

2018年度(平成30年度)版

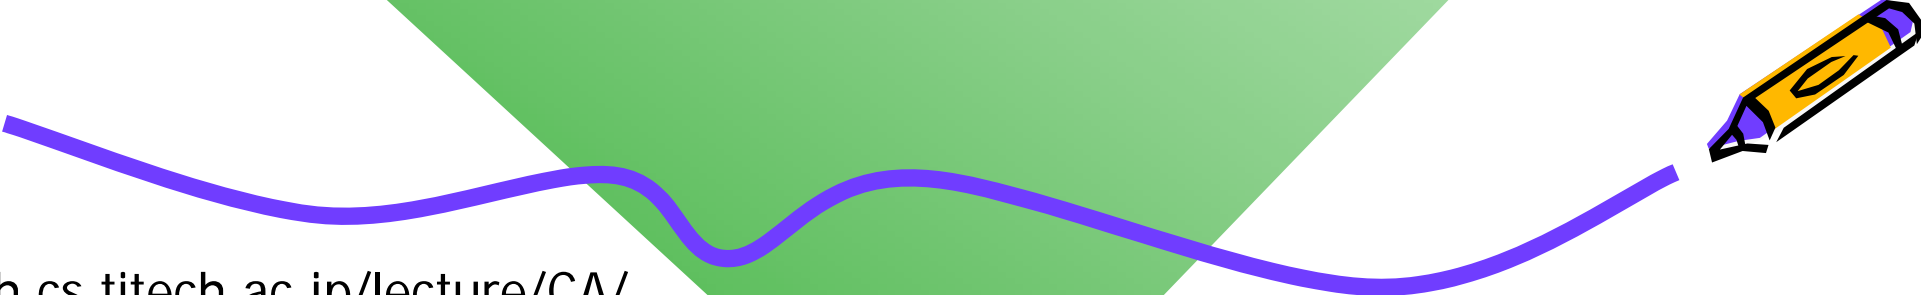
Ver. 2018-11-09a

Course number: CSC.T363



# コンピュータアーキテクチャ Computer Architecture

## 12. 相互接続ネットワークとNoC Interconnection Network and NoC

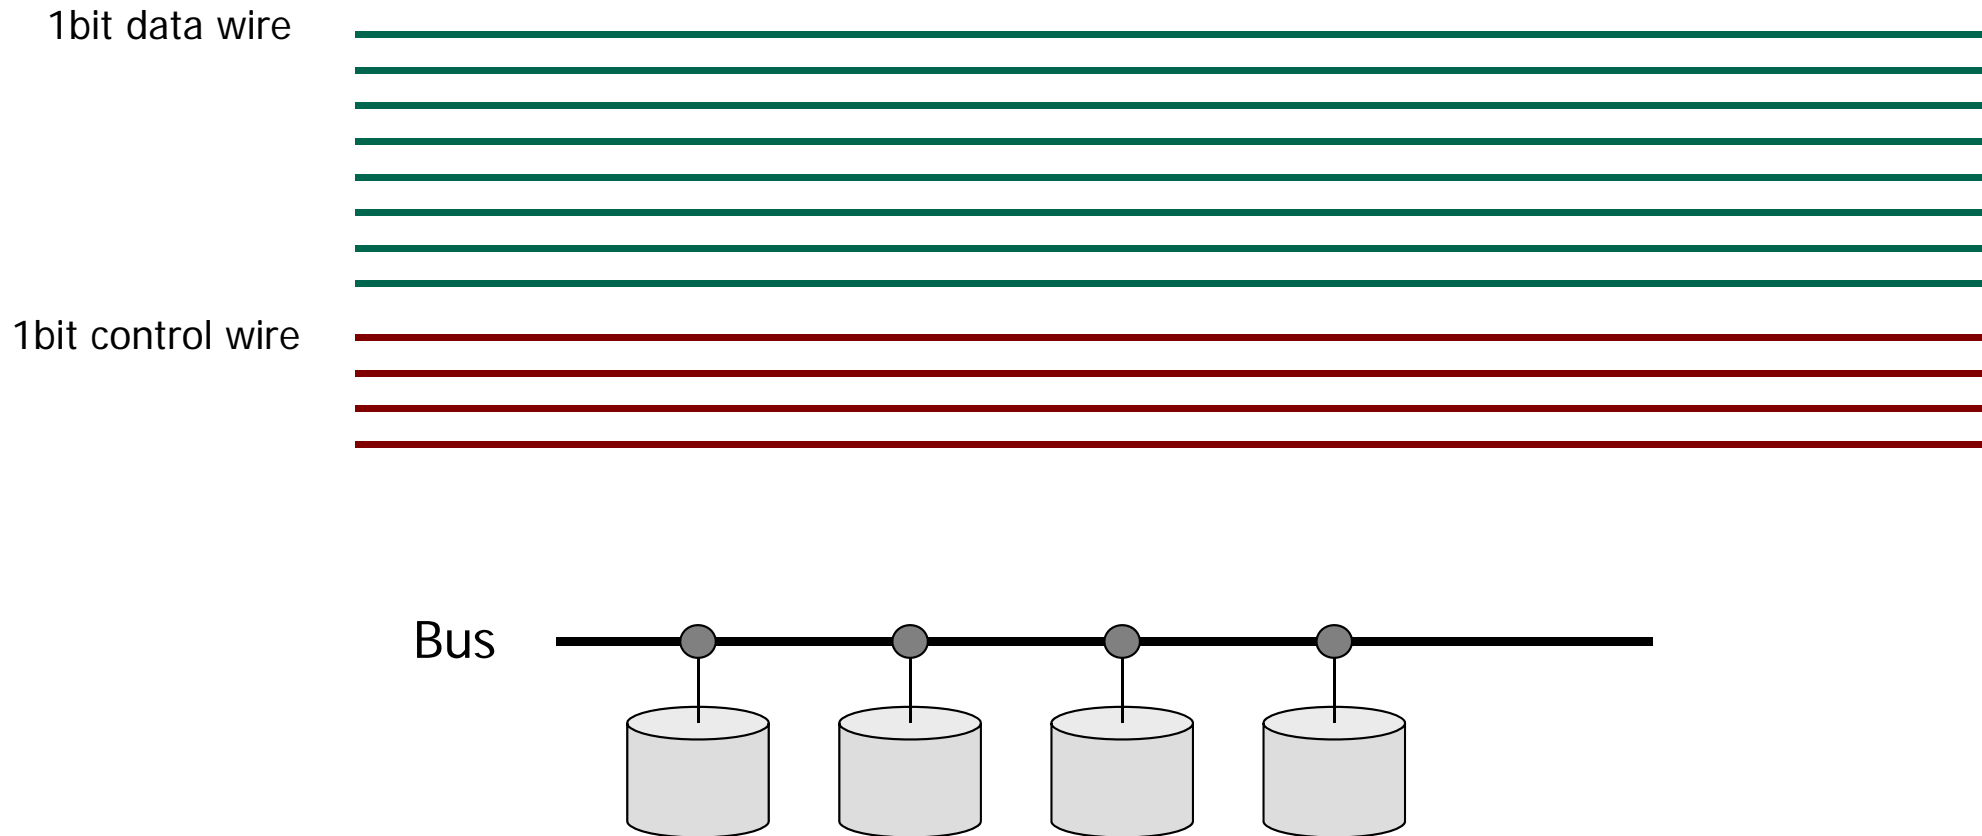


[www.arch.cs.titech.ac.jp/lecture/CA/](http://www.arch.cs.titech.ac.jp/lecture/CA/)  
Room No.W321  
Tue 13:20-16:20, Fri 13:20-14:50

吉瀬 謙二 情報工学系  
Kenji Kise, Department of Computer Science  
[kise\\_at\\_c.titech.ac.jp](mailto:kise_at_c.titech.ac.jp)


# Bus, I/O System Interconnect

- A **bus** is a **shared communication link**



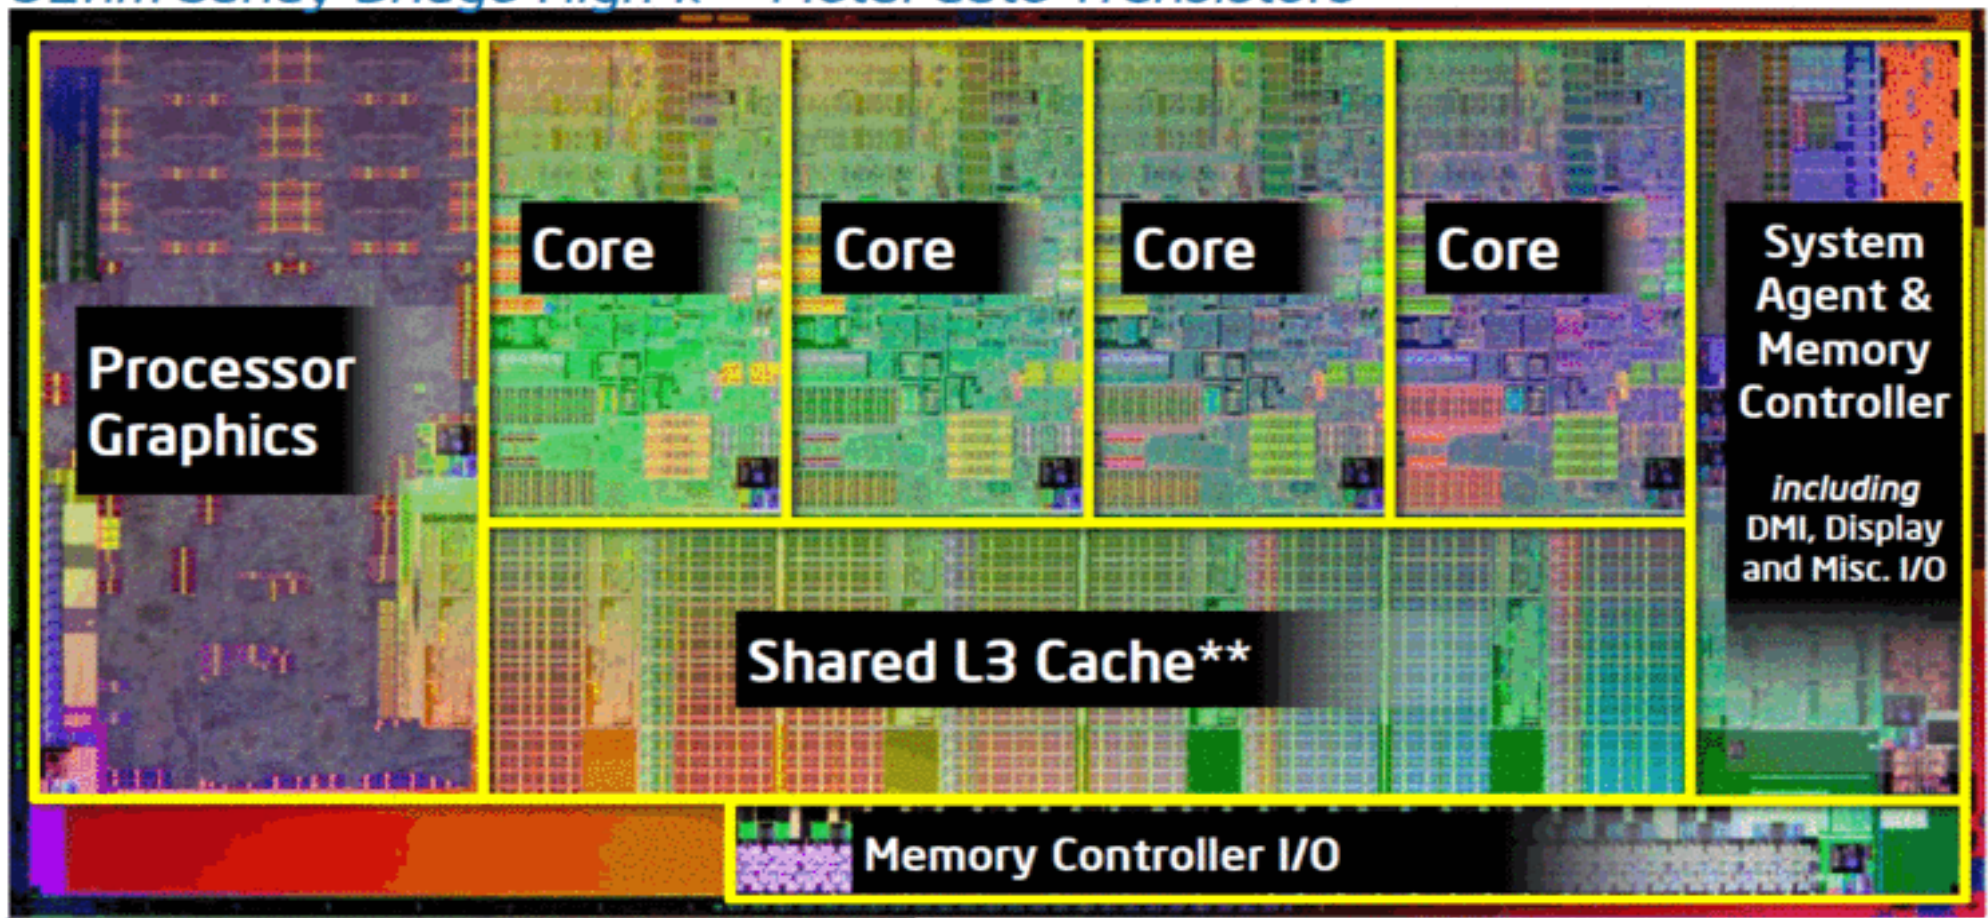
# Bus, I/O System Interconnect

- A **bus** is a **shared communication link** (a single set of wires used to connect multiple subsystems)
  - Advantages
    - Low cost – a single set of wires is shared in multiple ways
    - Versatile (多目的) – new devices can be added easily and can be moved between computer systems that use the same **bus standard**
  - Disadvantages
    - Creates a communication bottleneck – bus bandwidth limits the maximum I/O throughput
- The maximum bus speed is largely limited by
  - The length of the bus
  - The number of devices on the bus



# Intel Sandy Bridge, January 2011

- 4 to 8 core



# From multi-core era to many-core era




Figure 1: Current and expected eras of Intel® processor architectures

Platform 2015: Intel® Processor and Platform Evolution for the Next Decade, 2005

# Performance Metrics of Interconnection Network


- Network cost
  - number of switches
  - number of links on a switch to connect to the network (plus one link to connect to the processor)
  - width in bits per link, length of link
- Network bandwidth (NB)
  - represents the best case
  - bandwidth of each link \* number of links
- Bisection bandwidth (BB)
  - represents the worst case
  - divide the machine in two parts, each with half the nodes and sum the bandwidth of the links that cross the dividing line






# Bus Network

- N processors, 1 switch (●), 1 link (the bus)
- Only 1 simultaneous transfer at a time
  - NB (best case) = link (bus) bandwidth \* 1
  - BB (worst case) = link (bus) bandwidth \* 1



# Ring Network

- N processors, N switches, 2 links/switch, N links
- N simultaneous transfers
  - NB (best case) = link bandwidth \* N
  - BB (worst case) = link bandwidth \* 2
- If a link is as fast as a bus, the ring is only twice as fast as a bus in the worst case, but is N times faster in the best case





# Cell Broadband Engine (2005)

- Cell Broadband Engine (2005)
  - 8 core (SPE) + 1 core (PPE)
    - each SPE has 256KB memory
  - PS3, IBM Roadrunner(12k)



PlayStation3 の写真は  
PlayStation.com (Japan) から




Figure 2. Element interconnect bus (EIB).

IEEE Micro, Cell Multiprocessor Communication Network: Built for Speed




Diagram created by IBM to promote the CBEP, ©2005 from WIKIPEDIA

# Intel Xeon Phi (2012)



Intel® Xeon Phi™ Coprocessor Block Diagram

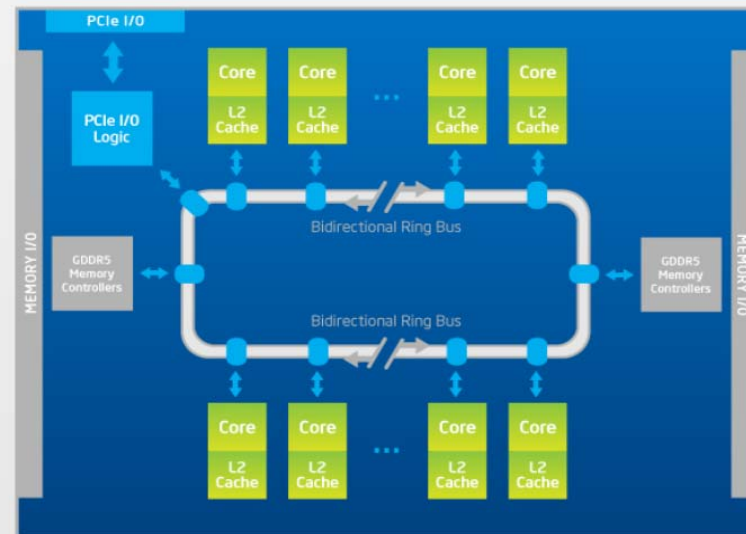

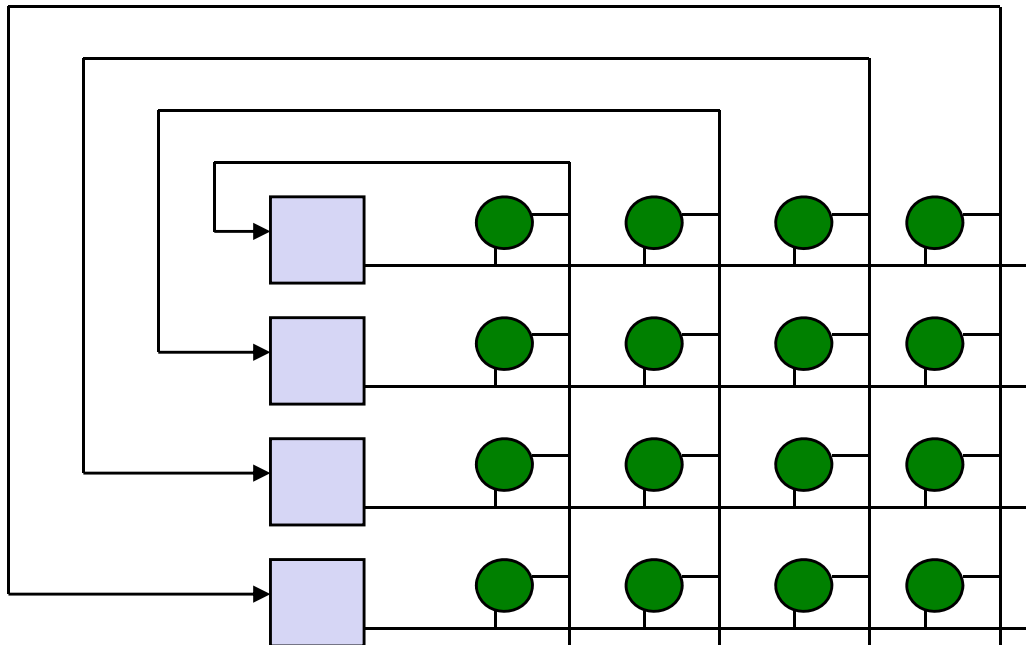


Table 2. Intel® Xeon Phi™ Product Family Specifications

PRODUCT NUMBER	FORM FACTOR & THERMAL SOLUTION*	BOARD TDP (WATTS)	NUMBER OF CORES	FREQUENCY (GHz)	PEAK DOUBLE PRECISION PERFORMANCE (GFLOP)	PEAK MEMORY BANDWIDTH (GB/s)	MEMORY CAPACITY (GB)	INTEL® TURBO BOOST TECHNOLOGY
3120P	PCIe, Passive	300	57	1.1	1003	240	6	N/A
3120A	PCIe, Active	300	57	1.1	1003	240	6	N/A
5110P	PCIe, Passive	225	60	1.053	1011	320	8	N/A
5120D	Dense form factor, None	245	60	1.053	1011	352	8	N/A
7110P	PCIe, Passive	300	61	1.238	1208	352	16	Peak turbo frequency: 1.33 GHz
7120X	PCIe, None	300	61	1.238	1208	352	16	

# Crossbar (Xbar) Network

- $N$  processors,  $N^2$  switches (unidirectional), 2 links/switch,  $N^2$  links
- $N$  simultaneous transfers
  - $NB = \text{link bandwidth} * N$
  - $BB = \text{link bandwidth} * N/2$




A symbol of Xbar



# Fat Tree (1)

- Trees are good structures. People in CS use them all the time. **Suppose we wanted to make a tree network.**
- Any time A wants to send to C, it ties up the upper links, so that B can't send to D.
  - The bisection bandwidth on a tree is horrible - 1 link, at all times
- The solution is to '**thicken**' the upper links.
  - More links as the tree gets thicker increases the bisection bandwidth




$N = 4$




# Fat Tree

- $N$  processors,  $\log(N-1) * \log N$  switches, 2 up + 4 down = 6 links/switch,  $N * \log N$  links
- $N$  simultaneous transfers
  - $NB = \text{link bandwidth} * N \log N$
  - $BB = \text{link bandwidth} * 4$



$N = 4$




$N = 8$




# Mesh Network

- $N$  processors,  $N$  switches, 4 links/switch,  $N * (N^{1/2} - 1)$  links
- $N$  simultaneous transfers
  - $NB = \text{link bandwidth} * 2N$
  - $BB = \text{link bandwidth} * N^{1/2}$




$N = 4$




$N = 16$




# 2D and 3D Mesh / Torus Network



2D Mesh



Torus




3D Mesh





# Intel Skylake-X, Core i9-7980XE, 2017

- 18 core




# Bus vs. Networks on Chip (NoC)






# NoC and Many-core


- NoC requirements: low latency, high throughput, low cost
  - Focus on mesh topology
- **Packet** based data transmission via **NoC routers** and XY-dimension order routing



PM: Processing Module or Core,  
R: Router



packet  
(tag + data)



# Packet organization (Flit encoding)

- A **flit** (flow control unit or flow control digit) is a link-level atomic piece that forms a network packet.
  - A packet has one head flit and some body flits.
- Assume that each flit has typical three fields:
  - payload(data) or route information(tag)
  - type : head, body, tail, etc.
  - virtual channel identifier



packet (tag + data)

Head flit

VC	Type		Route info
----	------	--	------------

Body flit

VC	Type	Payload
----	------	---------


Head and body flit formats

Head flit

Body flit


Body flit

Tail flit




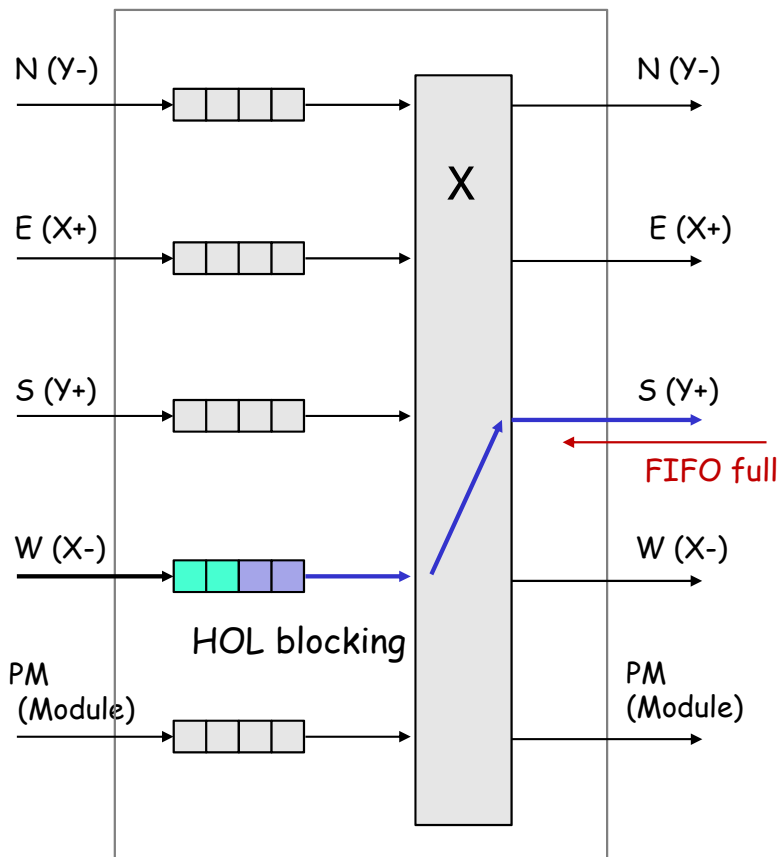
# Simple NoC router architecture

- Input buffer, routing, arbitration, flow control



# Datapath of Virtual Channel (VC) NoC router



- To mitigate **head-of-line (HOL) blocking**, virtual channels are used
- Pipelining



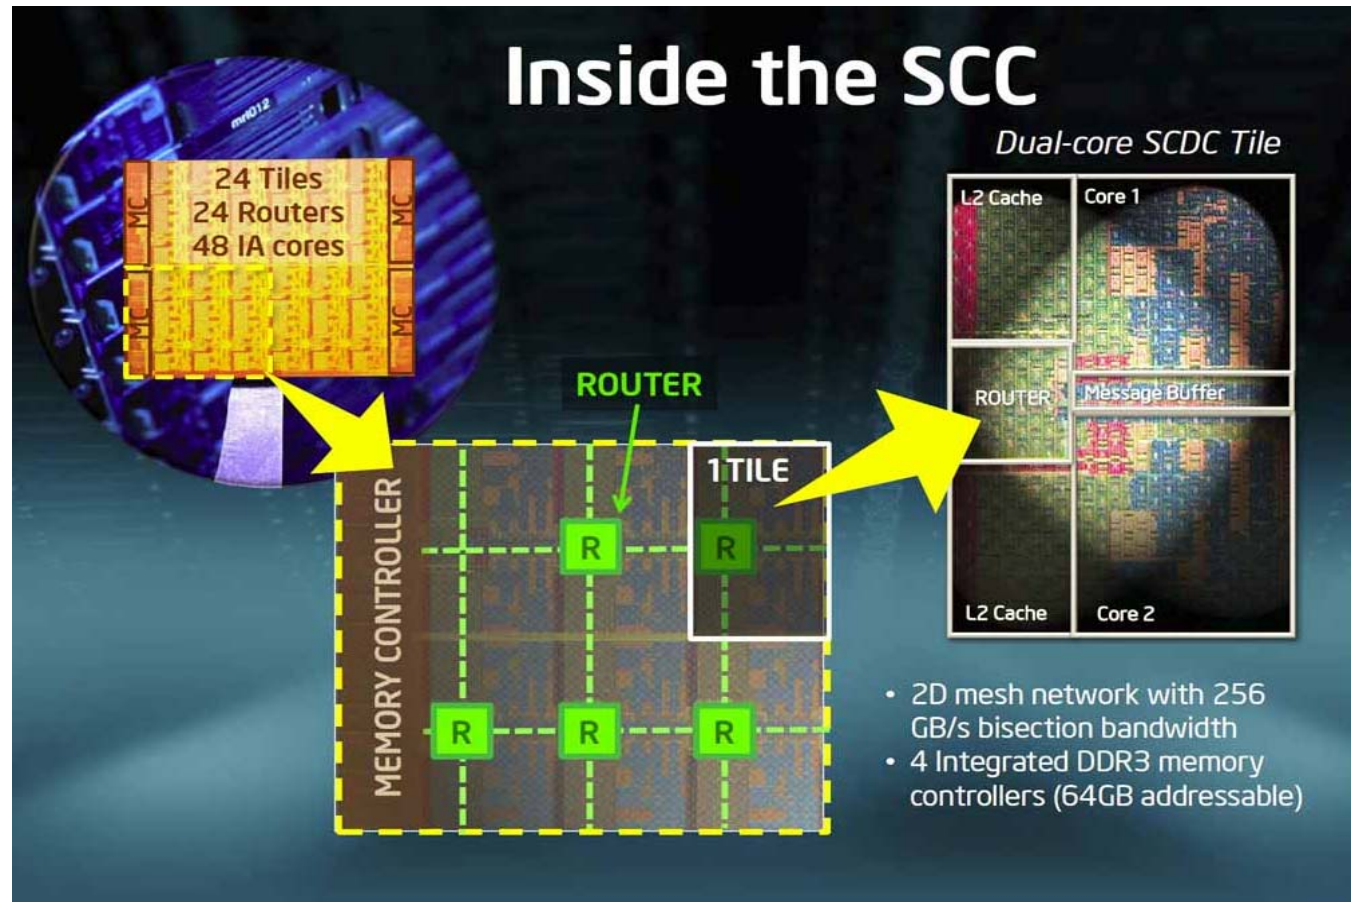
VC NoC router

# Intel Single-Chip Cloud Computer (2009)

- To research multi-core processors and parallel processing.



**A many-core architecture  
with 2D Mesh NoC**



**Intel Single-Chip Cloud Computer (48 Core)**





# Bus vs. Networks on Chip (NoC)

