Explicit Communication and Synchronization in SARC

M. Katevenis, V. Papaefstathiou, S. Kavadias, G. Nikiforos,
D. Pnevmatikatos, D. Nikolopoulos, C. Kachris

Features:

- Direct Access Address Space Extension:
  - SW Allocates part of *private cache* as *scratchpads*
- Supports both:
  - **Implicit Communication** using caches,
  - **Explicit Communication** via directly accessible local memories (scratchpad)
- Merged **Cache controller** (CC) and **Network Interface** for reduced latency

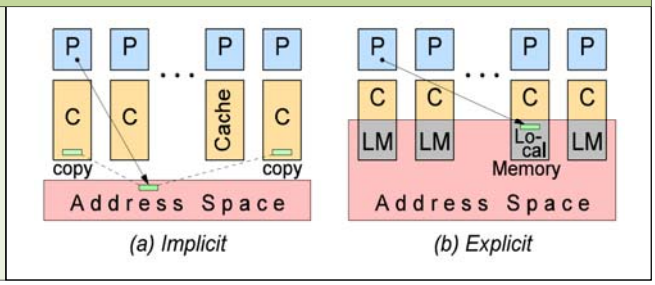

Explicit Communication Mechanisms:

- **Remote Stores**
  - Fast non-blocking direct transfers
  - Can use write combining to improve data comm. efficiency
  - Store fences may be required in SW
- **R-DMA**
  - Bulk asynchronous transfers
  - Completion notifications: combination of Acks and Counter

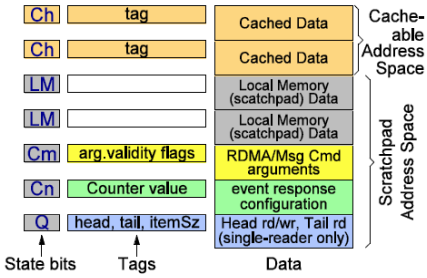

NI Communication & Synchronization primitives:

- **Command Buffers**
- **Counters**: Atomic add-on-store
  - Event response, send notification when counter=0
- **Queues**: Atomic enqueue and dequeue
  - Singe Reader: Many-to-One communication
  - Multiple Reader: Many-to-Many communication
- **Locks** (using queues)

Cluster Nodes:

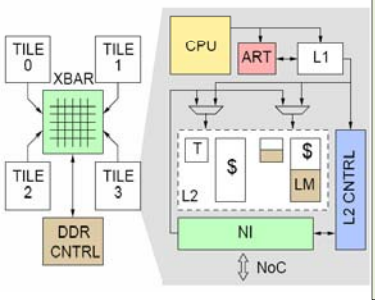

- Four MicroBlaze processors (32-bit, in-order, 5-stage pipeline)
- Configurable L2 cache/scratchpad
- Merged NI and CC
- Communication using a 64-bit, 5-port bufferless crossbar switch

FPGA Prototype:

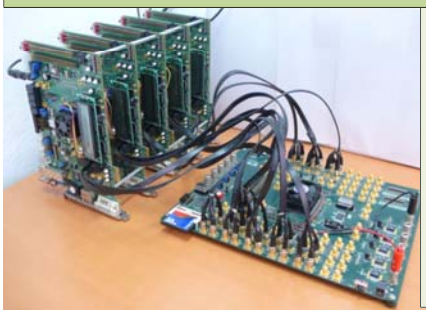

- Multi-FPGA Board (16 nodes based on Xilinx Virtex5 FPGA, Crossbar switch based on Xilinx VirtexII)
- Each node consists of 4 MicroBlaze processors and one DRAM (256 MBytes)

Performance Evaluation:

- Three Diverse applications
  Jacobi, Bitonic, FFT
- Configurations:
  - Plain Hardware managed caches
  - Hardware managed caches with strided hardware prefetching
  - Scratchpad memory with RDMA and Remote Stores
- 10%-40% higher speedup on 64 cores
- 2x-4x reduced network traffic
- 35%-70% Energy reduction
- 50%-90% Energy Delay Product (EDP) reduction
- 15%-30% Power consumption reduction