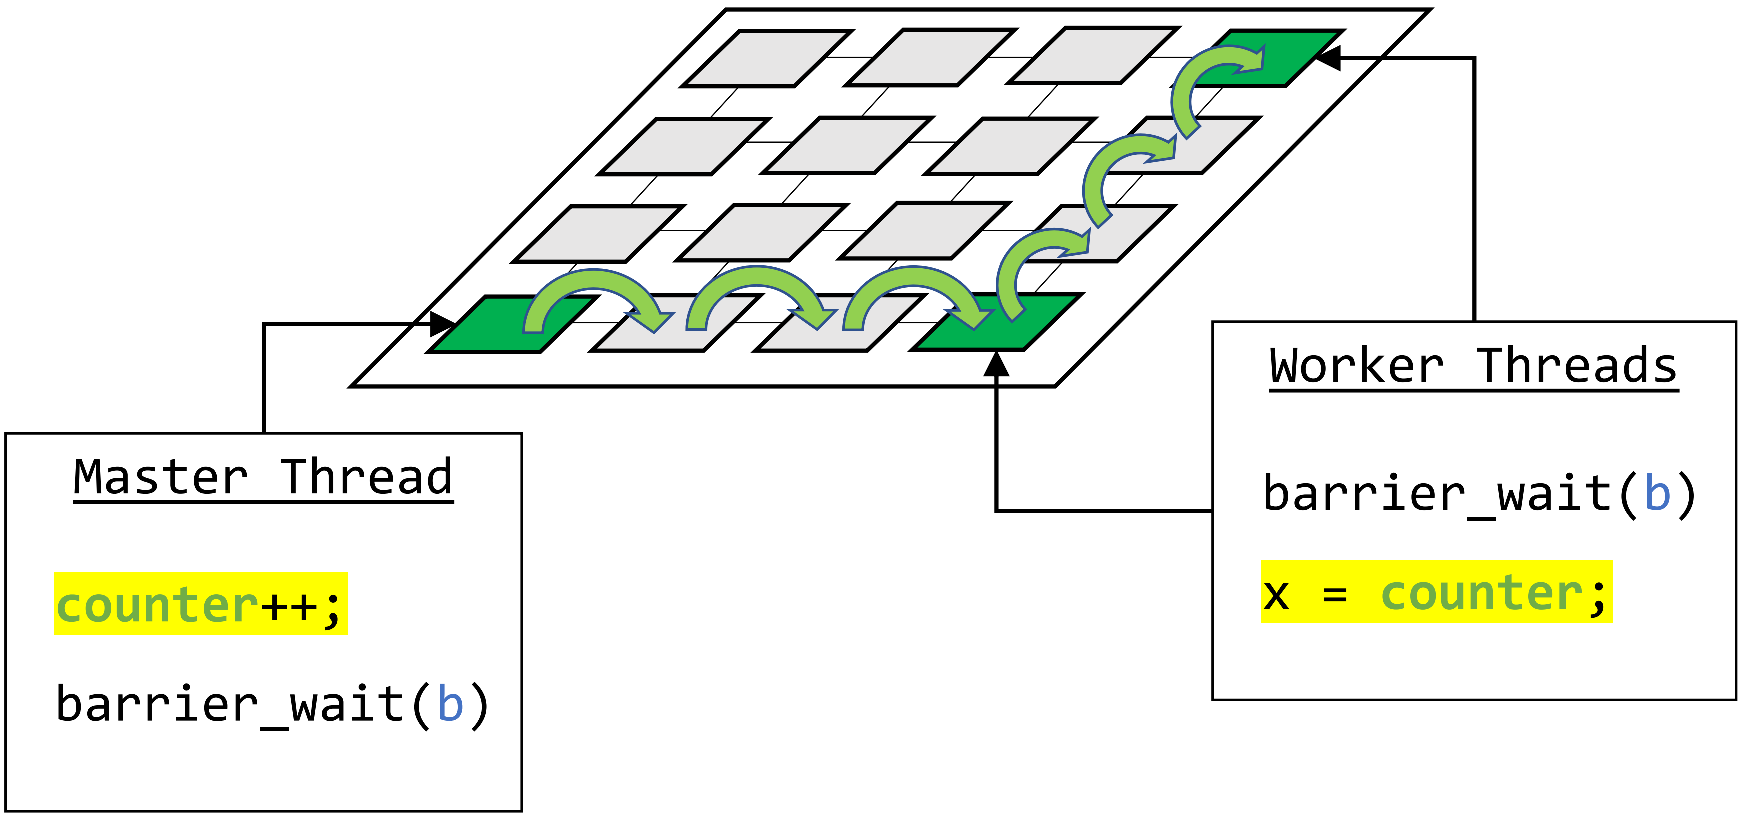

Manycore with a Network-on-Chip

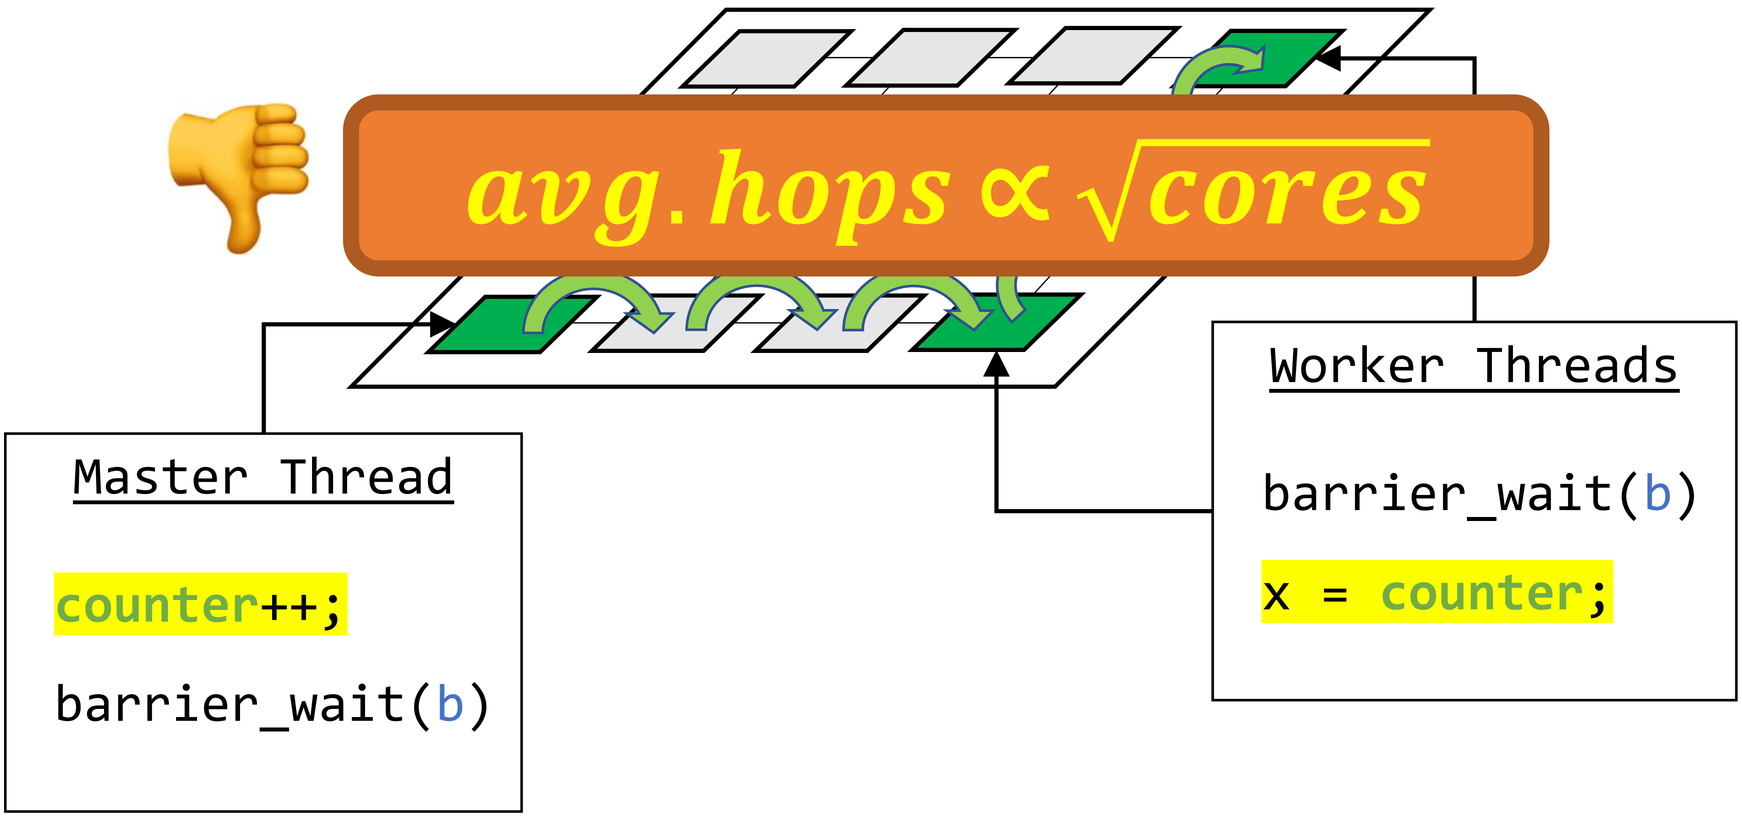

Manycore with a **Wireless** Network-on-Chip

- Low latency (speed of light → ≈ 0.1 *ns* to cross 3×3 *cm* chip)
- Inherently broadcast (can update all cores at once)
- Flexible (e.g. different apps → different channels)

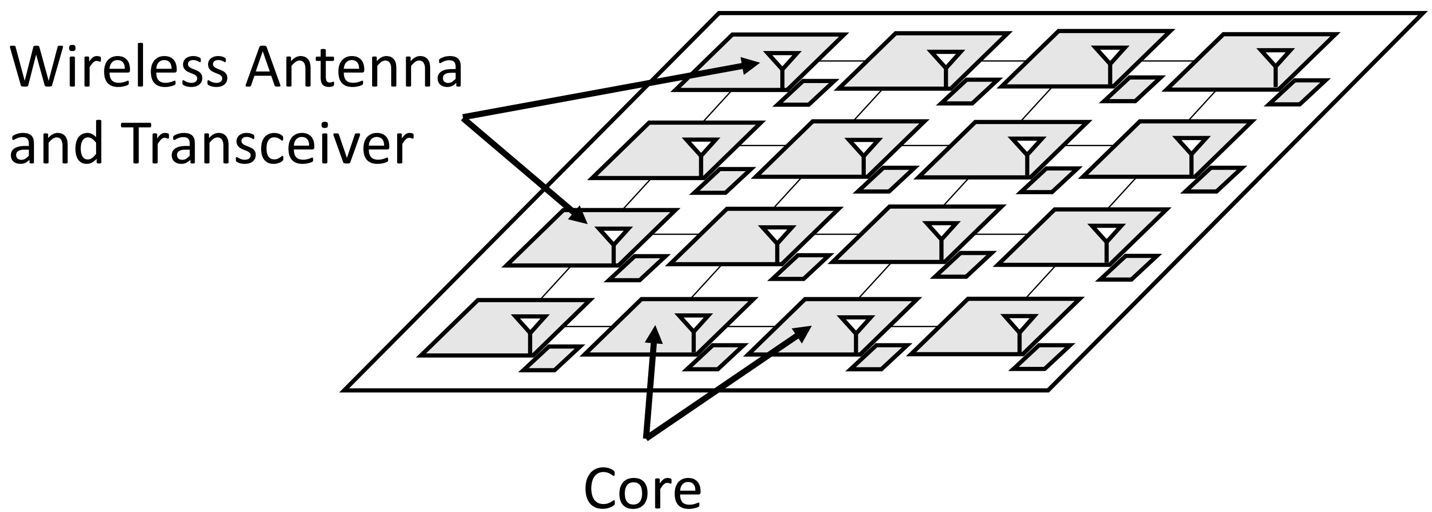

[IEEE MICRO '15] *"Broadcast-Enabled Massive Multicore Architectures: A Wireless RF Approach", Abadal et al.*

Replica: On-chip wireless communication for latency critical and highly shared data

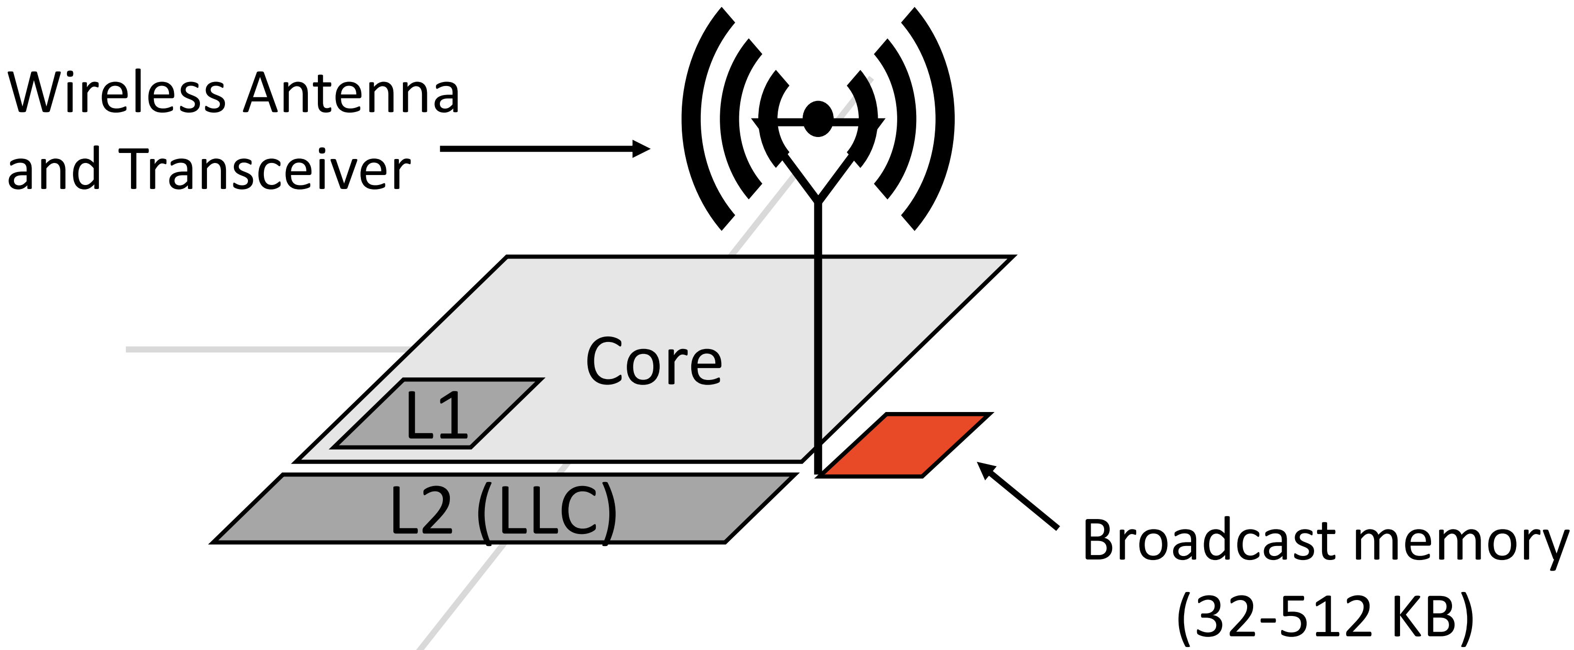

[ASPLOS ‘19] *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.“, Fernando et al.*

Replica: On-chip wireless communication for latency critical and highly shared data

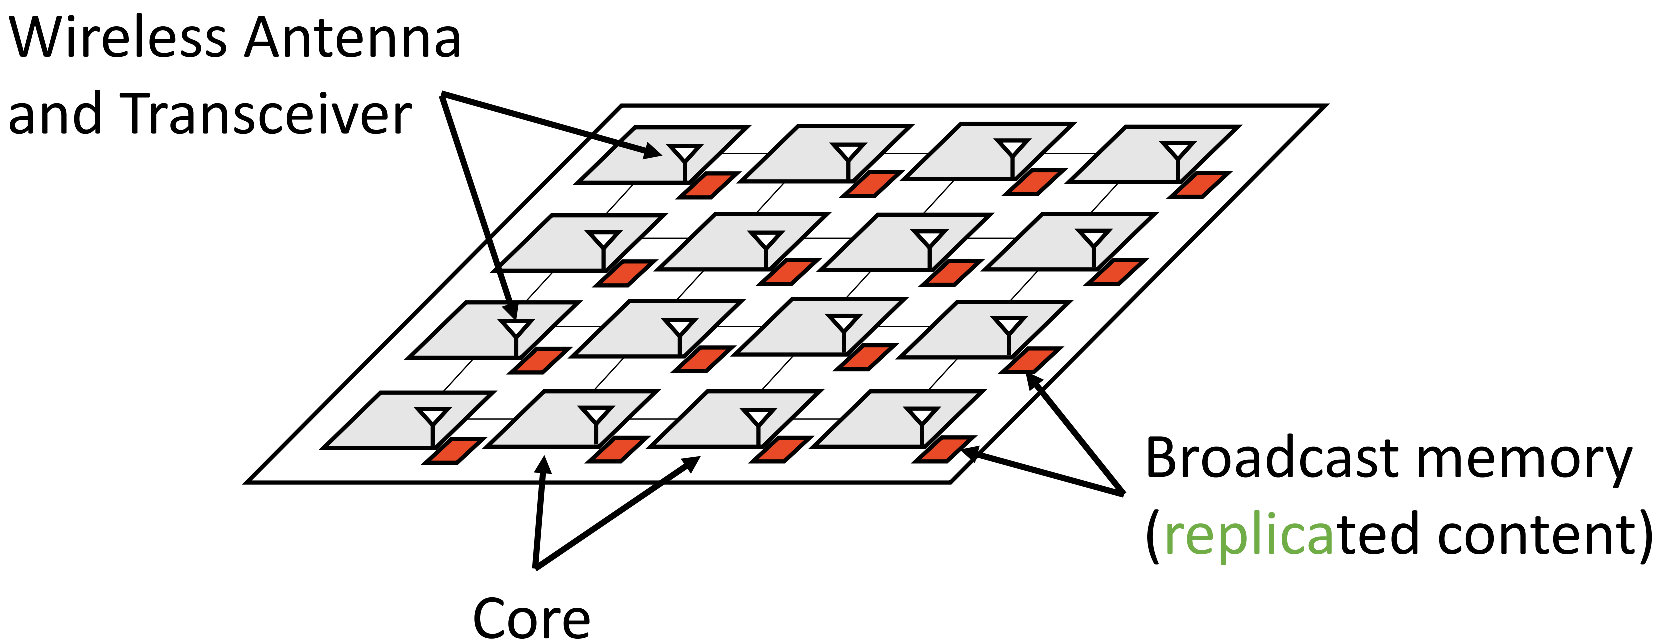

[ASPLOS '19] *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.", Fernando et al.*

Replica: Architecture

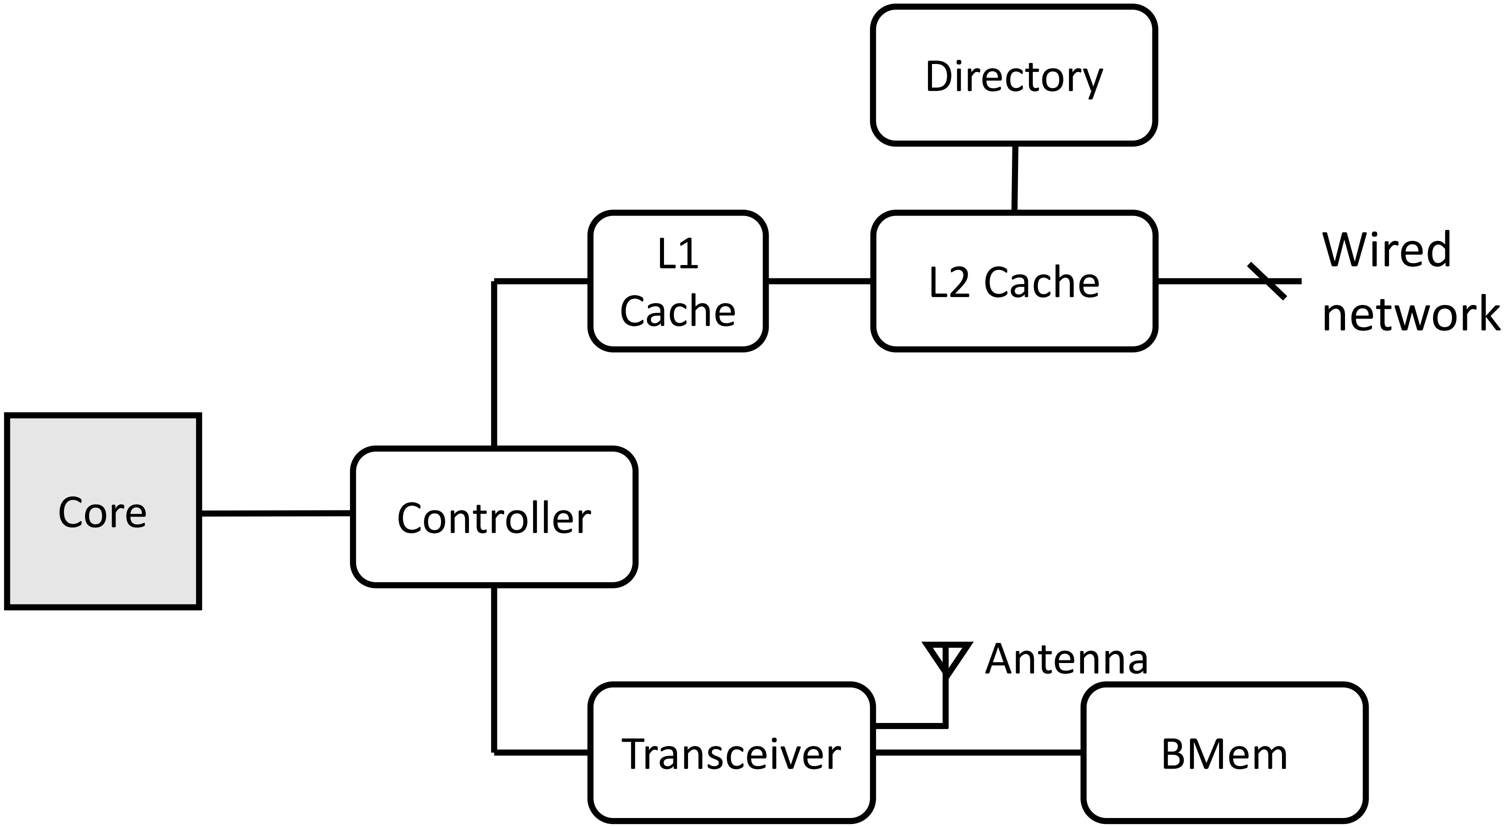

Replica: Example

```
int* counter = (int*) wireless_malloc(size)
```

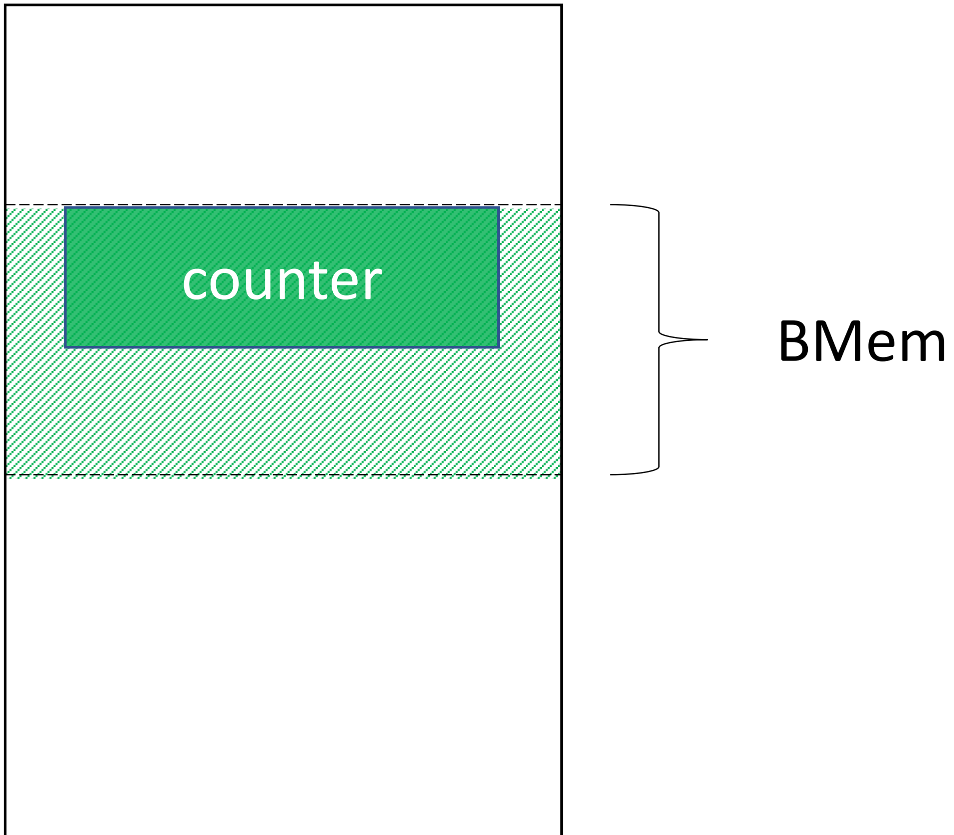

Replica: Write

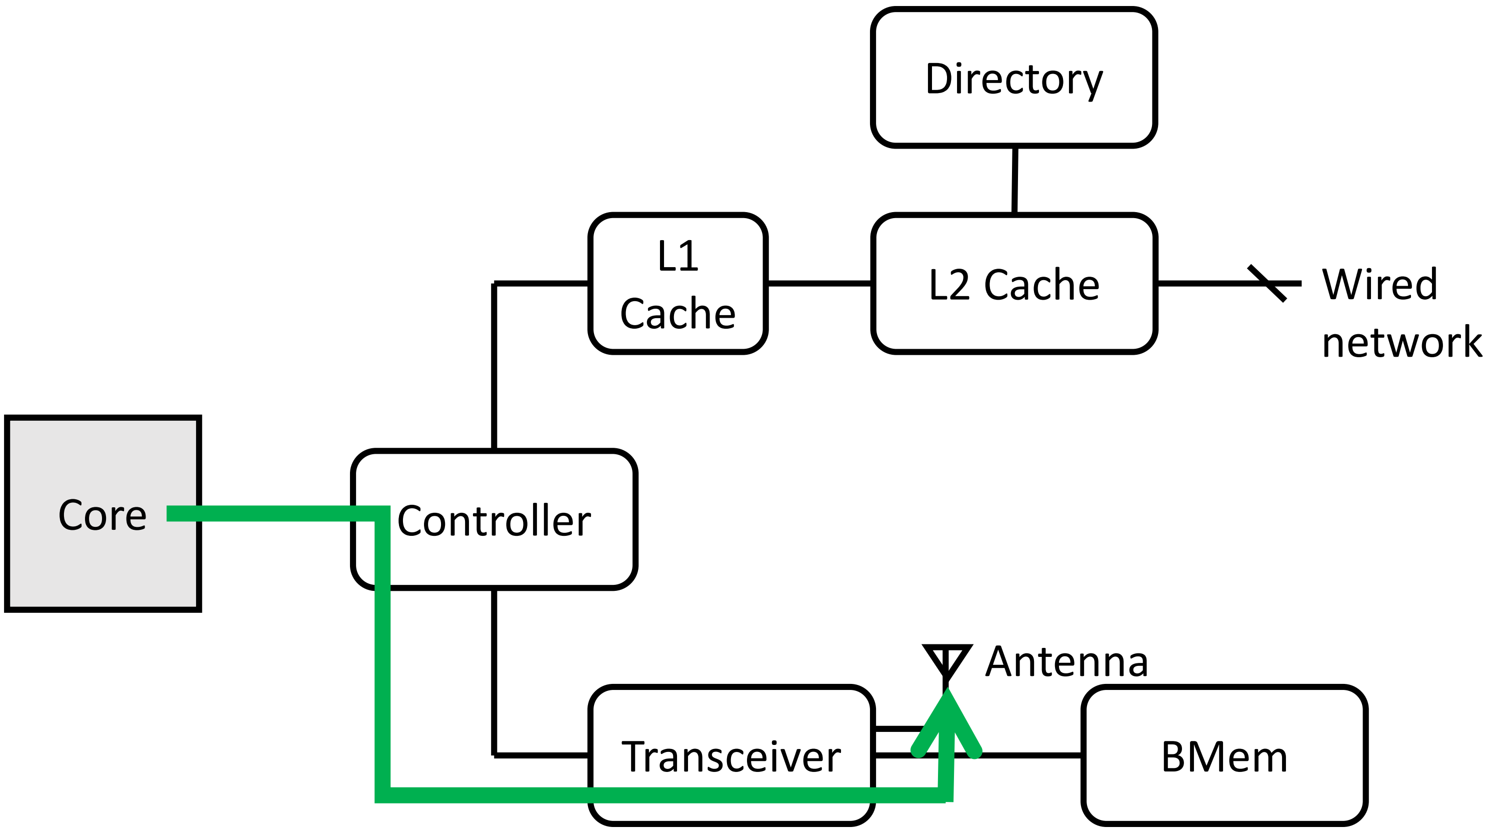

Replica: Write

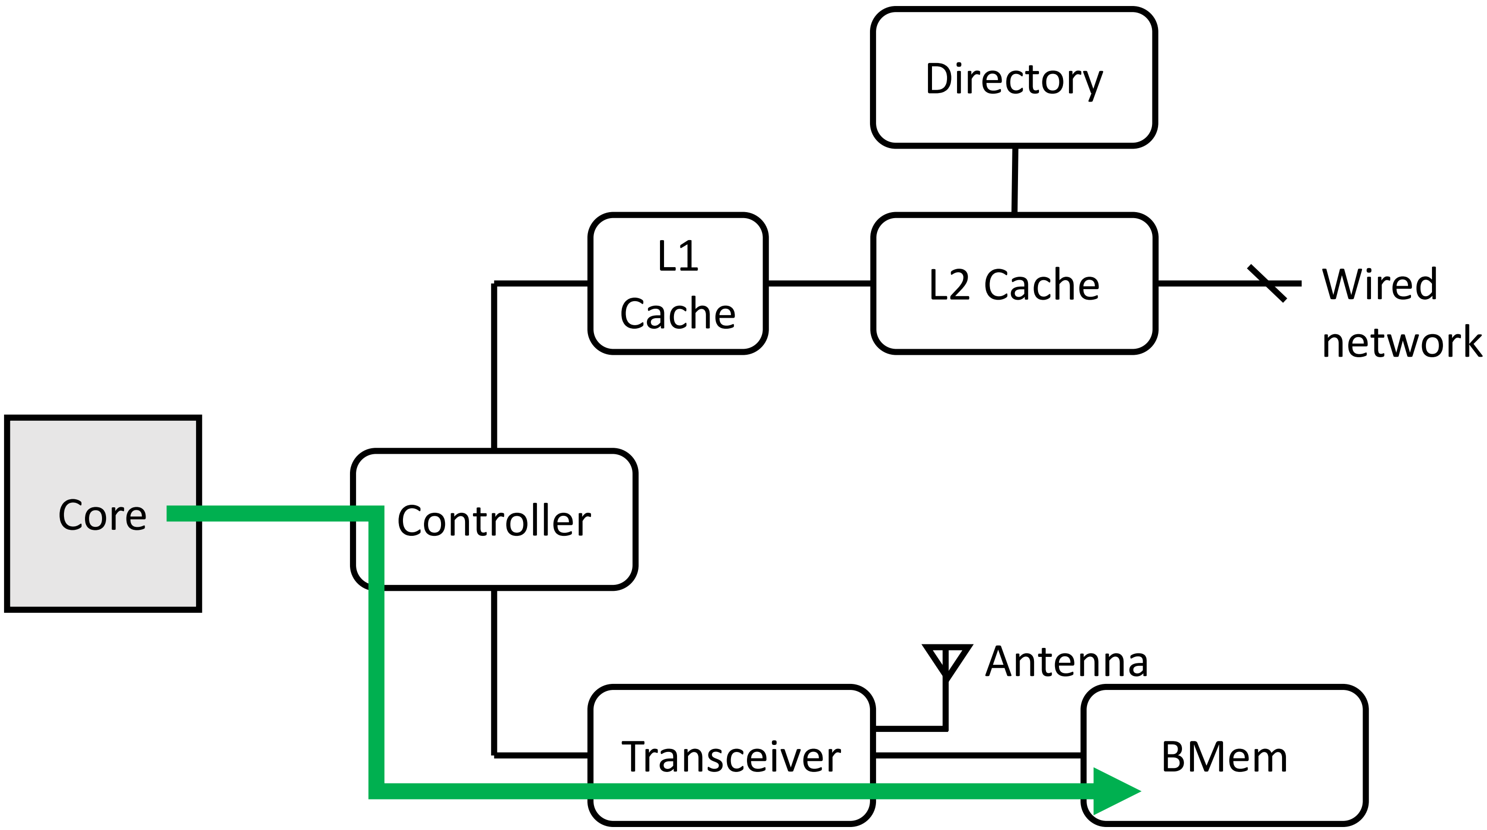

Atomic update of local and all remote BMems

Replica: Write - Example

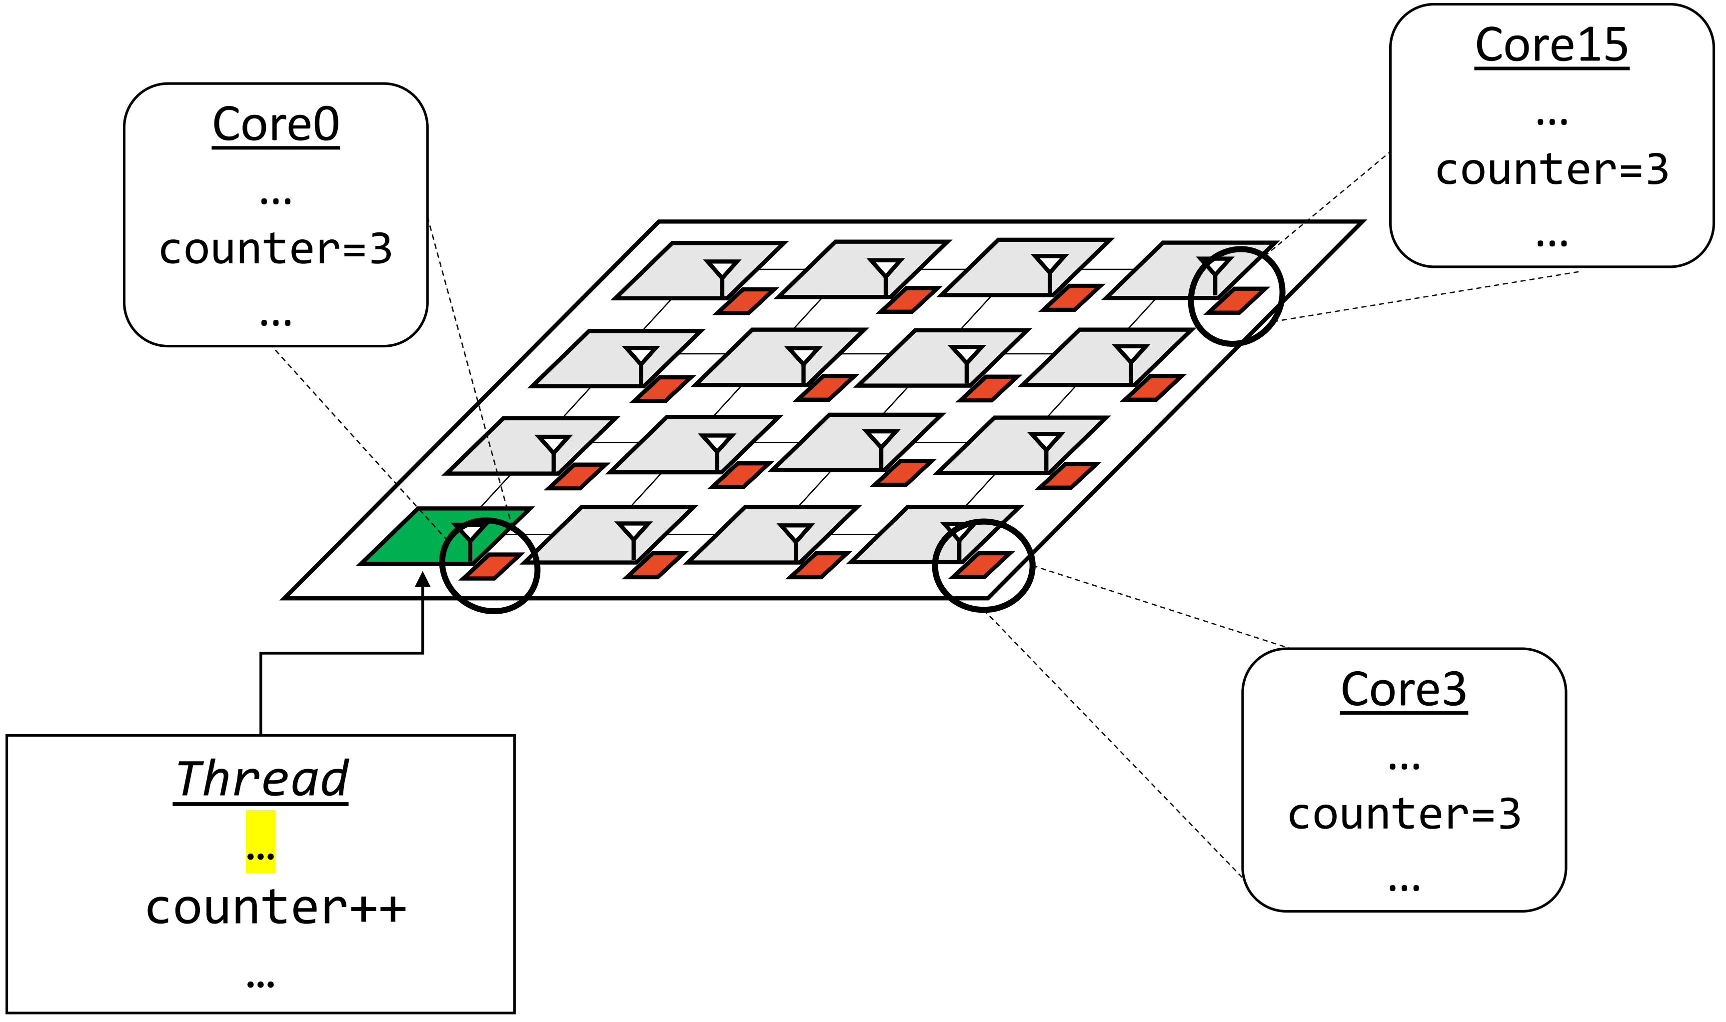

Replica: Write - Example

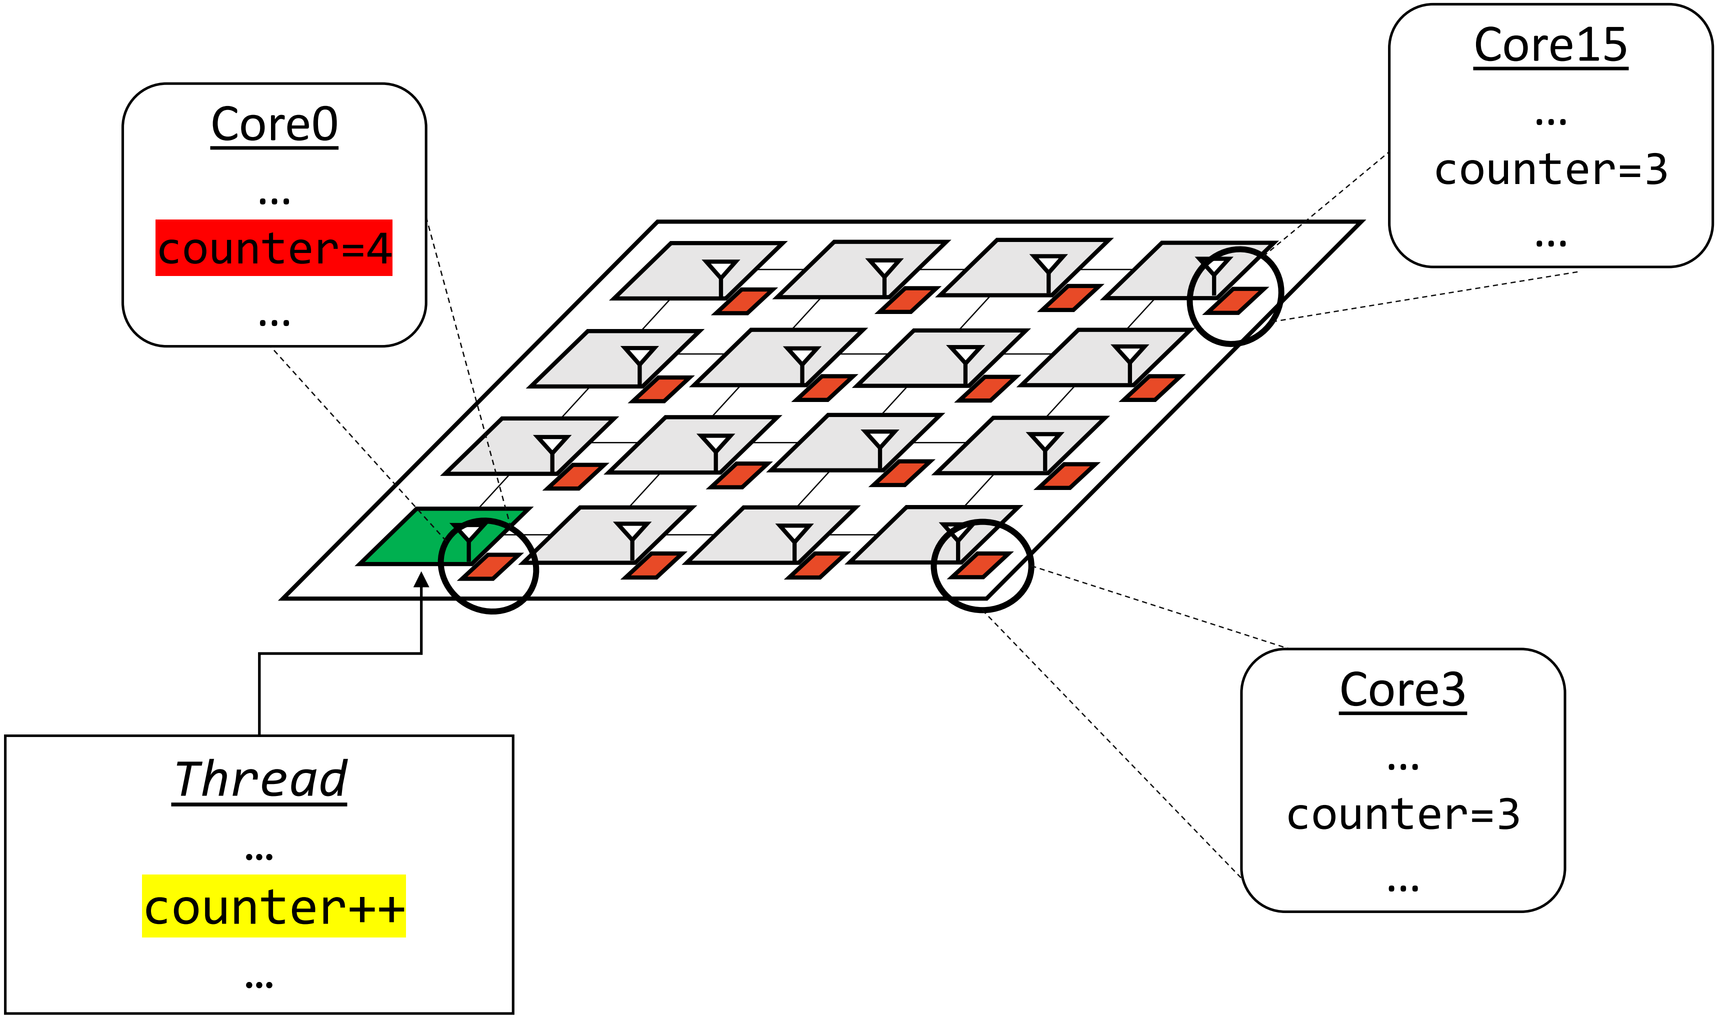

Replica: Write - Example

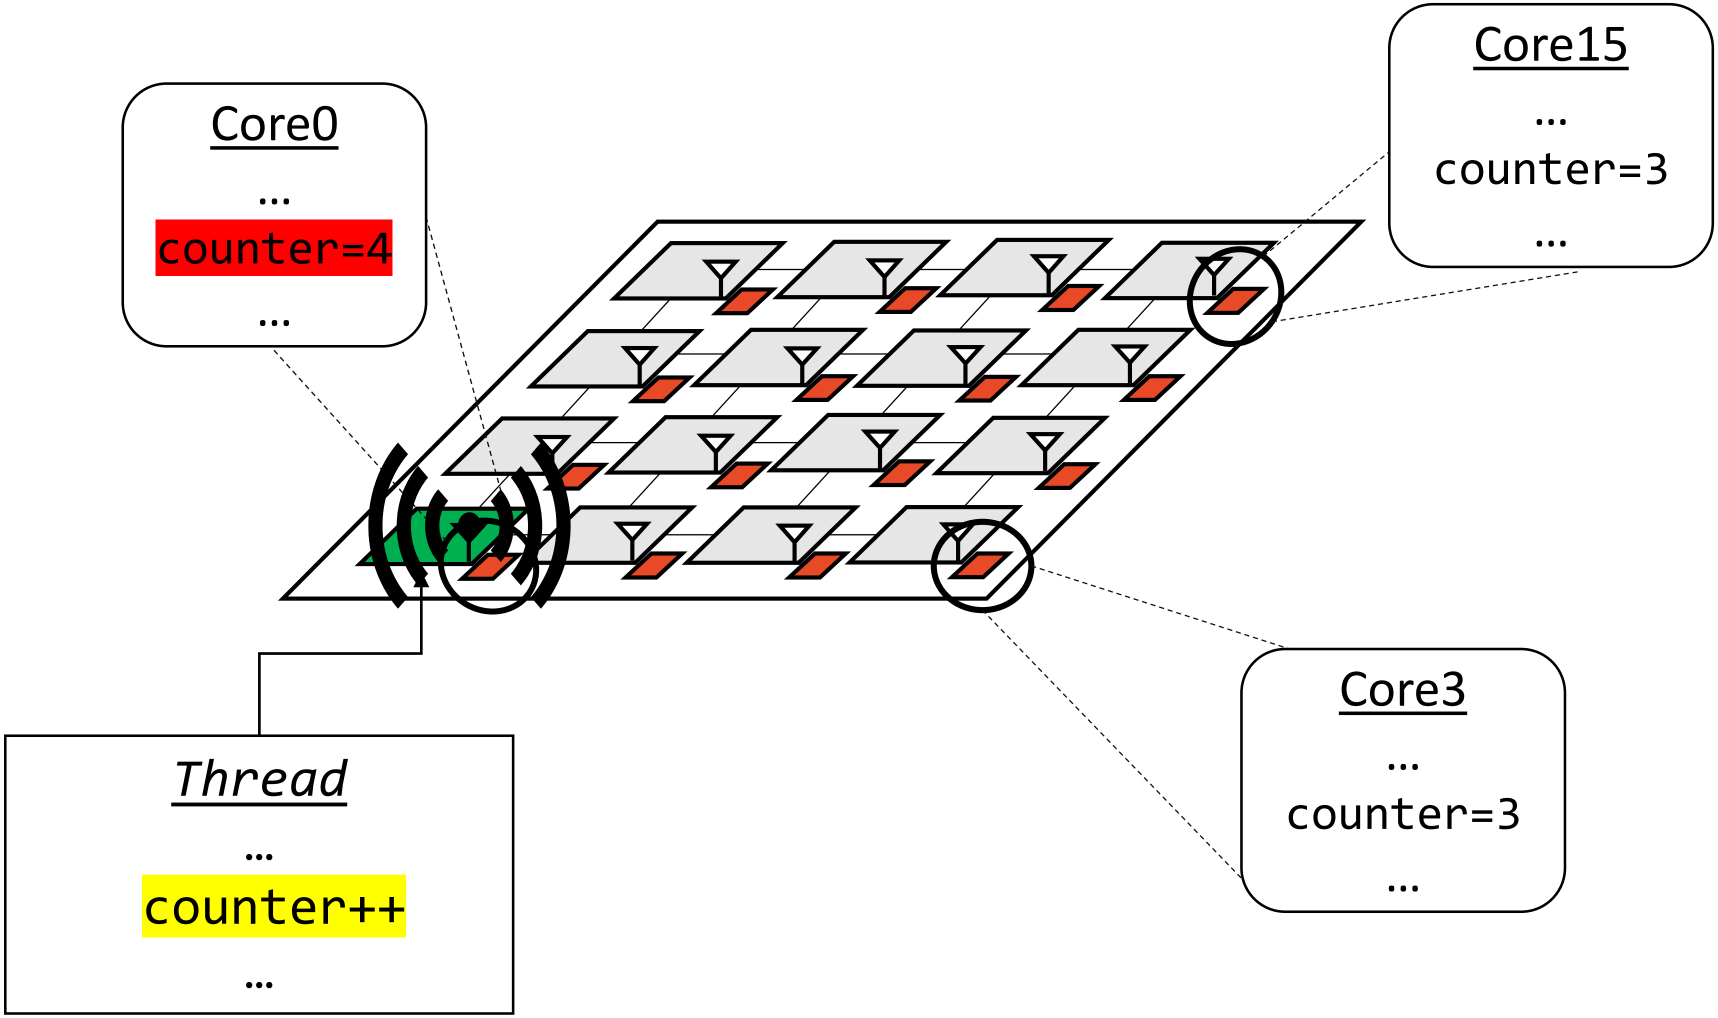

Replica: Write - Example

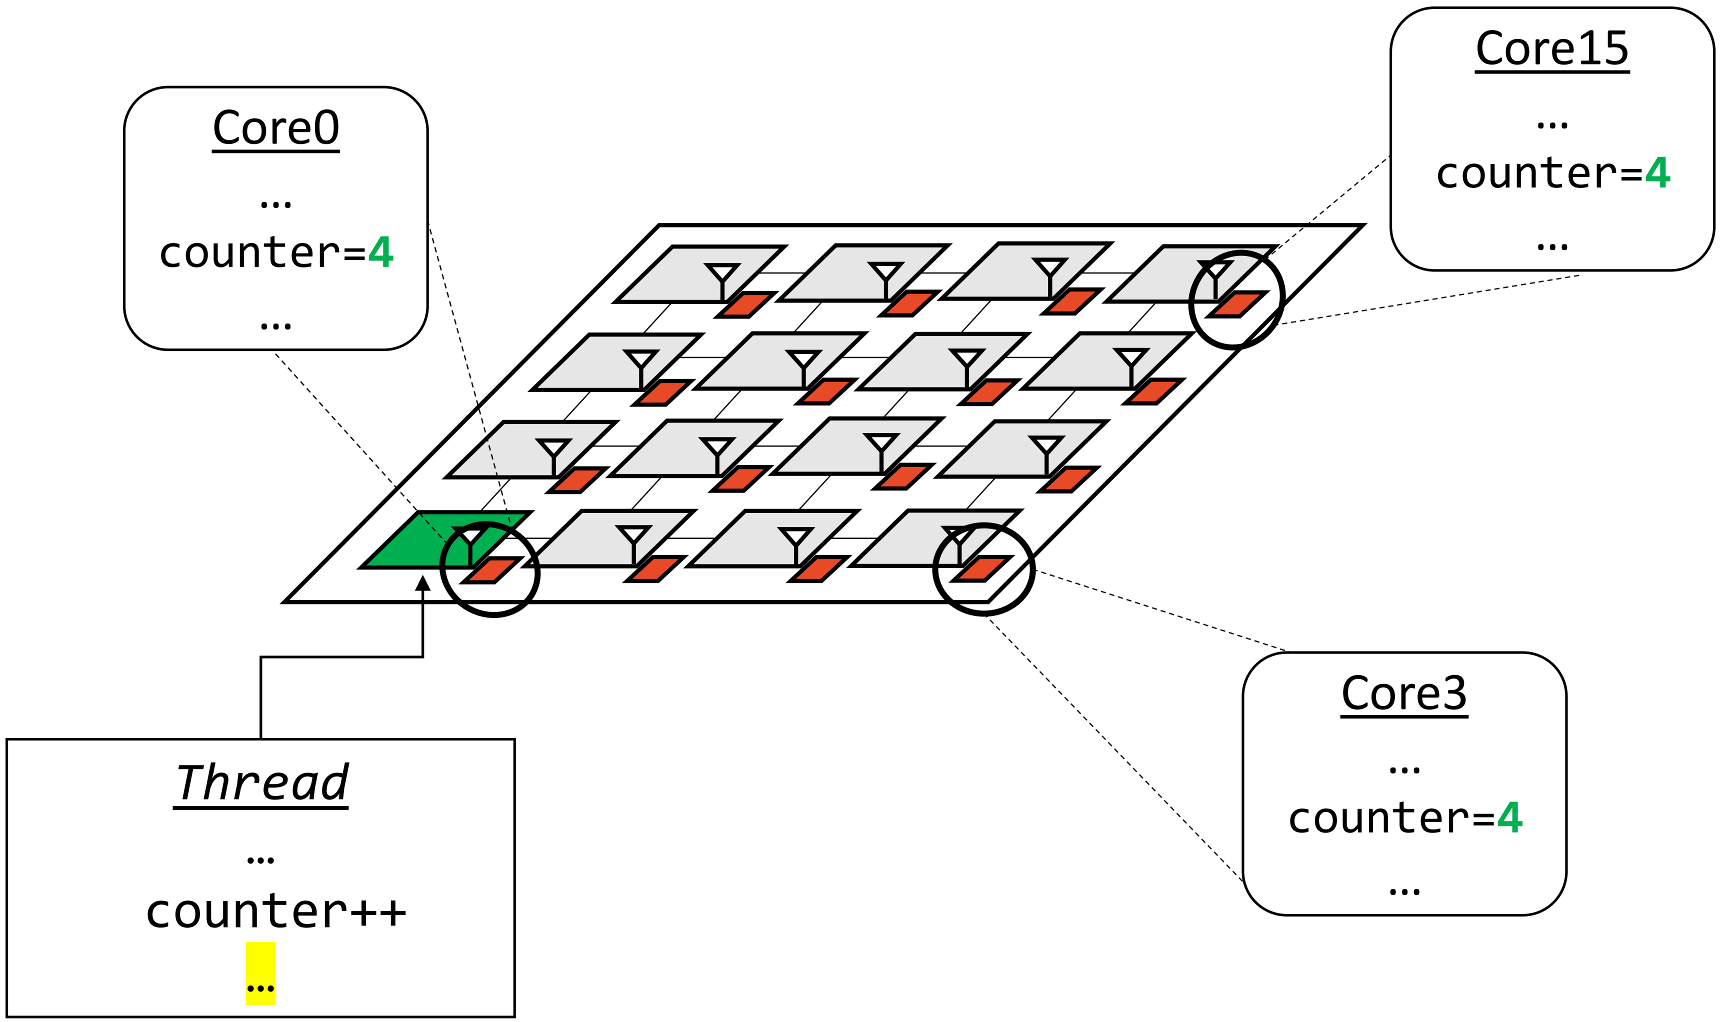

Replica: Reads

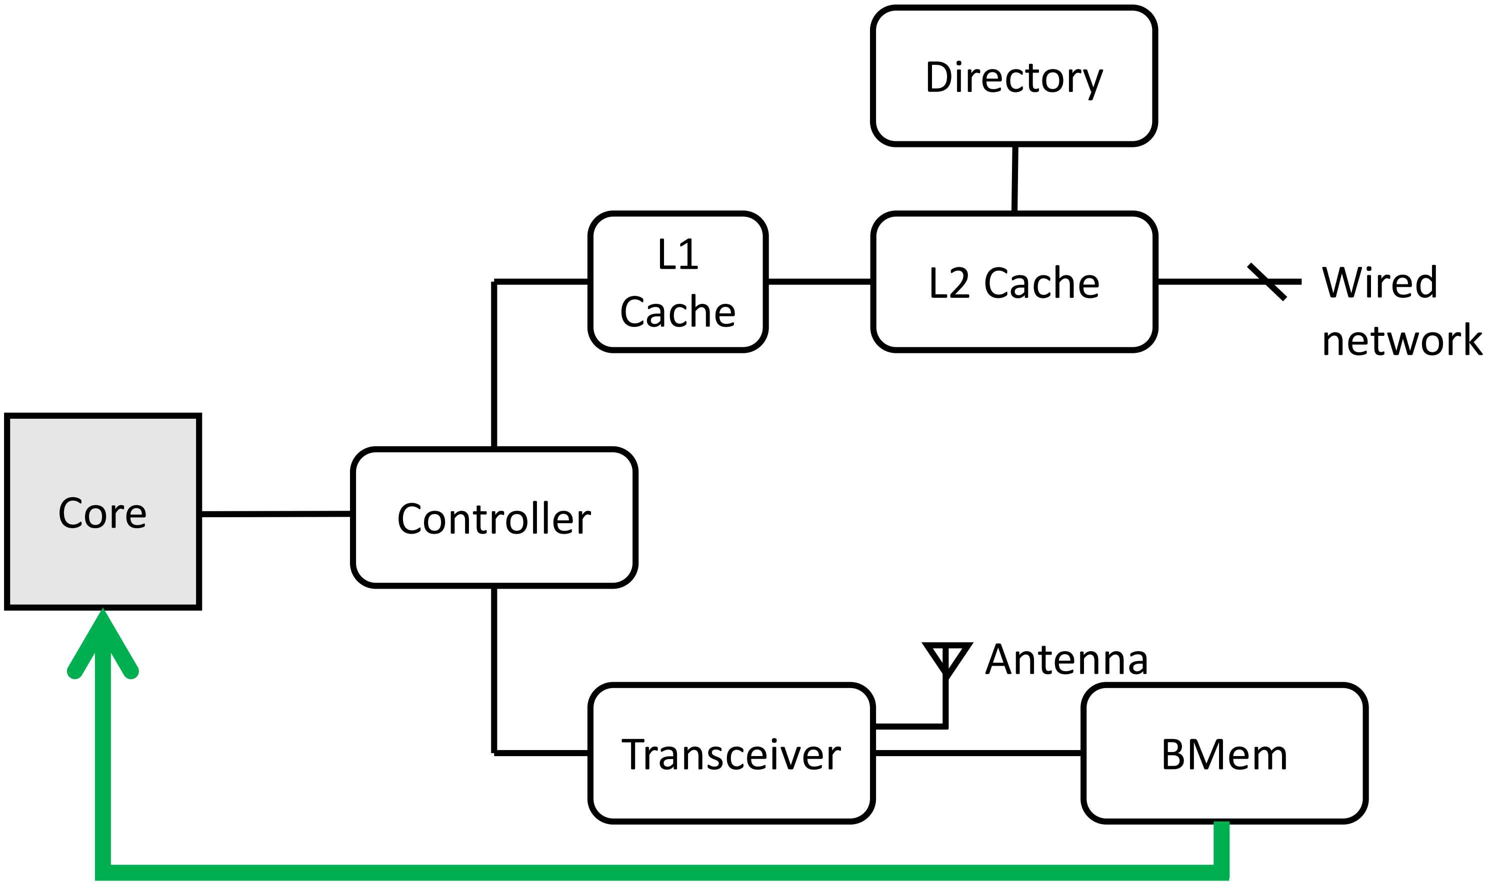

Read: Local access (6-cycle RT)

Replica: Wireless channel

- One channel, shared by all cores
- Everyone receives what one core transmits
- Only one core can transmit at a given time
  - Ensures the same order of updates across all BMems

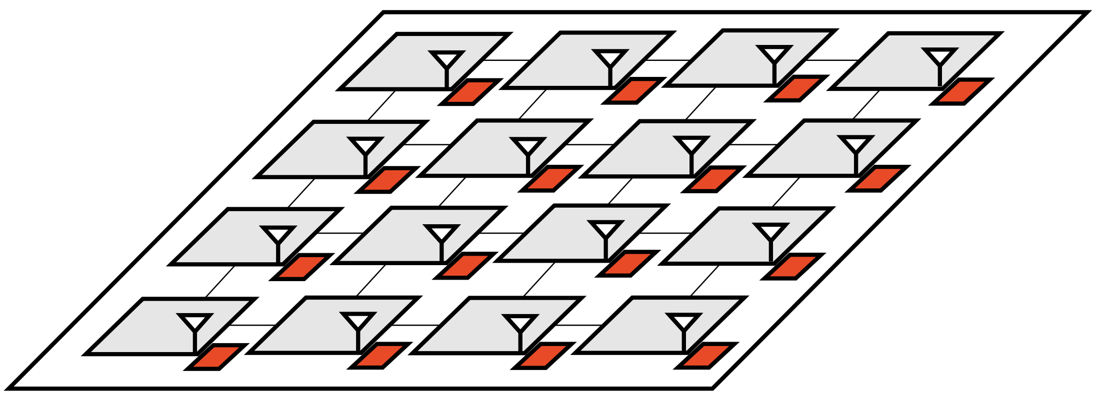

Layers of a Wireless NoC architecture

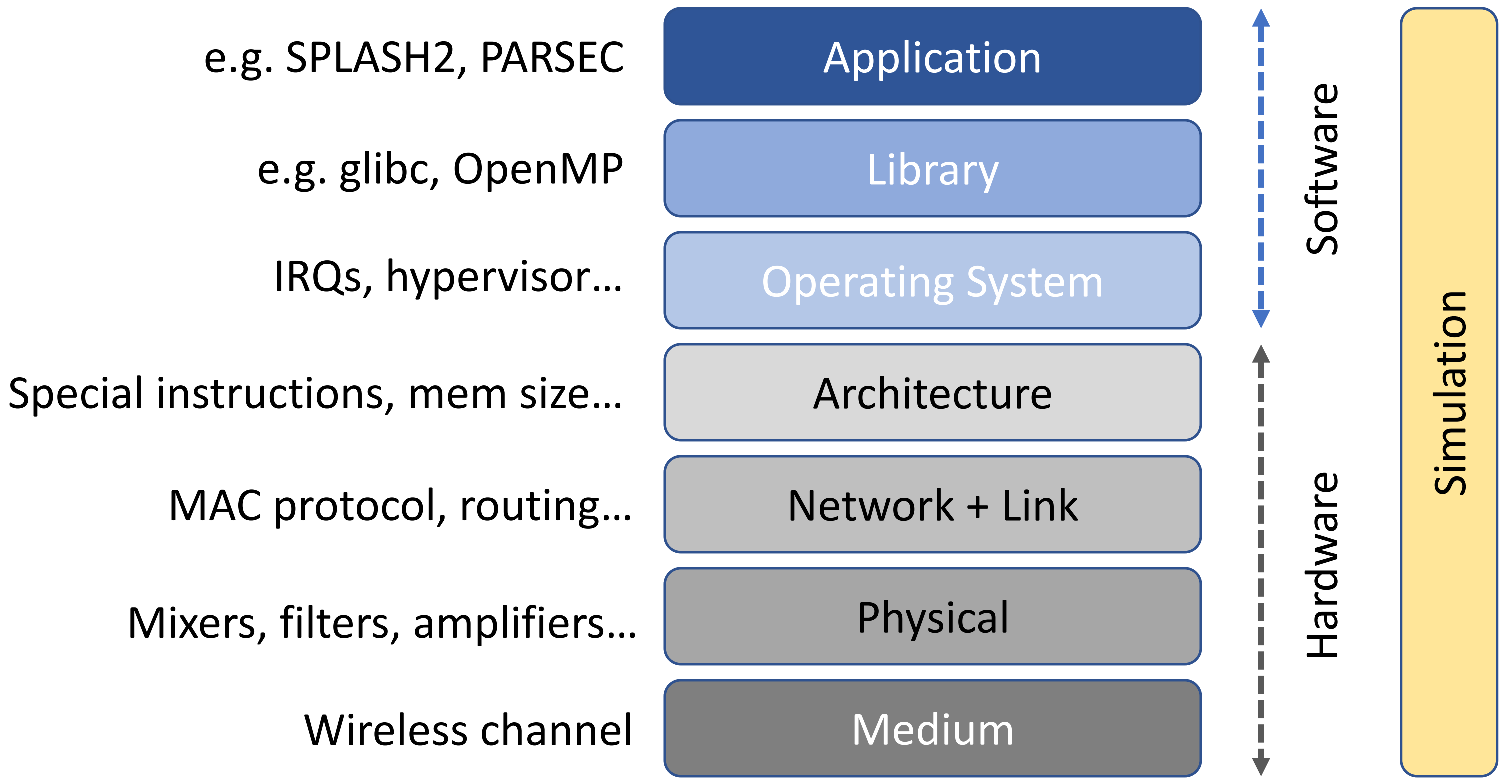

Challenges: Application

- Limited wireless bandwidth:
  - Greedy usage of wireless → poor overall performance
  - How do we split data between wired and wireless NoC?
- Traditional apps written to minimize data sharing
  - Historically, multicast/broadcast communication has been expensive → architecture designed to hide latency
  - Hard to find off-the-shelf algorithms that exploit unique properties of wireless

Opportunities: Application

- Tools to identify/autotune highly-shared data
- Approx. resilient apps → Selective message dropping
  - BER (Bit Error Rate) $\propto \frac{1}{\text{Power tx.}}$
  - Can also apply to baseline, but has higher impact on wireless

```
if (update < Threshold) {
      approx_store(x, update);
} else {
      x += update;
}
```

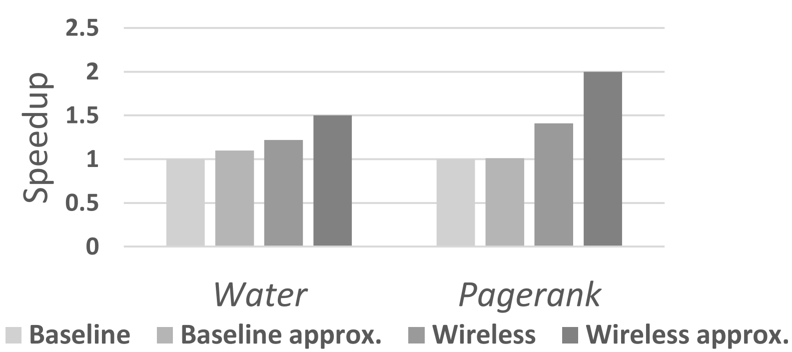

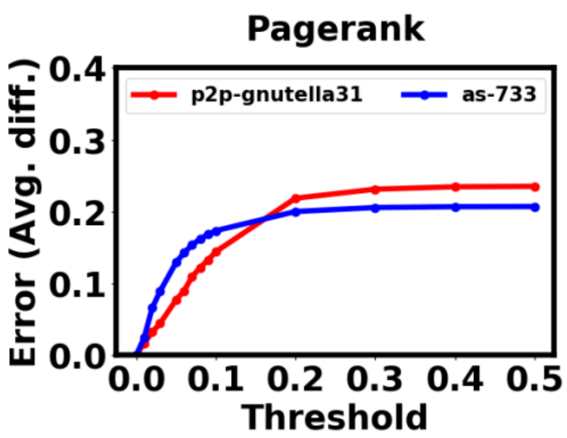

[ASPLOS '19] *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.", Fernando et al.*

Opportunities: Application

- Multithreaded benchmarks w/ data sharing
  - SPLASH-2, PARSEC, CRONO… (scientific simulations, computer vision, graph applications…)

Benchmark	Sharing Pattern	Approximations
Water	Broadcast	Precision reduction
BFS	Irregular: many-to-many	Approximate Stores
Bodytrack	One-to-many	Approximate Stores
SSSP	Irregular: many-to-many	Approximate Stores
CC	Irregular: many-to-many	Approximate Stores
Streamcluster	One-to-many, reduction	Cyclic collection updates
Pagerank	Irregular: many-to-many	Skipping negligible updates
Community	Irregular: many-to-many	Approximate Stores
	…	

- What about machine learning algorithms?

Challenges/Opportunities: Library

- Wireless capabilities should be exposed to users
  - If too cumbersome to use → technology will not catch
  - Libraries should take the burden, not the users
- Need to rethink and rewrite implementation of popular parallel programming libraries/frameworks
  - …and the popular apps that use them too
  - E.g. *wireless_malloc**, locks/barriers, etc.

OpenMP Scala MPI Charmworks akka

***[ASPLOS '19]** *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.", Fernando et al.*

Challenges: Operating System

- Memory management might require changes
  - Separate between standard and wireless memory
- Historically, thread scheduling exploited locality to minimize communication cost between threads
  - This requirement might no longer hold (constant propagation time)

Opportunities: Operating System

- Can we relax the thread scheduling policies?
- Speed up OS synchronization?
- Broadcast interrupt requests (IRQs)
  - e.g. core wake-up

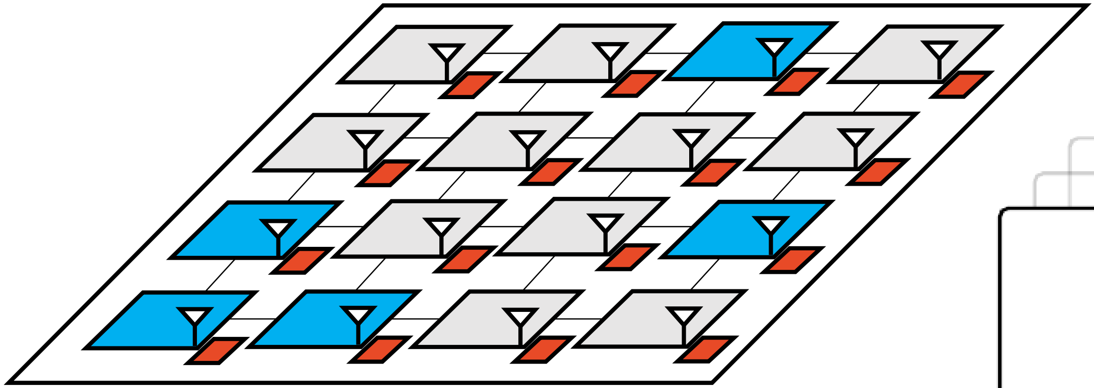

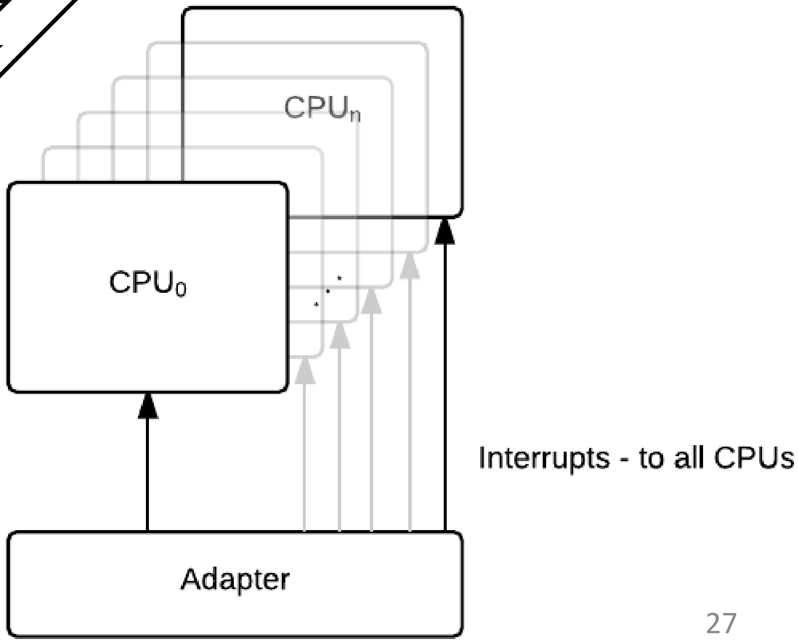

Challenges: Architecture

- Area/Energy overhead
  - Dennard scaling no longer holds → power is scarce
- Multiple frequencies/channels
  - More channels → More data parallelism
  - How to distribute data across channels?
- Separate memory for regular/wireless data?

Manycore architecture (Piton*) **Integrated RF Transceivers**

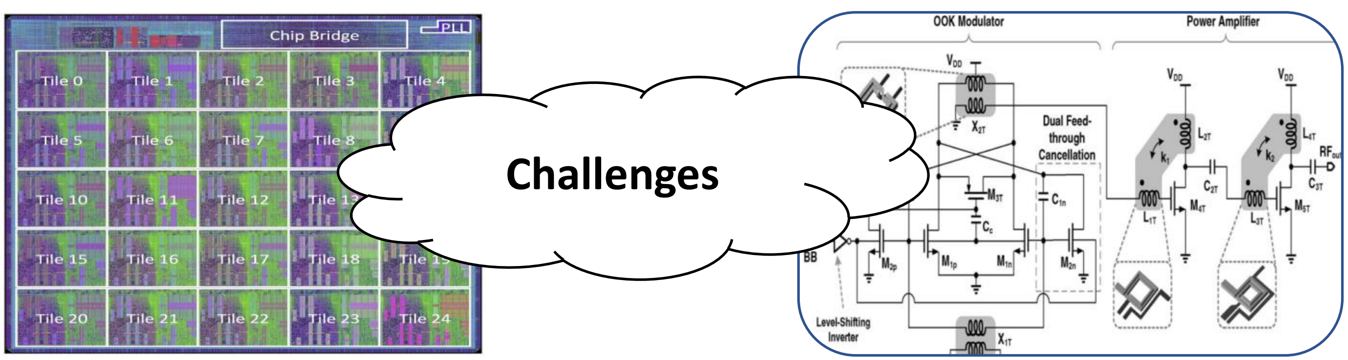

***[HotChips '16]** *"Piton: A 25-core Academic Manycore Research Processor", McKeown et al.*

Opportunities: Architecture

- Chiplets paradigm
  - Die-to-die wireless communication?
- Instruction Set Architecture (ISA)
  - Extensions for wireless-specific operations?
- Scalability with number of cores
  - Clustering of antennas?
  - Broadcast?
  - Point-to-point + routing?
- Layout sensitivity analysis
  - Closer to core → lower latency, smaller transc. size
  - Farther from core → higher latency, bigger transc. size

Wisync: BMem for synchronization

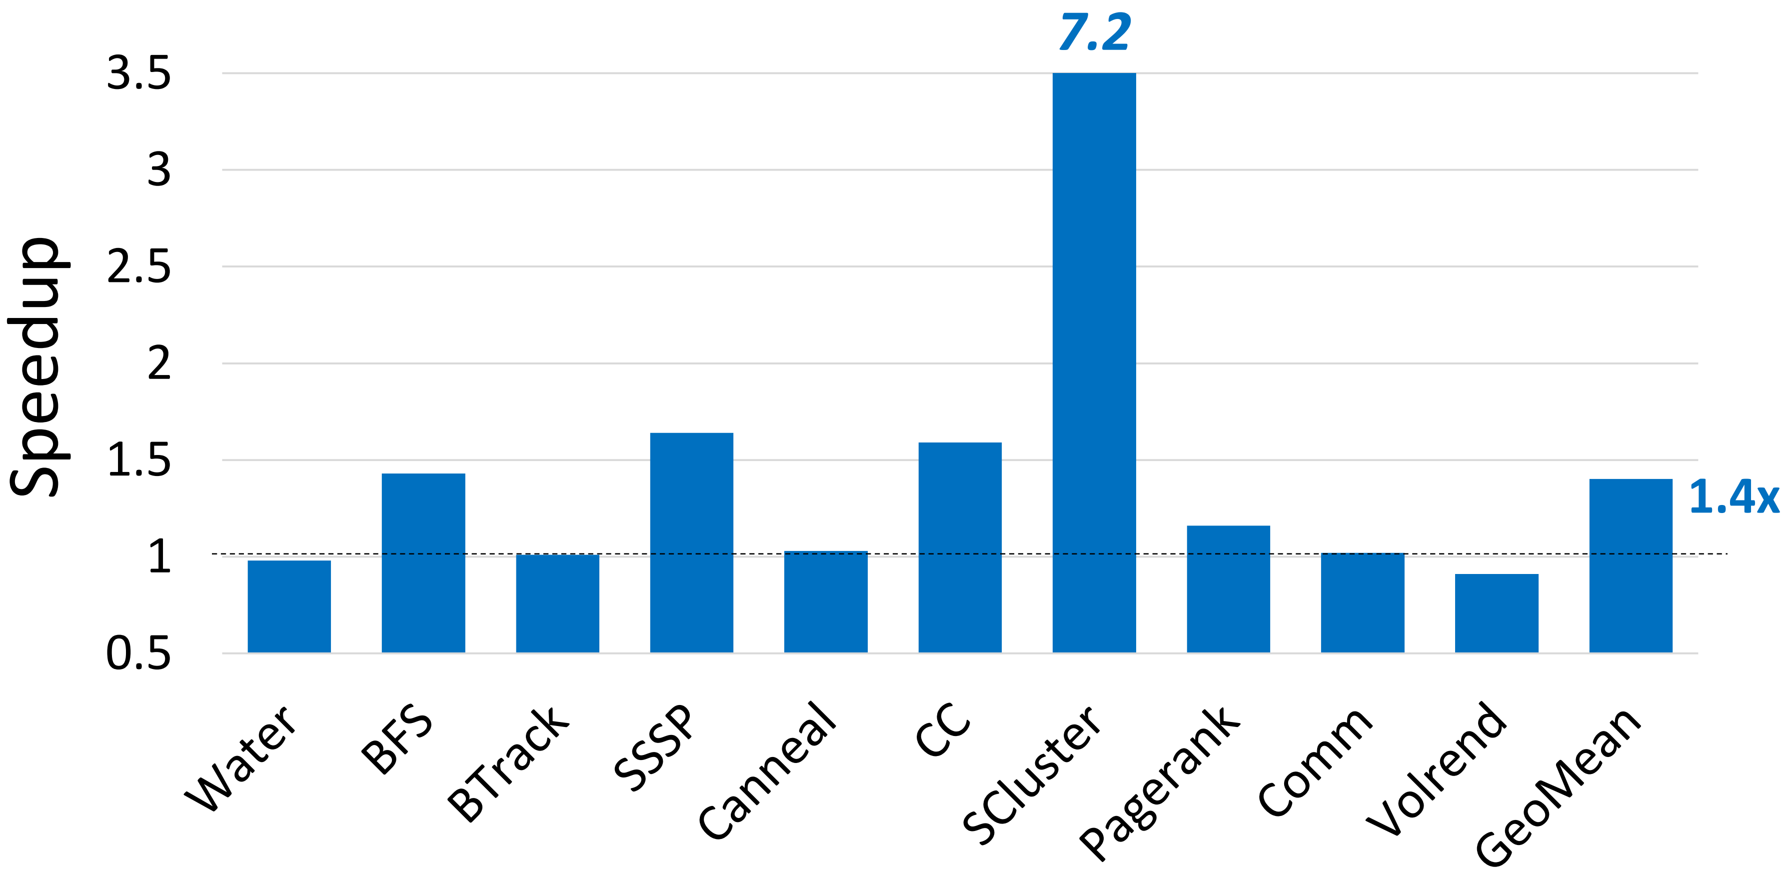

1.4x speed up over conventional wired multicore (Geometric Mean)

[ASPLOS '16] *"WiSync: An Architecture for Fast Synchronization through On-Chip Wireless Communication", Abadal et al.*

Replica: BMem for synch + shared data

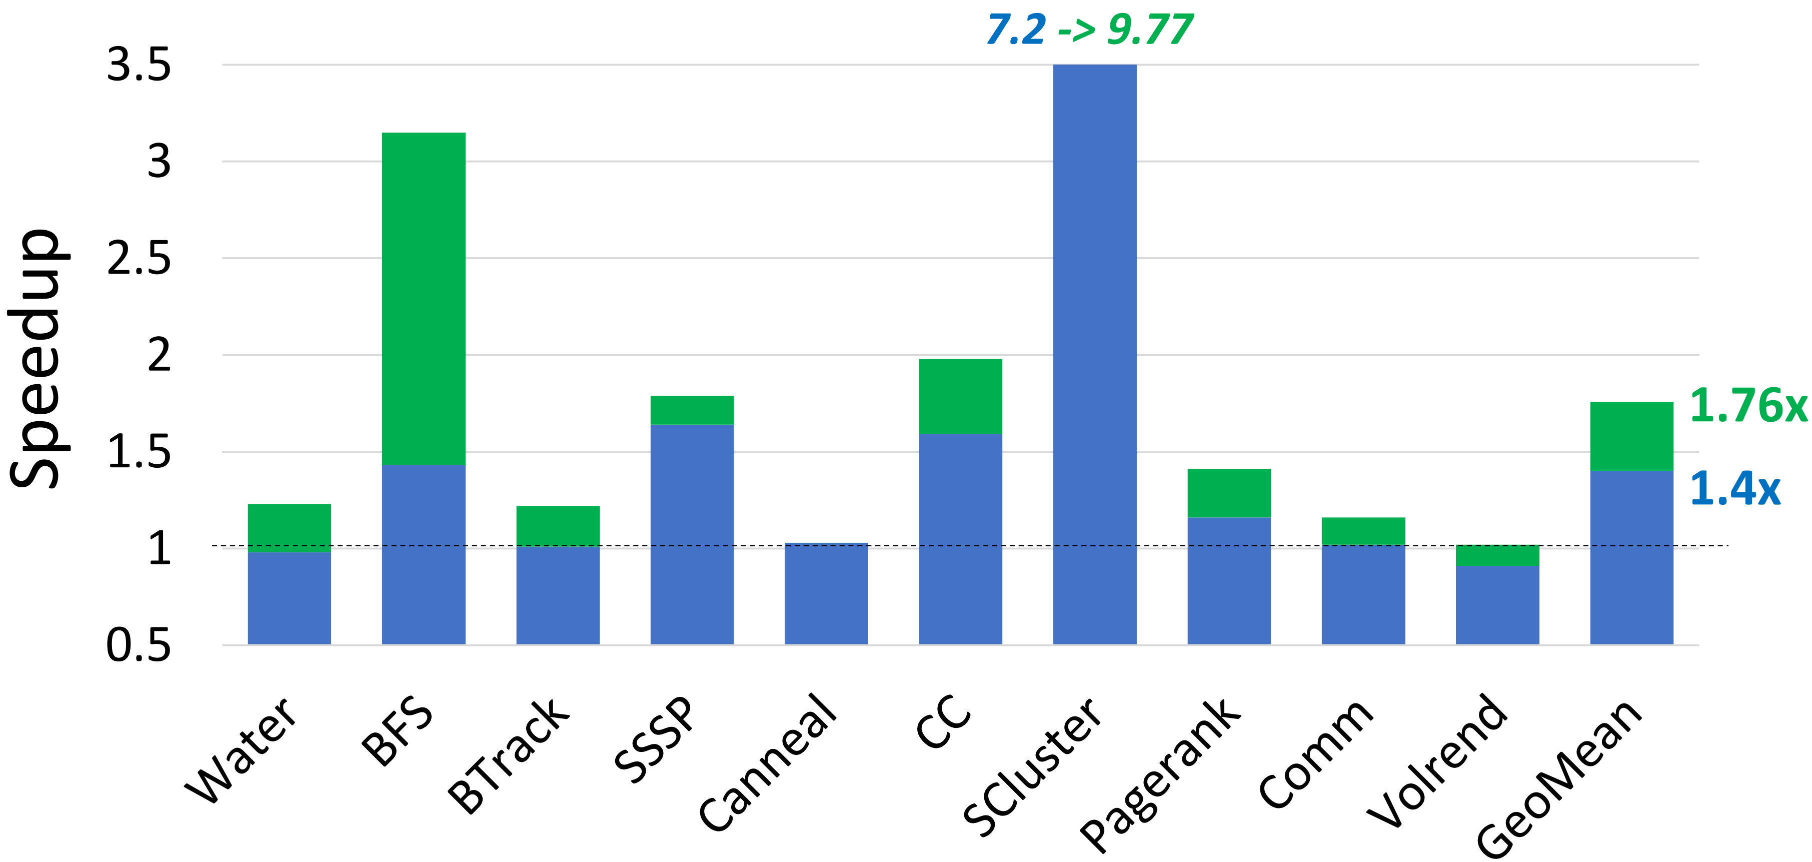

1.76x speed up (Geometric Mean)

[ASPLOS '19] *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.", Fernando et al.*

Replica: BMem for shared data + approx

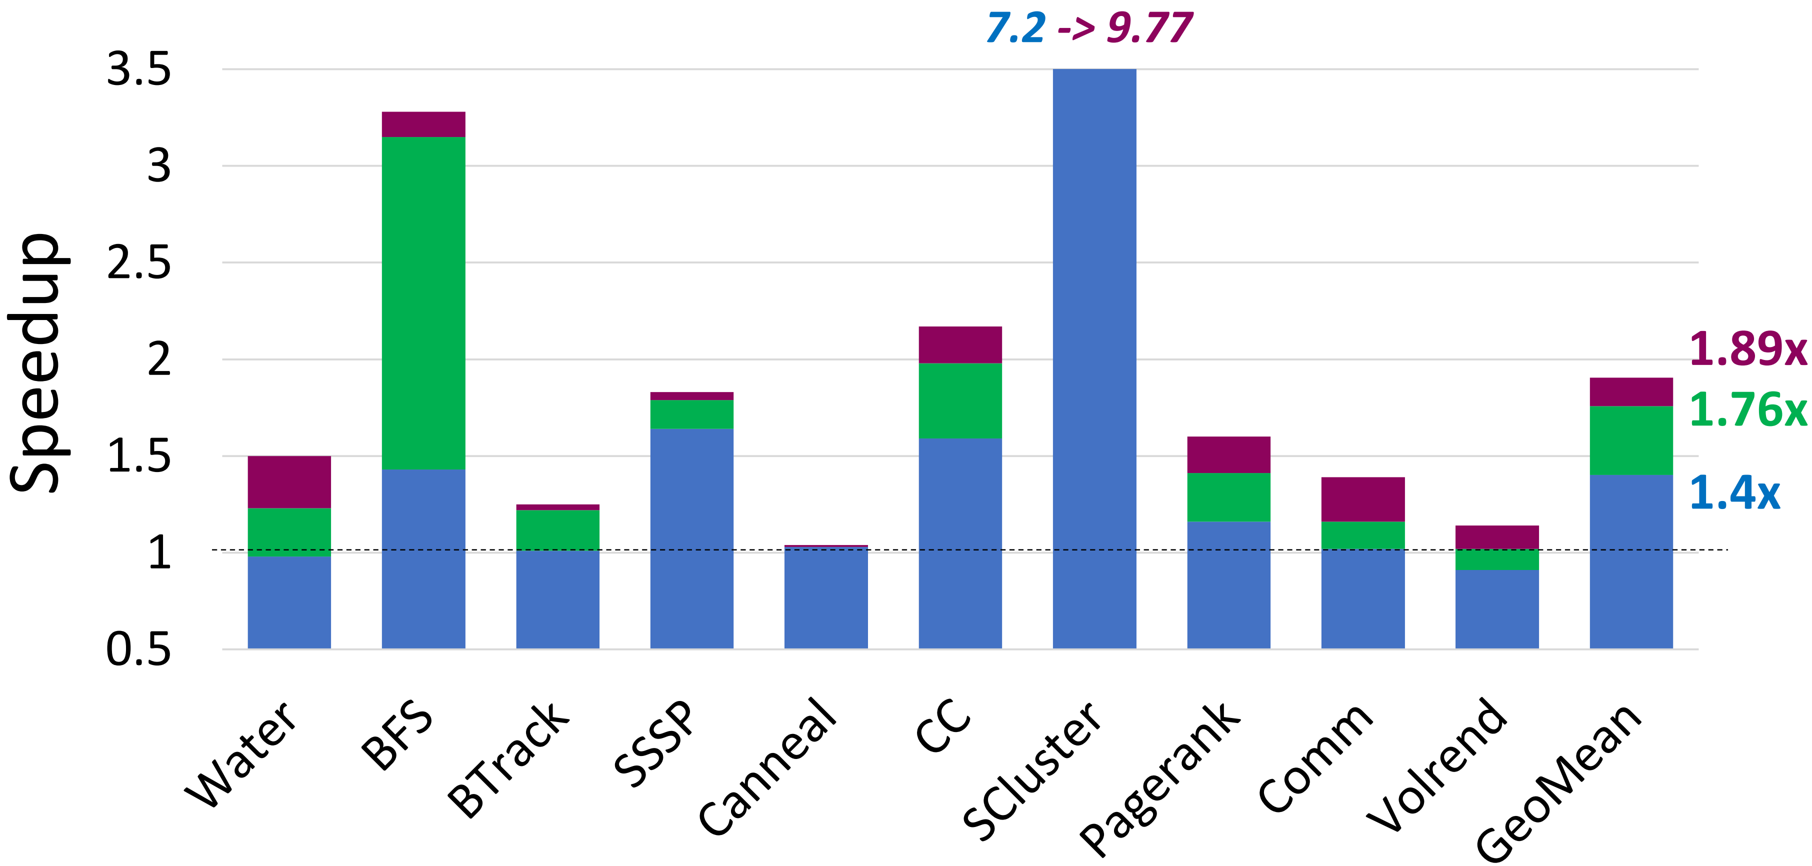

1.89x speed up (Geometric Mean)

[ASPLOS '19] *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.", Fernando et al.*

Energy and area

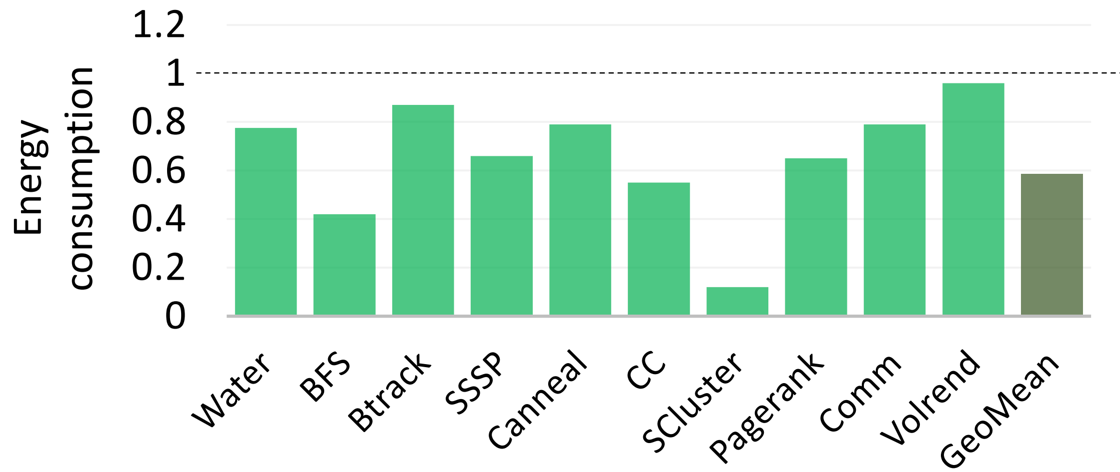

- Faster execution → 33% energy reduction
- Replica components: 9% of total energy consumed

- 15% area overhead (over standalone wired NoC)
  - 11% from the BMem + 4% from the transceiver/antenna
  - Using the same area to increase the L2 cache has little impact on performance (1.04x speedup)

Challenges: Network + Link

- Medium Access Control (MAC):
  - Different applications → different sharing patterns
  - Traffic varies within and across applications
  - Needs to deal with bursts and hotspots
  - Contention based (CSMA) vs. controlled access (Token)

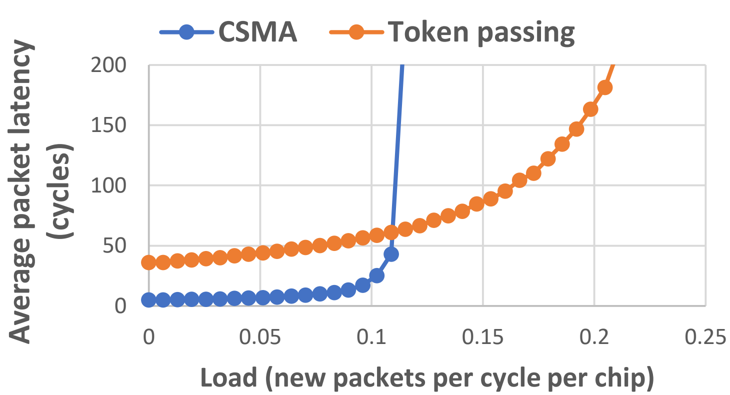

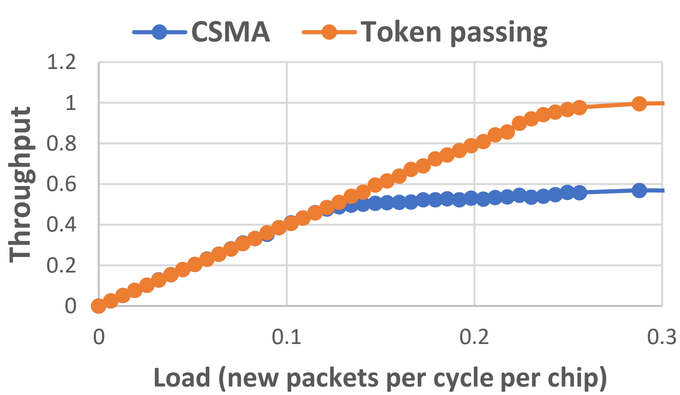

Opportunities: Network + Link

- Medium Access Control (MAC):
  - Static environment → no hidden terminal problem
  - Need low overhead → nanoseconds matter!
    - *Example:* **Broadcast Reliability Sensing protocol (BRS)***

- Wireless NoC is multipurpose:
  - Number of participants: broadcast vs point-to-point+
  - Latency-critical data (send critical word first, rest later)
  - Multiplexing techniques? E.g. FDMA, SDMA, TDMA…

***[NoCArc '16]** *"A MAC protocol for Reliable Broadcast Communications in Wireless Network-on-Chip", Mestres et al.*
+[IEEE Emerging '12] *"Wireless NoC as Interconnection Backbone for Multicore Chips: Promises and Challenges", Deb et al.*

Replica : Wireless Message

- 80 bits* (13-bit address + 64-bit data + 3-bit checksum/other)

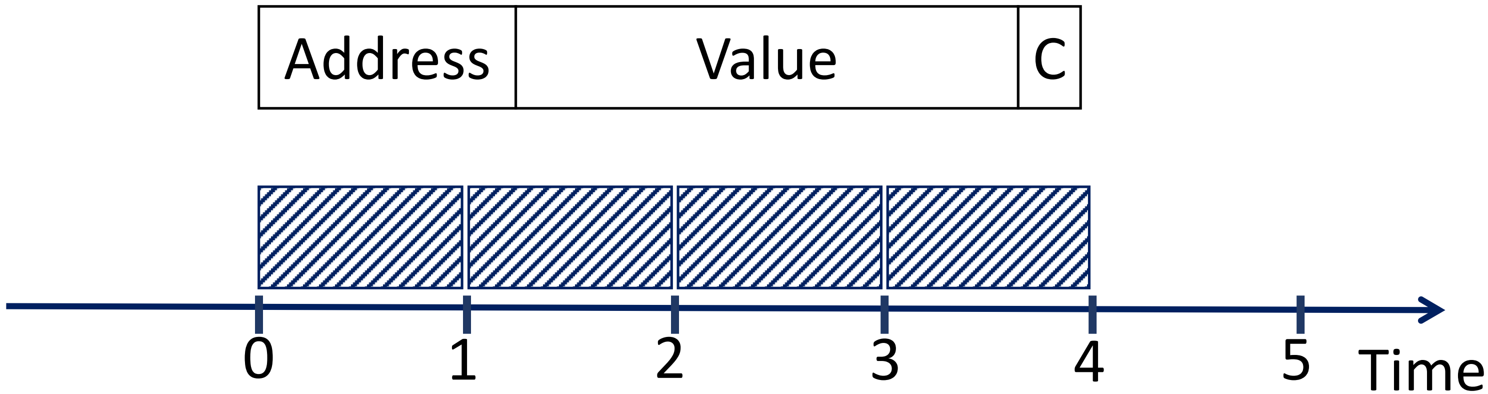

4 cycles at 20Gb/s†

*[ASPLOS '19] *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.", Fernando et al.*

†[IEEE Design & Test '14] *"Architecture and Design of Mul5-Channel Millimeter-Wave Wireless Network-on-Chip", Yu et al.*

Broadcast **R**eliability **S**ensing Protocol

- Collision detection with half-duplex radio & minimal latency
- An ACK-free version of CSMA → ACKs are too expensive!

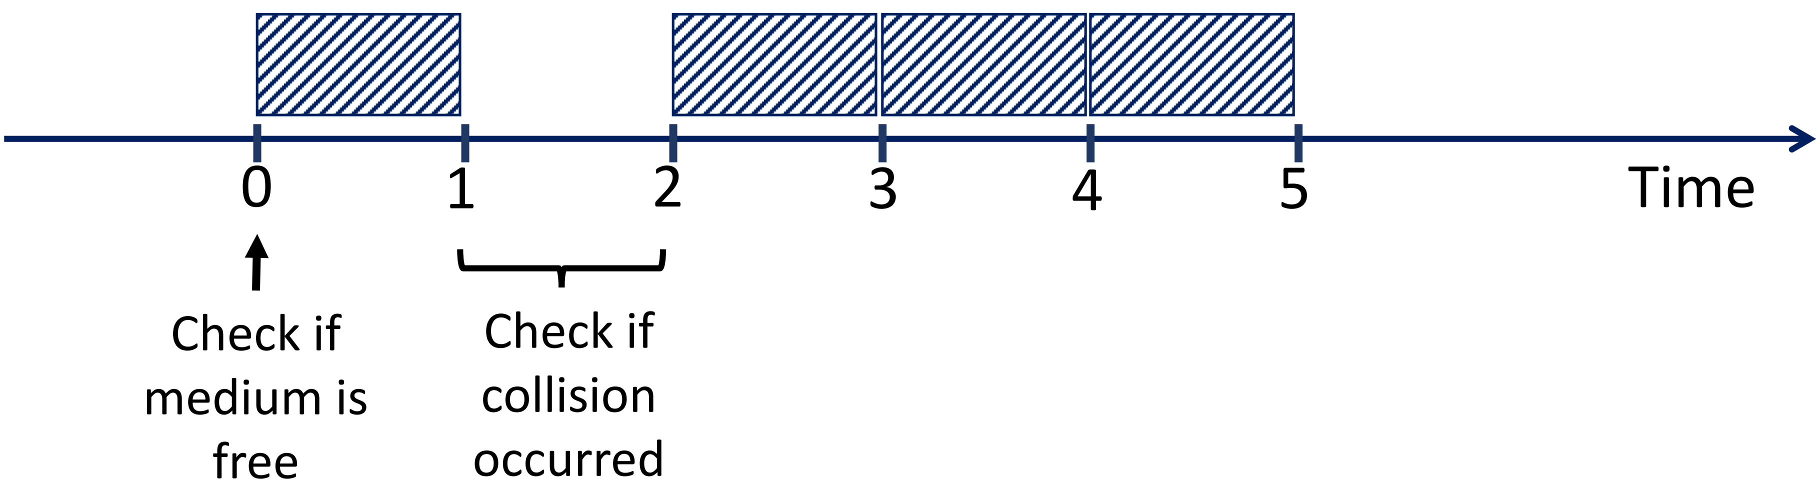

[NoCArc '16] *"A MAC protocol for Reliable Broadcast Communications in Wireless Network-on-Chip", Mestres et al.*

Opportunities: Network + Link

- *Another example:* **Replica Adaptive Wireless Protocol**
  - Sparse traffic → use BRS (transmit ASAP)
  - Bursty traffic → use Token-Ring (avoid starvation)
  - Runtime: switch BRS↔Token by observing behavior
    - # of collisions
    - # of skipped token slots

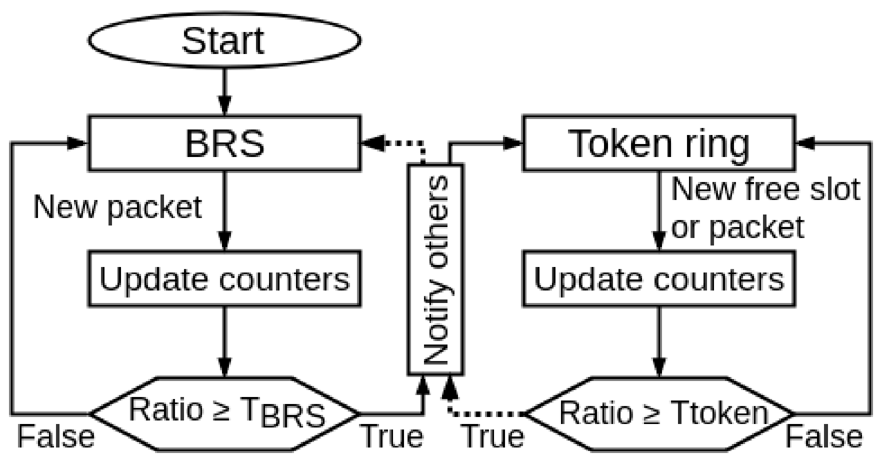

[ASPLOS '19] *"Replica: A Wireless Manycore for Communication-Intensive and Approximate Data.", Fernando et al.*

Challenges: Physical

- Transceiver trade-offs: performance vs area/energy
  - # of channels → parallelism vs. wasted bandwidth
  - Spectral efficiency vs. Complexity of modulation
- Antenna efficiency
  - Shape must fit chip package → potentially sub-optimal size
- Tx power vs. Bit Error Rate (also determined by the medium)

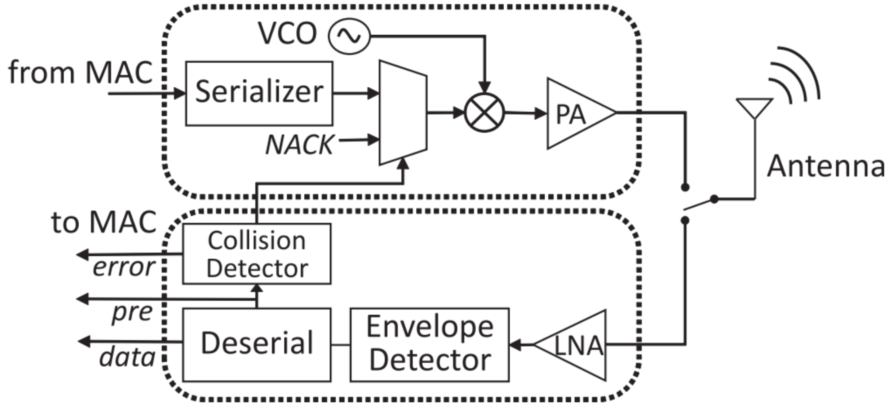

[IEEE Transactions '18] *"OrthoNoC: A Broadcast-Oriented Dual-Plane Wireless Network-on-Chip Architecture", Abadal et al.*

Opportunities: Physical

- WNoC → Many antennas next to each other
  - Undesired coupling effects?
  - Controlled environment → Opportunistic beamforming!

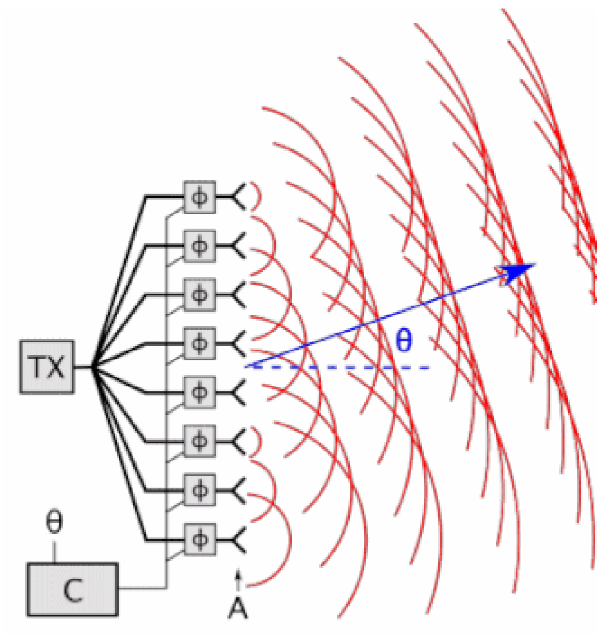

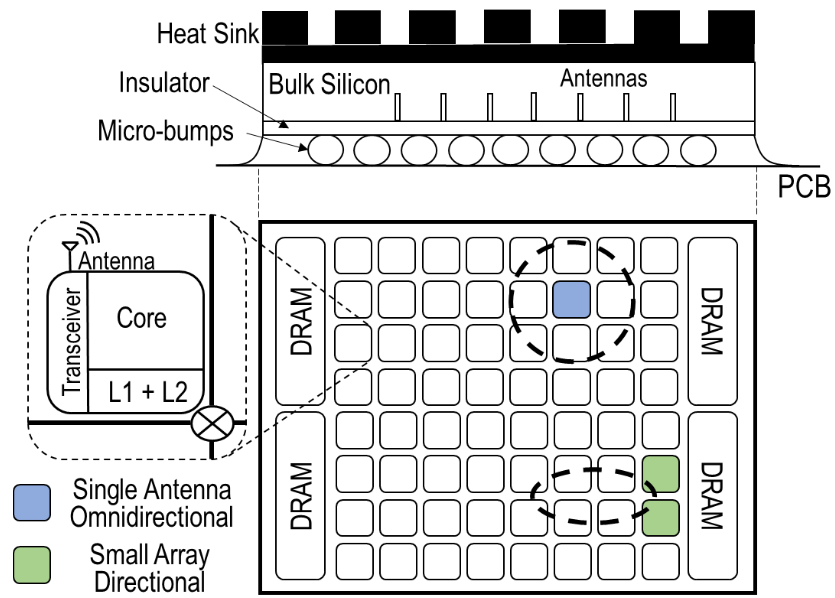

[ISCAS '19] *"Opportunistic Beamforming in Wireless Network-on-Chip", Abadal et al.*

Opportunities: Physical

- WNoC → Many antennas next to each other
  - Undesired coupling effects?
  - Controlled environment → Space-Division Multiplexing!

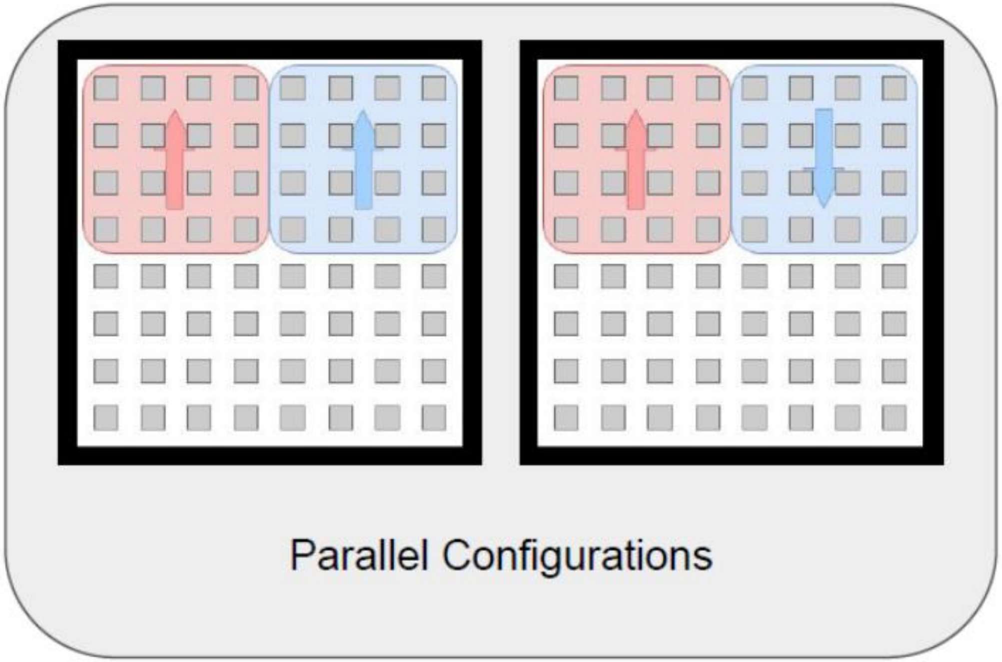

[ISCAS '19] *"Opportunistic Beamforming in Wireless Network-on-Chip", Abadal et al.*

Challenges: Medium

- The intra-chip wireless channel is relatively unknown
  - Characterization is essential for transceiver implementation
  - Free-space conditions?
  - Same rules as off-chip transmission?
- There is a need to understand the chip package effects on propagation

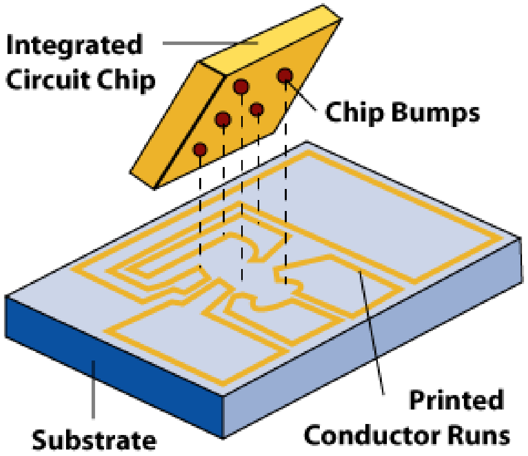

Opportunities: Medium

- Make a realistic chip package model

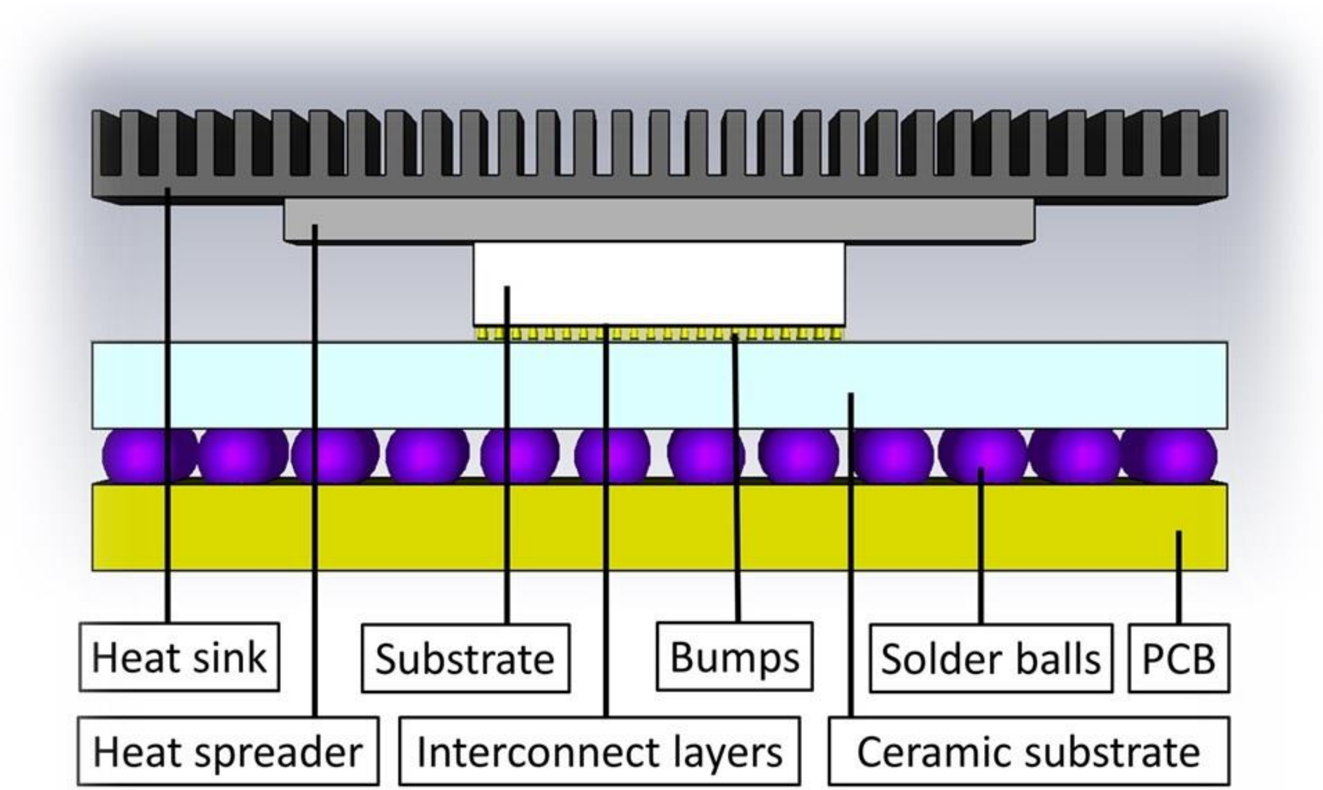

[ISCAS '18] *"Millimeter-Wave Propagation within a Computer Chip Package", Timoneda et al.*

Opportunities: Medium

- Investigate wave propagation with full-wave solver (CST)

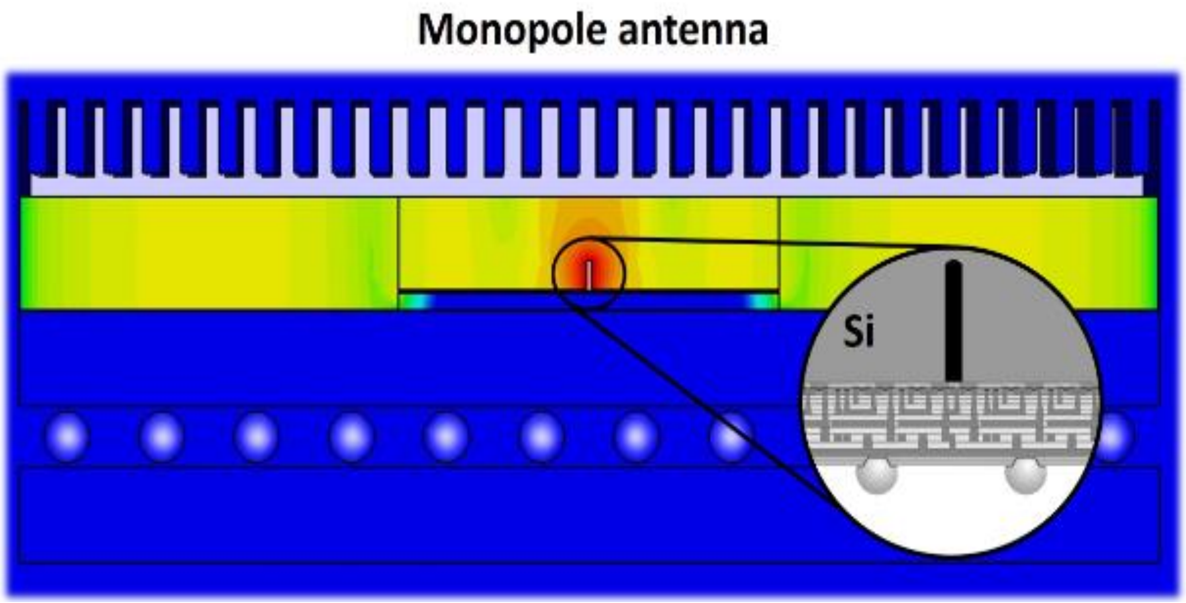

[ISCAS '18] *"Millimeter-Wave Propagation within a Computer Chip Package", Timoneda et al.*

Opportunities: Medium

- Tweak package parameters to minimize path loss
  - *AIN*: Aluminium Nitride (Dielectric)
  - *Si*: Silicon
  - S_{min}: worst-case path loss among all freqs. in bandwidth

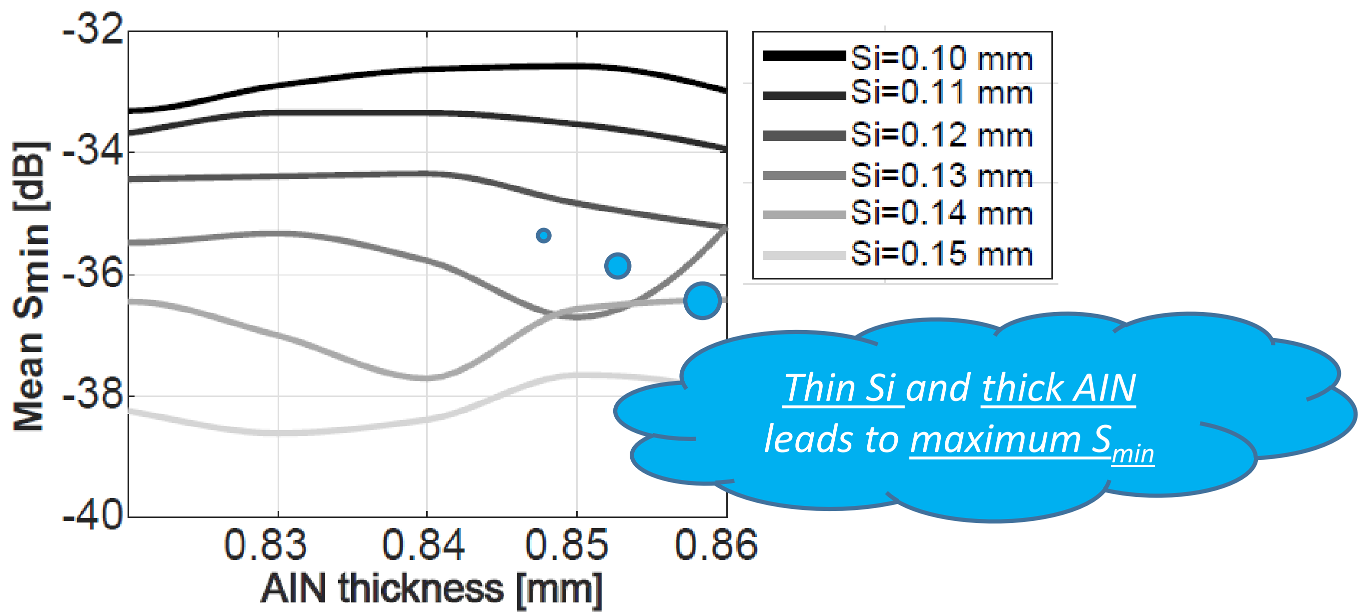

[ISCAS '18] *"Millimeter-Wave Propagation within a Computer Chip Package", Timoneda et al.*

Challenges/Opportunities: Simulation

- Wireless NoCs usually have > 64 cores:
  - Architecture simulators do not scale well (long sim. time)
  - If they scale, they do not model network in detail
- Architecture simulators:
  - Gem5, Multi2sim, Zsim, Prime, Sniper, etc.
- Network simulators (good scalability):
  - Noxim (includes WNoC capabilities)
  - Garnet2.0 (integrated into Gem5)
  - Booksim2, ATLAS, DARSIM, FlexNoC, etc.
- **Can we get scalable integrated arch+net simulators?**

Thank you!

WNoC unique points:

- Multi layered system with many co-design opportunities
- Ability to break long-standing architecture assumptions
- Static environment opens up door for redesign of many off-chip wireless protocols and schemes

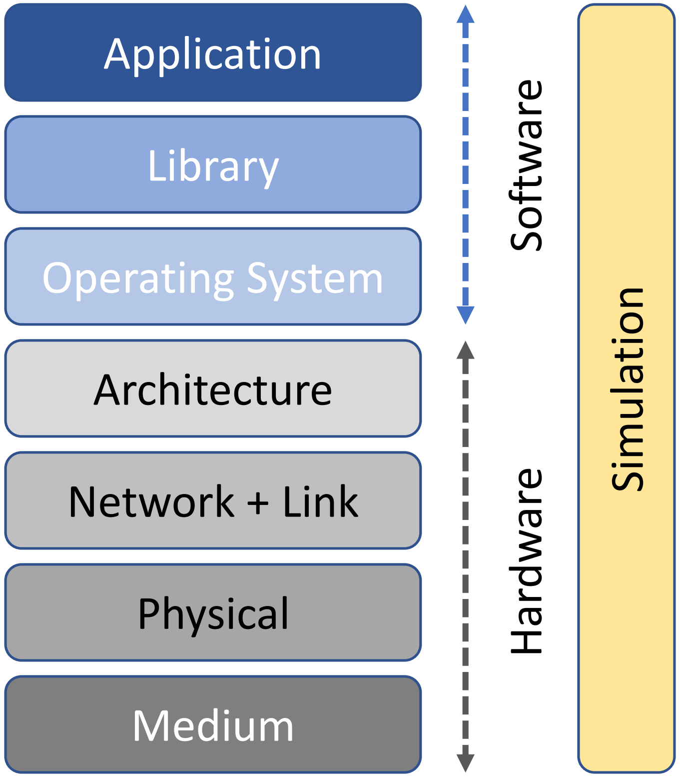

Extra slides

Opportunistic beamforming

- The solution is opportunistic because
  - It exploits already existing antennas
  - Cores are, by definition, tightly synchronized through the system clock
  - Data may be already present in several cores due to architectural/software mechanisms
  - With few directions (row/column) the scheme is already relevant; coarse-grained steering is enough. This simplifies the phase shifters

Opportunistic Beamforming

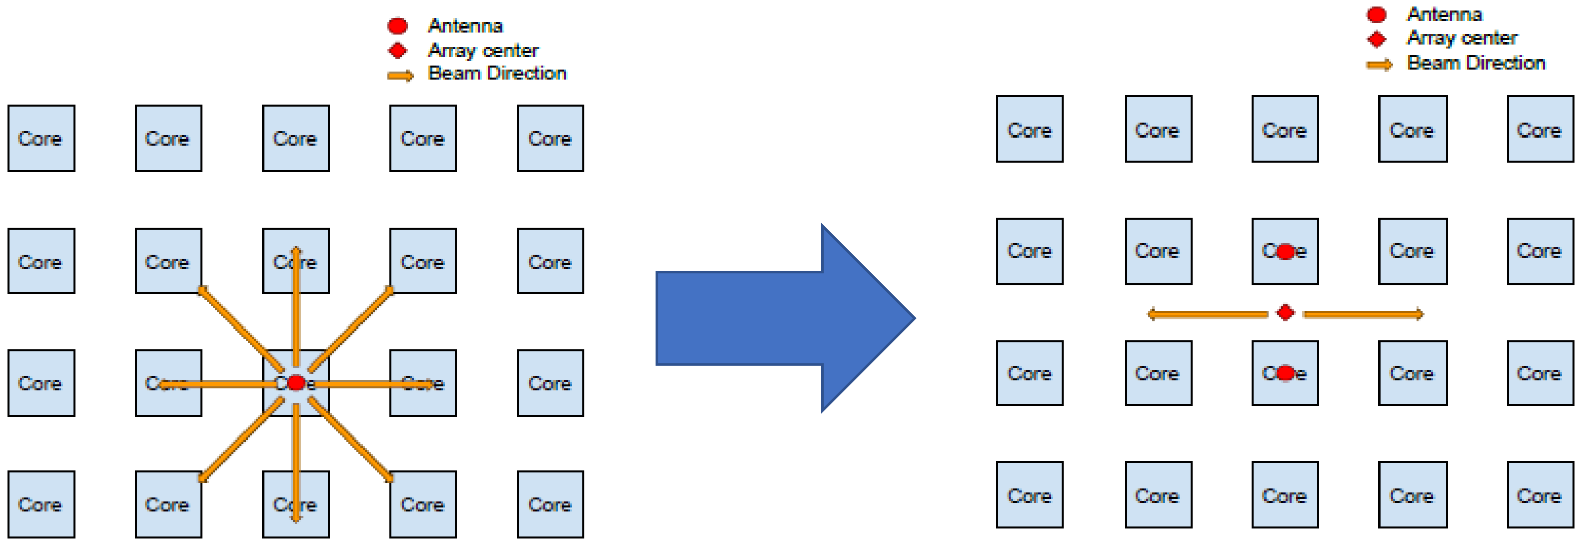

- Each core has its antenna and transceiver
- By default, broadcasts through one antenna
- If needed, two or more antennas are activated simultaneously through a shared controller

Walkthrough example: 2x2 cluster

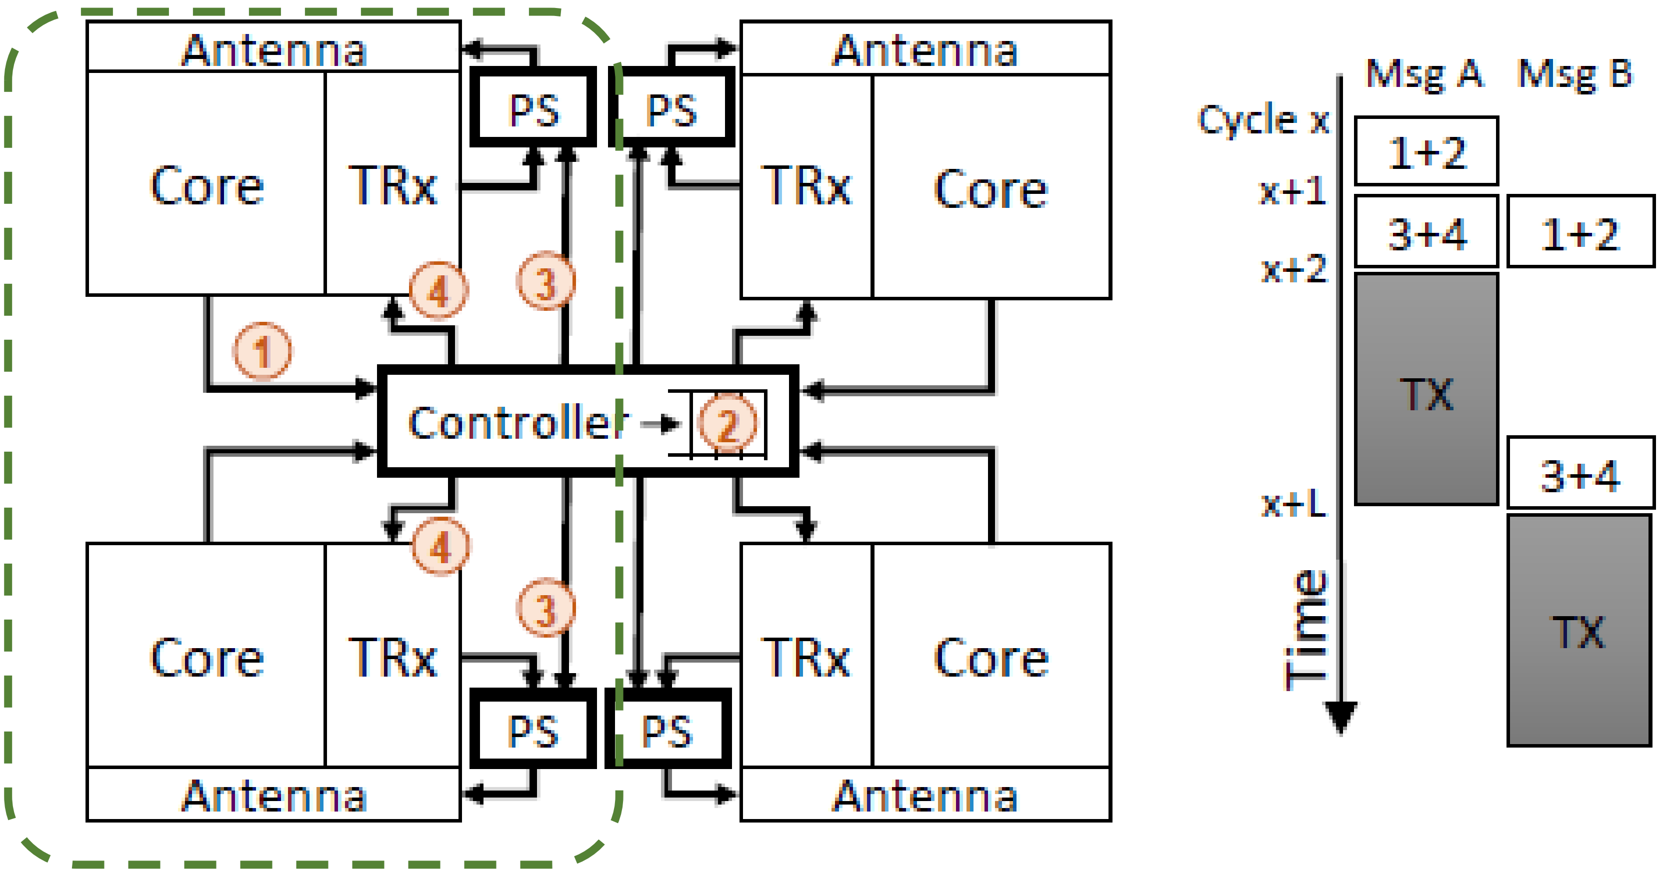

Opportunistic Beamforming

- The separation between antennas is chosen to be λ/4
  - At mmWave, compatible with core sizes
  - It avoids grating lobes by being < λ
  - Allows to obtain single-sided patterns

Integrated RF Transceivers

- For the WNoC idea to make sense, we need
  - 10+ Gbps
  - ~1 pJ/bit
  - ~100 Gb/s/mm^2
- Transceiver design is getting up to speed

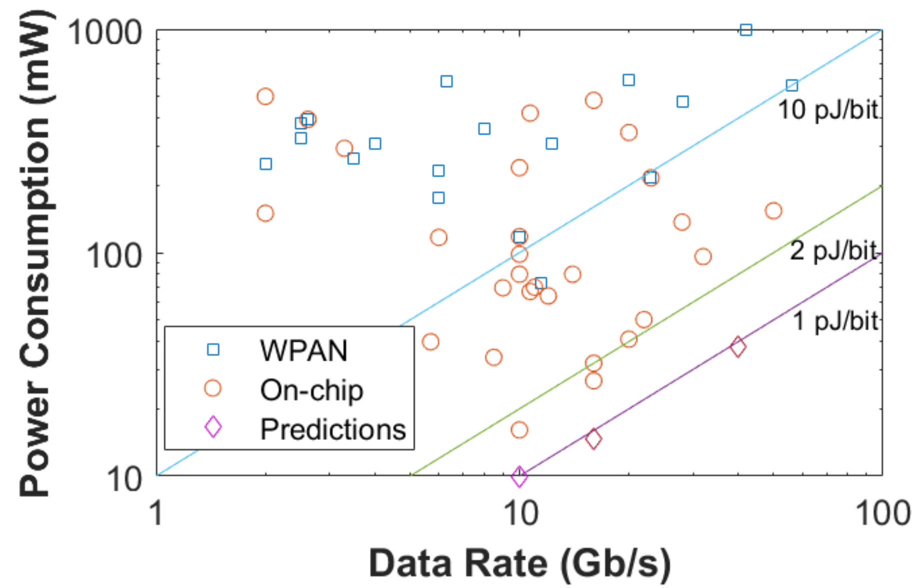

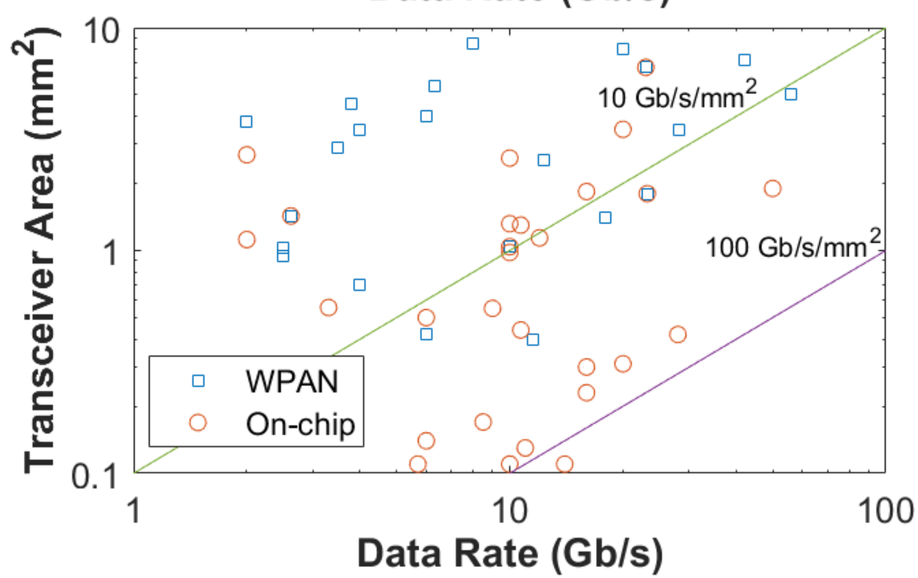

S. Abadal, M. Iannazzo, M. Nemirovsky, A. Cabellos-Aparicio, H. Lee, E. Alarcón,"On the Area and Energy Scalability of Wireless Network-on-Chip: A Model-based Benchmarked Design Space Exploration," IEEE/ACM Transactions on Networking, vol. 23, no. 5, pp. 1501-1513, Oct. 2015.