Lecture 28: Board Notes: On-chip IC NWs

Consider the following "sea" of processor cores and routers

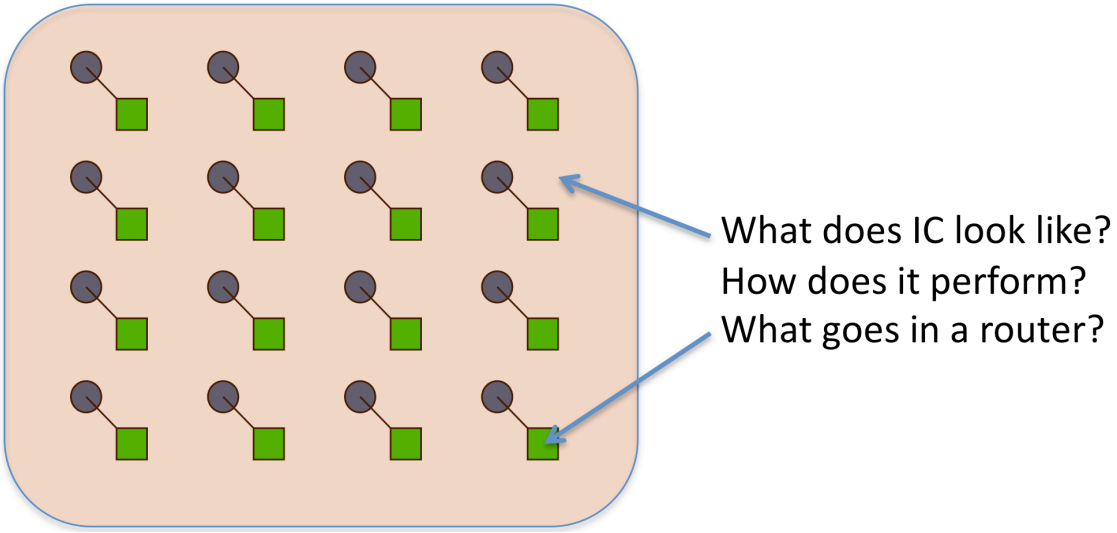

Let's look inside of a router first…

- Router has 2 main components:
  1. Datapath:
     - Handles storage and movement of a packet's payload
     - Consists of input buffers, switch, & output buffers
  2. Control
     - Logic to coordinate packet resource allocation

- I'm going to talk about a "Virtual Channel Router"
  - Virtual channel router requires extra resources (HW), but can help overcome blocking issues
    - (Might see blocking issues with wormhole routing)
    - (VC allows packets to pass a blocked packet and make better use of idle bandwidth)

Example:

1. Packet B enters node #1 from the network; B acquires channel *p* from node #1 → node #2
2. A 2nd packet A has entered node #1 from the wst and needs to be routed east to node #3
3. Meanwhile, B wants to leave node #2 and go south, but is blocked
4. Now channels p and q are idle .. but cannot be used
   a. Packet A is blocked in node #1
   b. It cannot acquire channel *p*
   c. B blocks

See figure:

Now, assume 2 VCs per physical channel:

1. B arrives at node #1 and acquires the bandwidth to go to channel *p*
2. A arrives from the east, B tries to leave node #2 and is blocked
3. A can use free bandwidth *p* and goto another VC on node #2
4. Can also proceed onto node #3

This is a better use of resources

- May have 1 physical channel, but more buffers

What happens during packet routing?

1. Let's start with a flit of a packet arriving at the input unit of a router
   - Input unit consists of a flit buffers to hold arriving flits until they can be forwarded
   - Input unit also maintains state of virtual channel
     i. I: Idle
     ii. R: Routing
     iii. V: Waiting for virtual channel
     iv. A: Active
   - Once packet in router, heed to perform route computation to see where it goes; can then go to VC for allocation
2. Each head flit must advance through 4 stages of routing computation
   - It's pipelined! Assume…
     - RC: Routing Computation
     - VA: Virtual Channel Allocation
     - SA: Switch Allocation
     - ST: Switch Traversal
   - Packet might move through like this:

	1	2	3	4	5	6	7
Head Flit	RC	VA	SA	ST			
Body Flit 1		**		SA	ST		
Body Flit 2					SA	ST	
Tail Flit						SA	ST

   - ** (second body flit arrives)

Important Points:

- t_r (time through a single router) does not equal 1!
  - (more like 5 or 6 at least)
- Routing and VC allocation are per packet functions
  - Nothing for body flits to do
  - With no stalls, need 3 input buffers (for 3 flits)
  - With stalls, need # of buffers = # of packets

Outlook:
- Ultimately, issues involved in routing process discussed above + router architecture + storage needed determine the bandwidth for the topology
    - Possibilities:
        - Even though you can devise a topology for ideal performance, it may not be feasible to implement
        - Or, 1 part may be technologically feasible (pitch) but another may not be (router or buffer)

Why can routers be hard to implement?

Consider the following picture:

Now, consider how connections would actually be made on chip:
- **Discuss metal stack**
- **Show cross-sectional die photo**
- **Draw lines for input and output**

Now, let's go back to our picture and made some observations:
1. No lines of the same color can touch (it would be an electrical short)
2. We draw 1 line, but really many (1 line for each bit)
3. Router areas are by no means insignificant!

How can on-chip IC NWs affect performance?

Want to know – for a given IC NW topology – how long it takes to send a message:
- Note → initial #s in the *absence* of contention → a bit more on this later

Time: (# of hops) x (time in router) +
time required for packet to traverse *all* channels +
serialization latency

(serialization latency = ceiling(length of message / bandwidth))

Therefore, if:
- Average # of hops = 6.25
- Average time for packet to traverse all channels = 5.3333
- Serialization latency = 3
- Time in router = 2
- Total time: = ~20.8

We'll conclude by looking at some real performance projections