


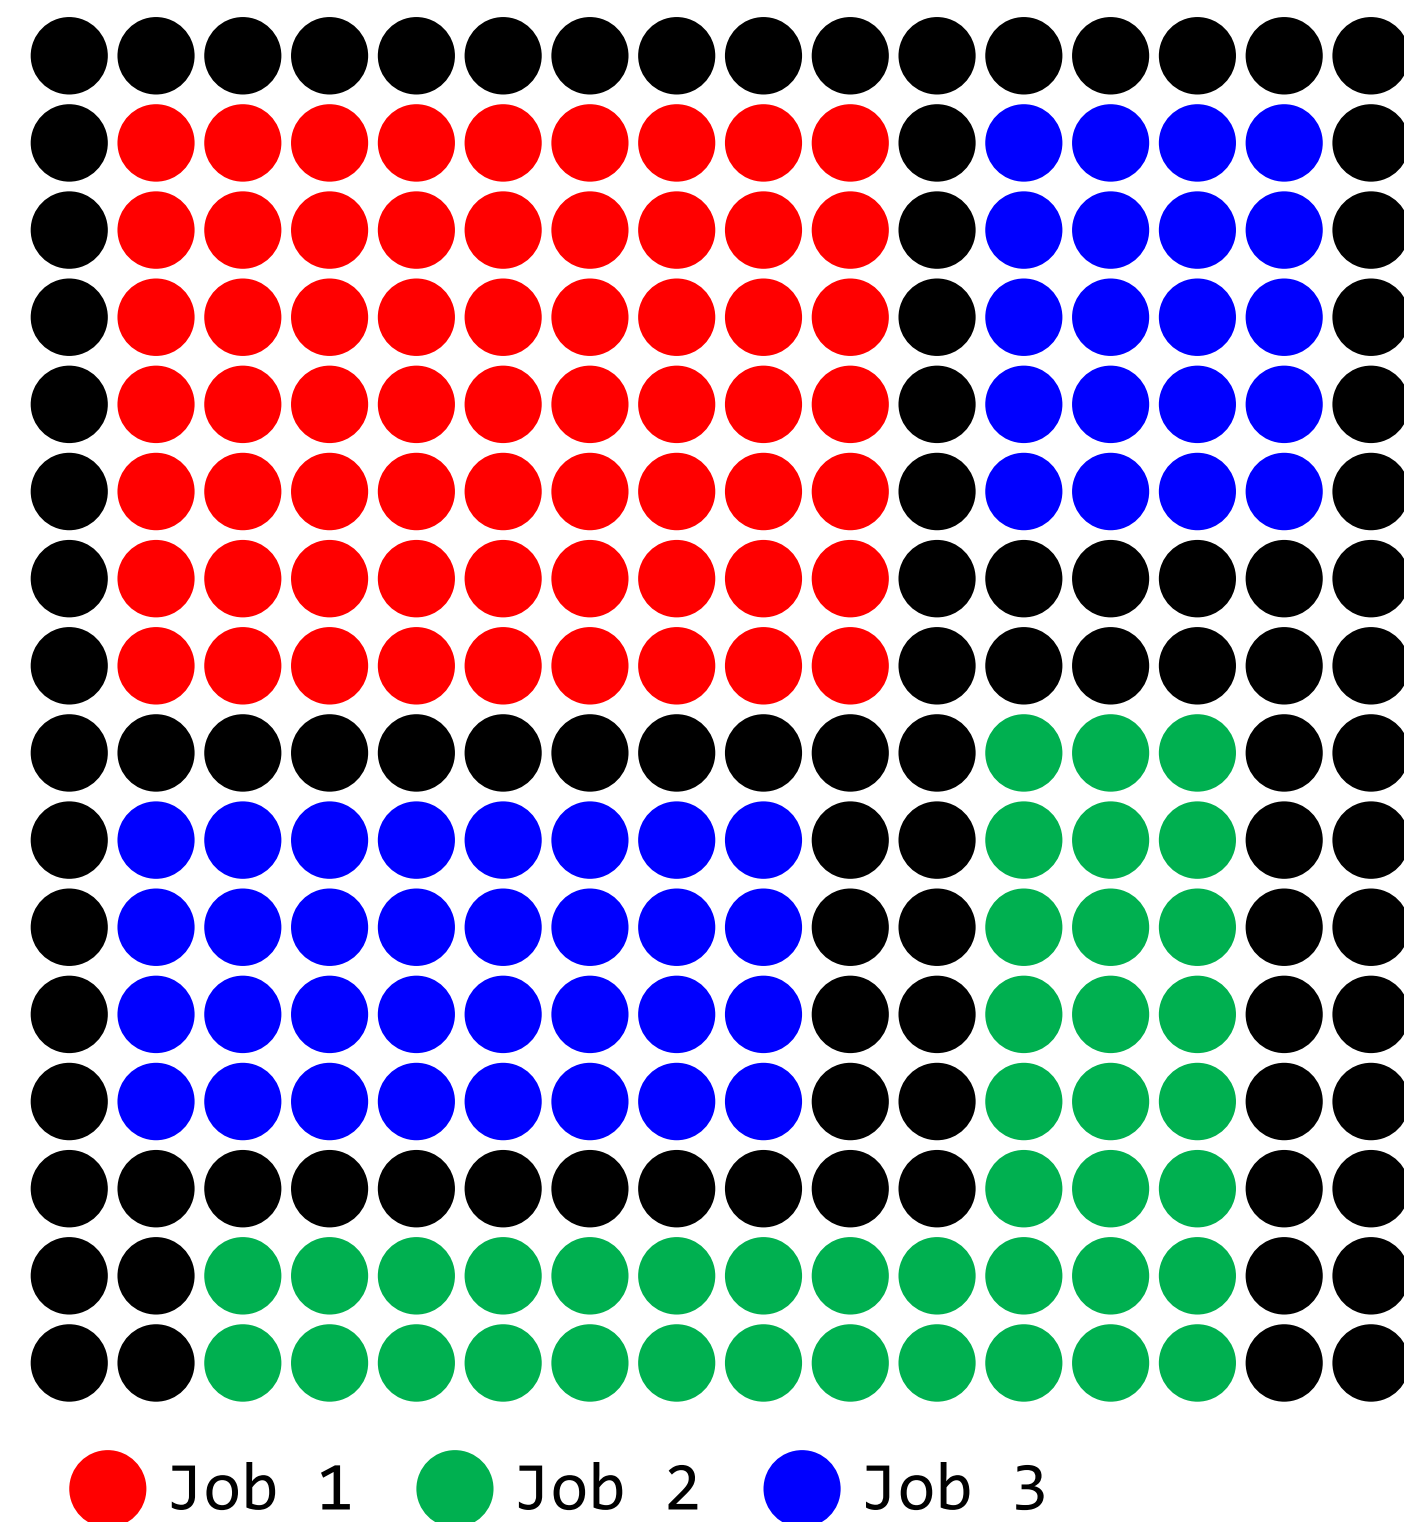
# Lightspeed Datacenter Networking

Nicholas McDonald, William J. Dally  
 Concurrent VLSI Architecture Group, Stanford University



## Network Level "Job" Awareness

A job is defined as a collection of processes correlated in the network to work together.




## Network Addressing

Physical addresses specify location.  
 Virtual addresses specify job node ID.  
 NIC performs address translation.

Benefits:

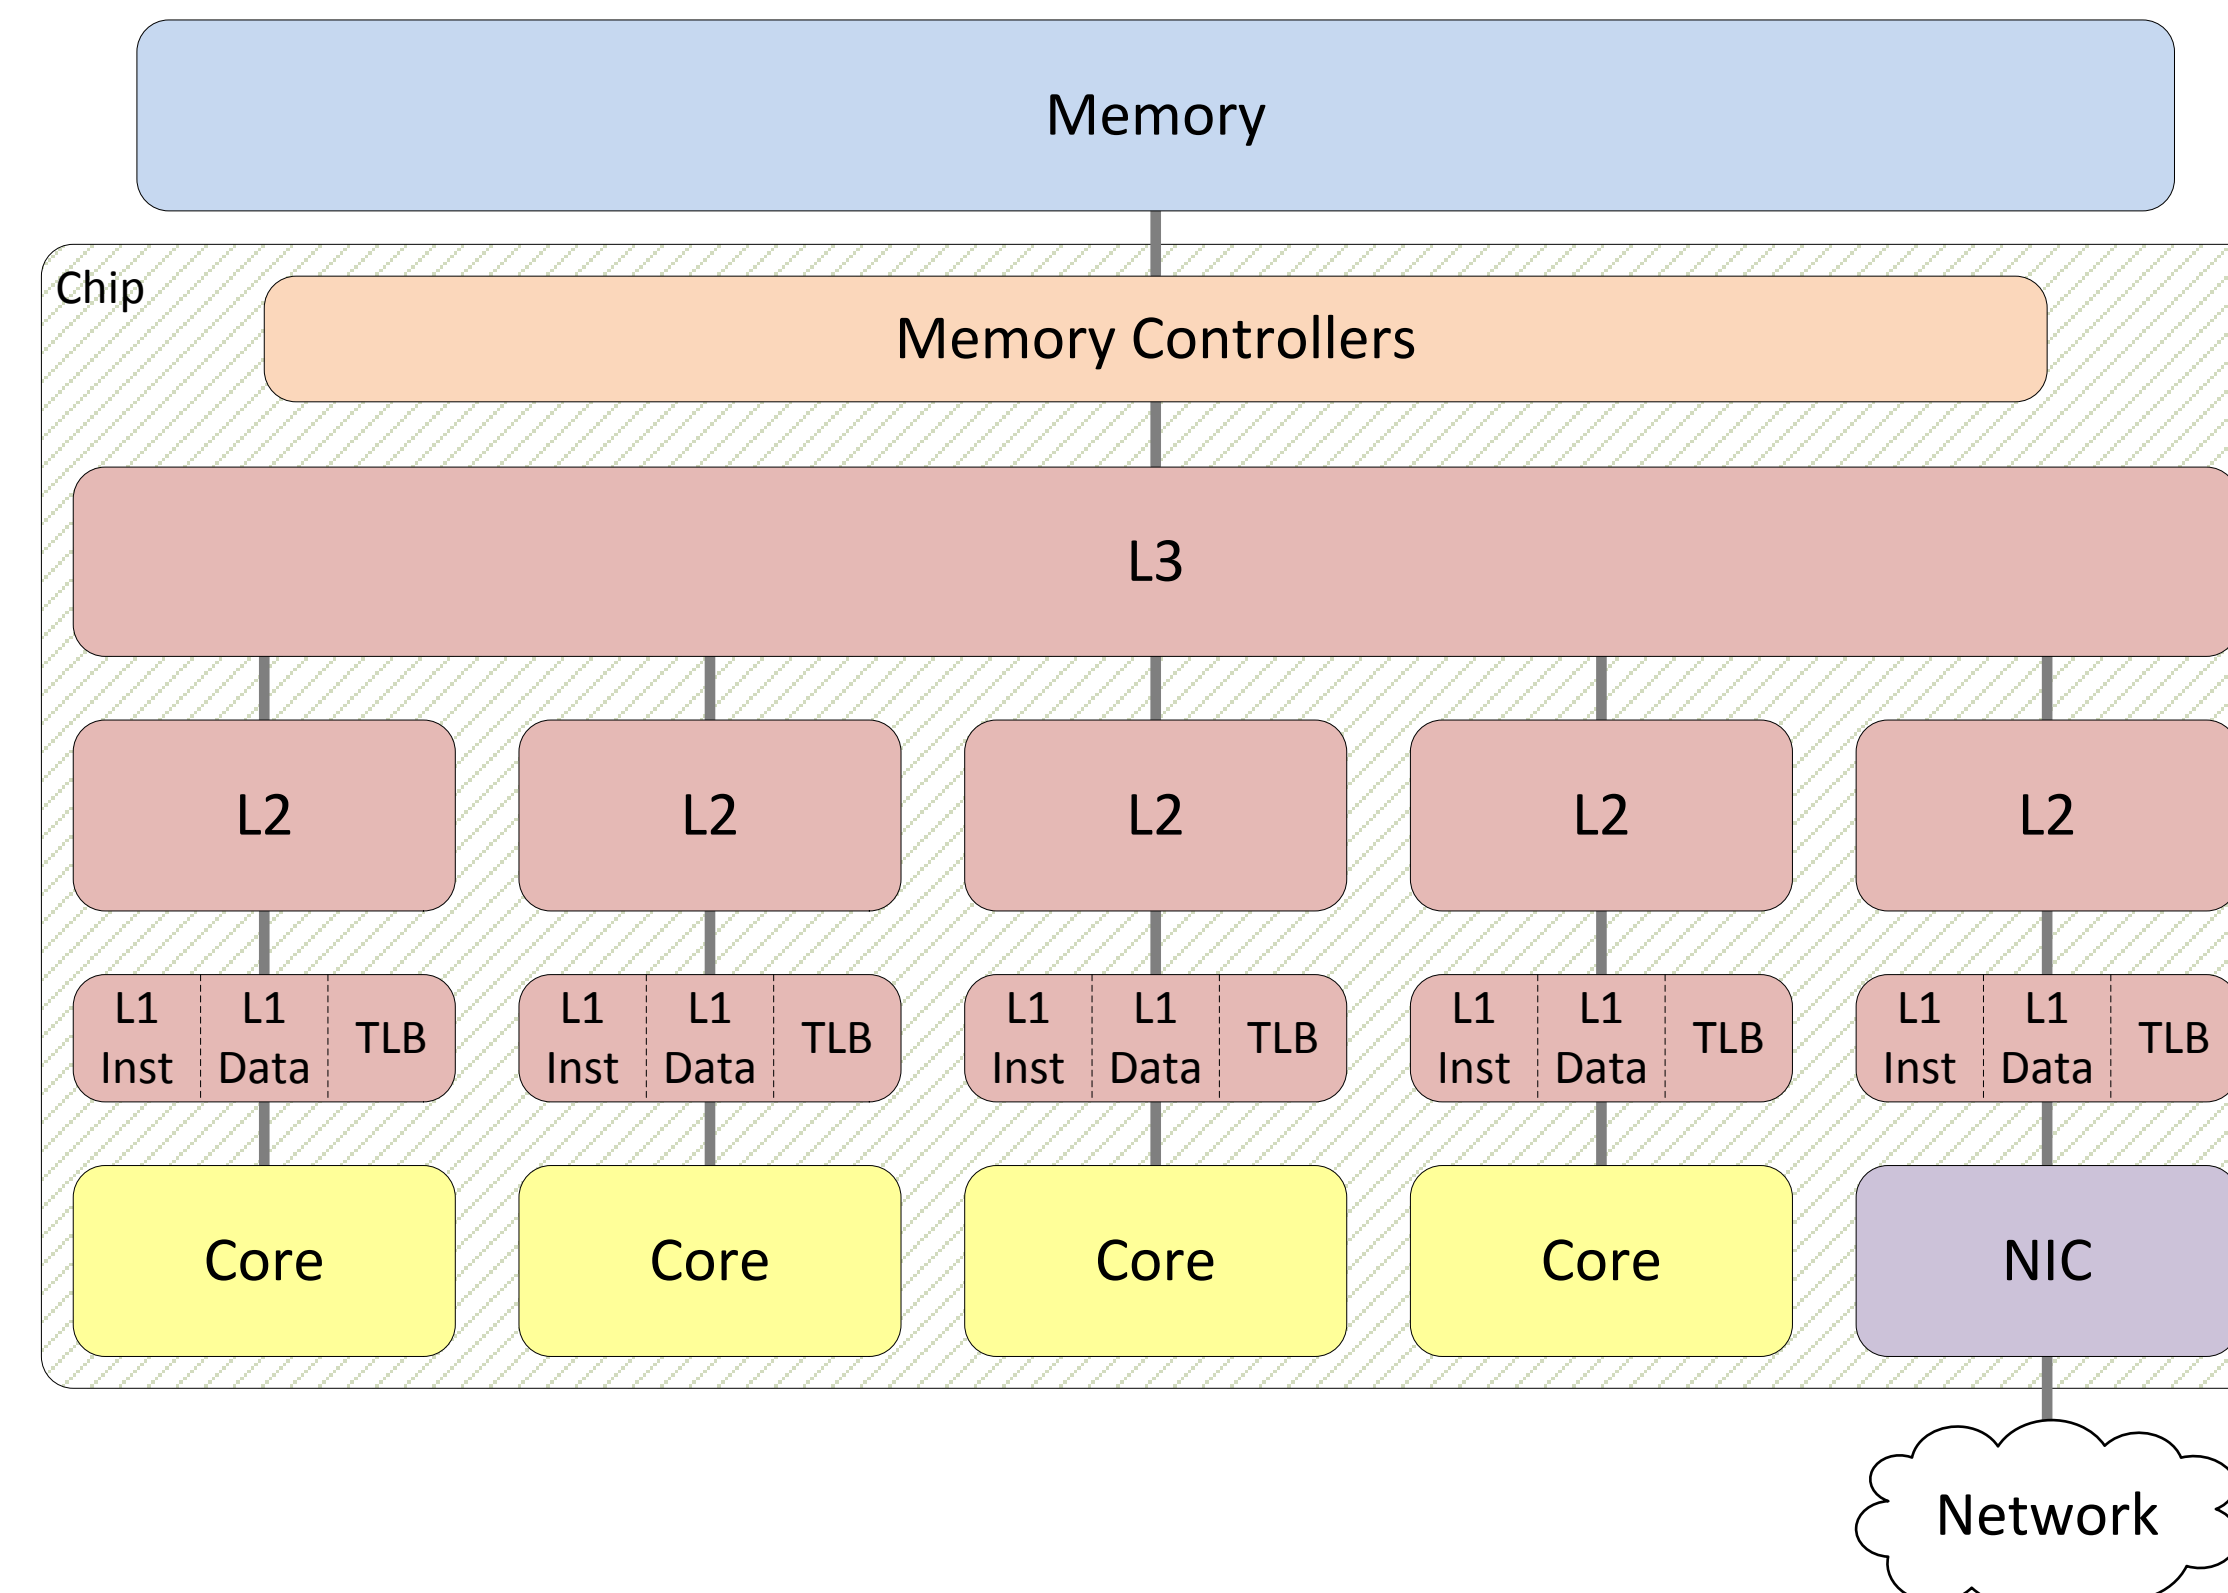
- Fast routing decisions
- Dense job location definitions
- Physical implementation abstract
- Flexible topology usage
- High scalability



Translation:  
 $\{ [A=2, B=3, C=5], [a, b, c], [n]=cnt \}$   
 $addressof(i) = [a + ((b + ((i+c)/C))/B), (b + ((i+c)/C)) \% B, (i+c) \% C]$   
 $addressof(2) = [0 + ((0 + ((2+2)/5))/3) = 0, (0 + ((2+2)/5)) \% 3 = 0, (2+2) \% 5 = 4]$   
 $addressof(9) = [0 + ((1 + ((9+3)/5))/3) = 1, (1 + ((9+3)/5)) \% 3 = 0, (9+3) \% 5 = 2]$   
 $addressof(1) = [1 + ((1 + ((1+4)/5))/3) = 1, (1 + ((1+4)/5)) \% 3 = 2, (1+4) \% 5 = 0]$

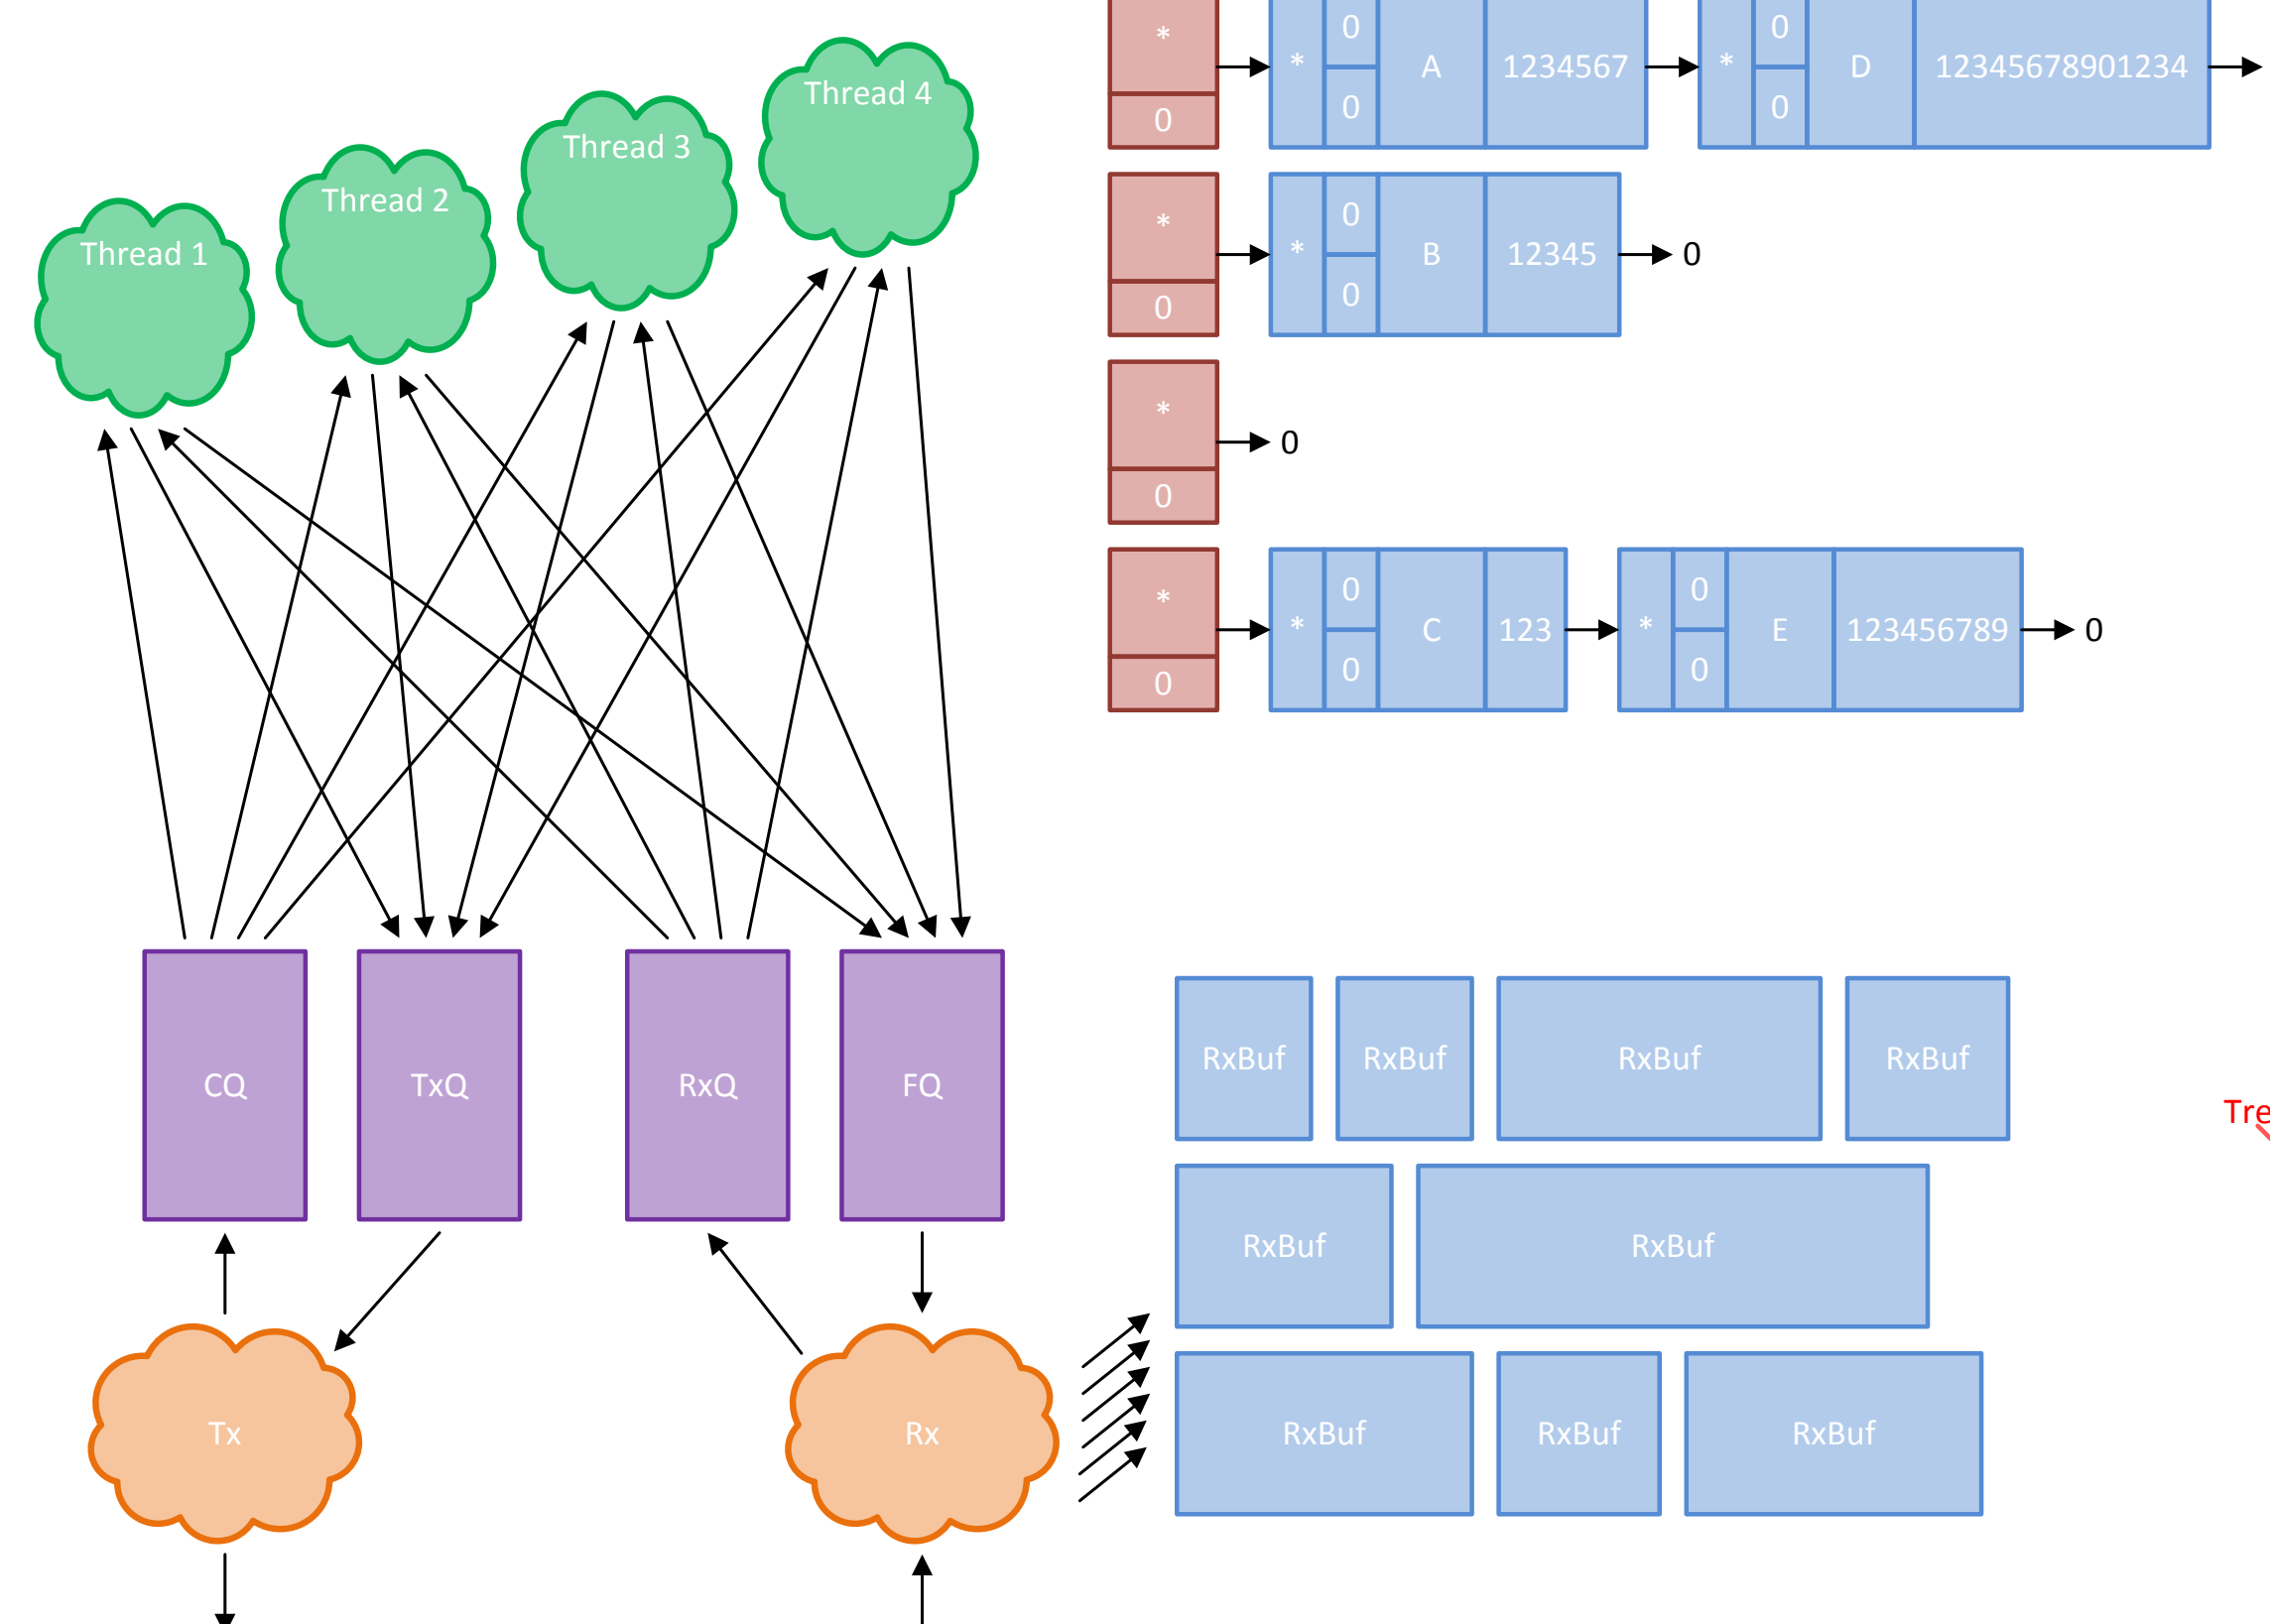
## On-Chip NIC Integration

- Integrate NIC into the memory system
- Provide fast user mode access to NIC



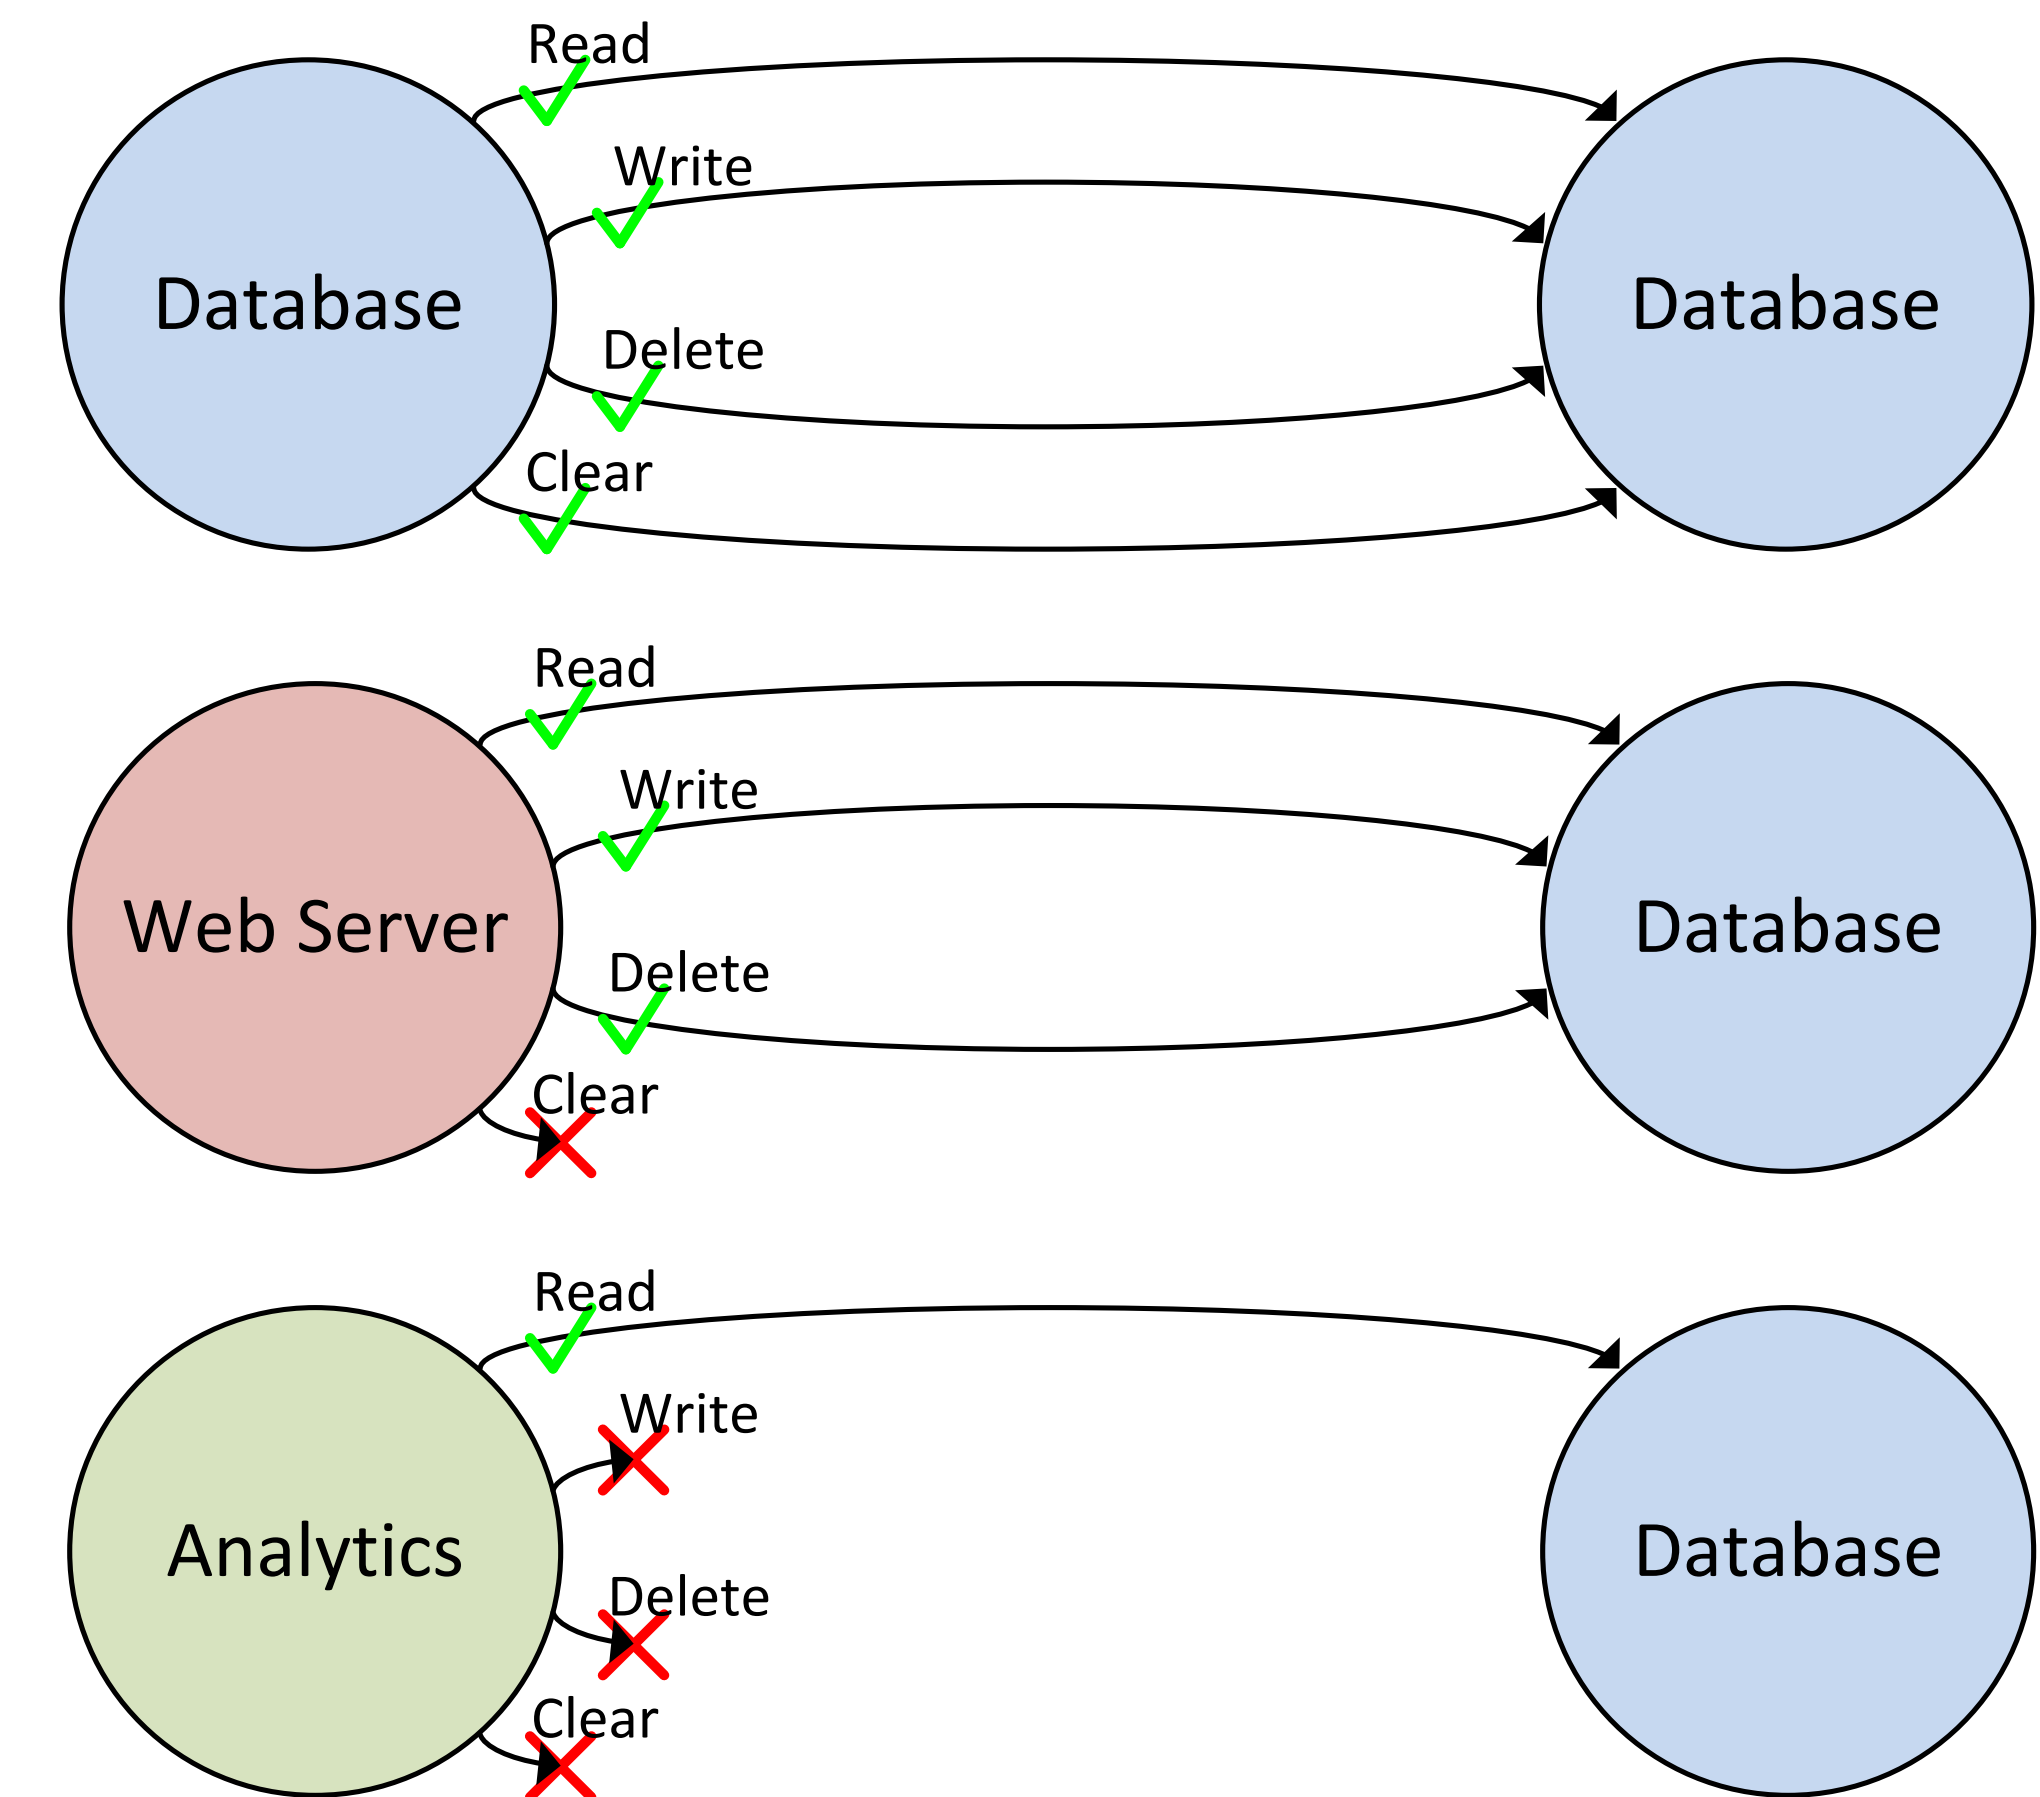
## High Performance Multi-Core "On-Load"

Memory efficient communication between the NIC and CPU cores.



## Secure Inter-Job Communication

Every network access specifies a uniquely numbered "communication type". Communication types can map to any application specific function.



## Distributed Application Model

Intelligent networking can provide distributed systems with the traditional application abstraction model.

