Interconnect IP for Network-on-Chip

J. Liu, M. G. Shen, L-R Zheng and H. Tenhunen
Labratory of Electronic and Computer Systems
Royal Institute of Technology (KTH)
Electrum 229, SE 164 40 Kista/Stockholm Sweden
{jianliu,mgshen,lrzheng,hannu}@imit.kth.se

Presentation Outline

- Why Network-on-Chip (NoC)
- NoC architecture
- Interconnect constraints in NoC
- Bandwidth optimizaiton in NoC
- Interconnect Intellectual Property (IIP)
- Conclusions

Trends and Challenges

- Communication vs. computation
- Deep submicron effects
- Global synchrony
- Design productivity gap
- Heterogeneity of functions

NoC Architecture

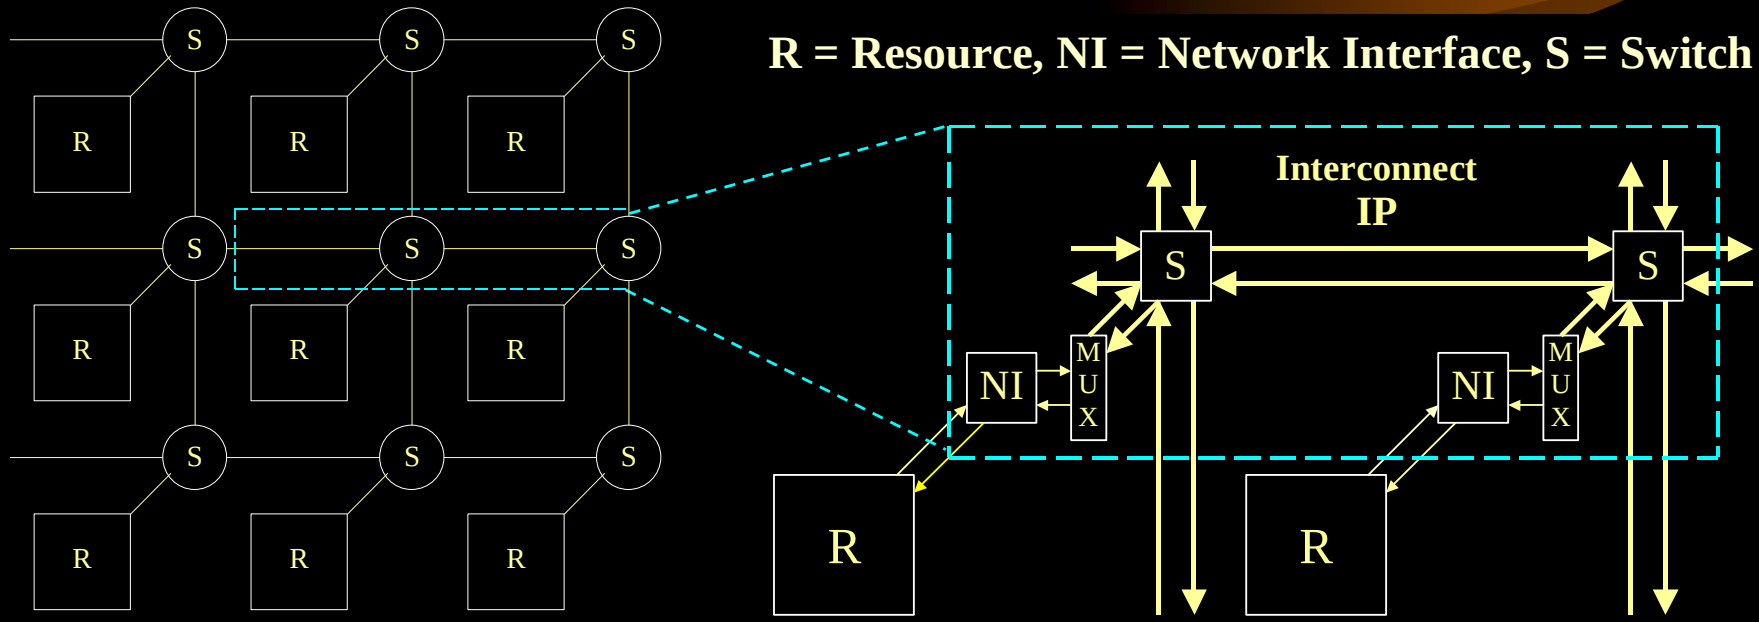

Simple 2D mesh, other topologies possible

Resource: processor, memory, IP core, FPGA...
Inter-resource communicate using data packets.

Advantages

- Reuse
  - Components and resources
  - Communication platform
  - Design and simulation environment
  - Verification
- Predictability
  - Communication performance
  - Electrical properties

Interconnect Constriants

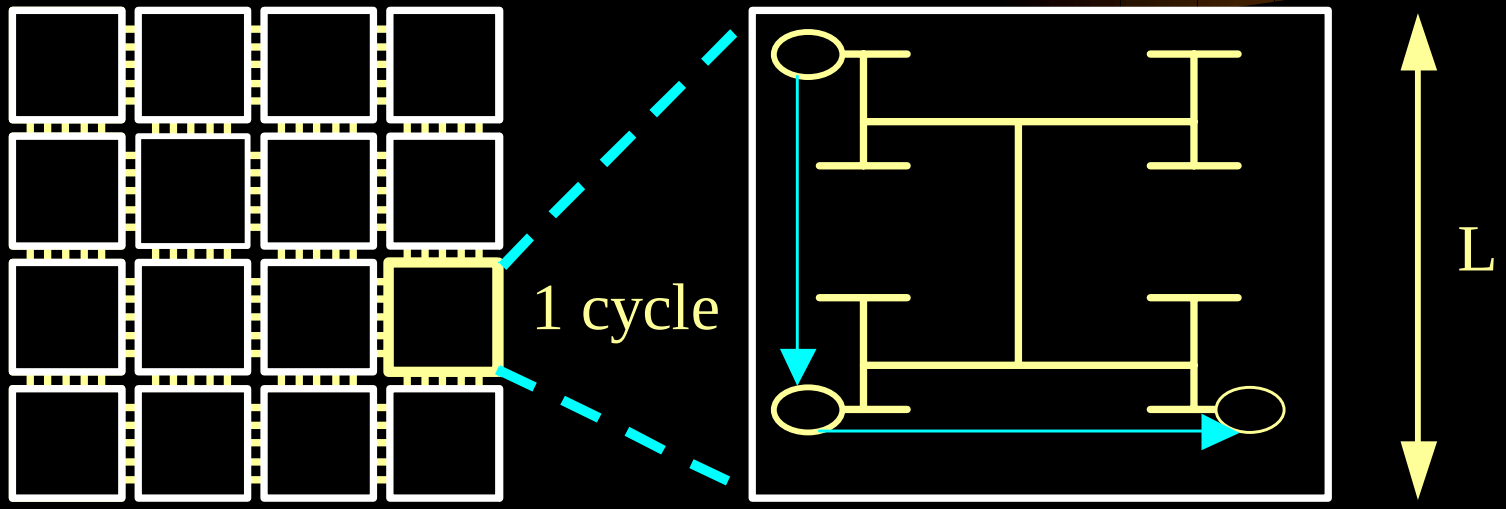

GALS = globally asynchronous, locally synchronous

Cycle time > wire delay of length 2L

RC delay for a single wire

Without repeaters

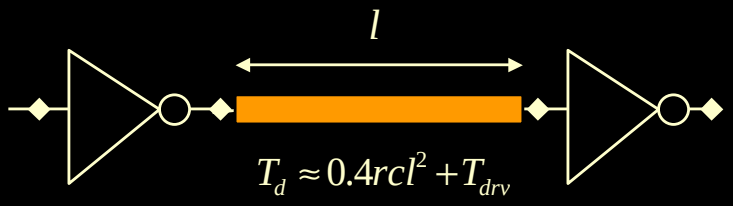

$$T_d = 0.377R_sC_s + 0.693(R_dC_{out} + R_dC_s + R_sC_{out})$$

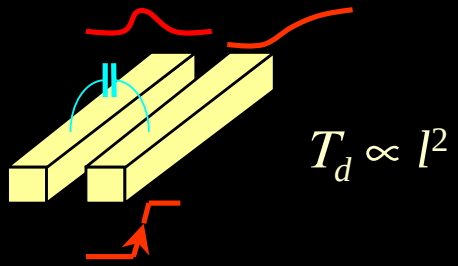

$$T_d = 2.13\sqrt{rcFO1} \Rightarrow 4.26L\sqrt{rcFO1} < 1\ clock\_cycle$$

With repeaters

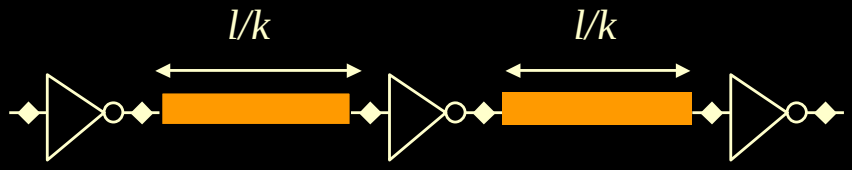

$$T_{total} = k \cdot (T_{drv} + 0.4 \cdot rc(l/k)^2)$$

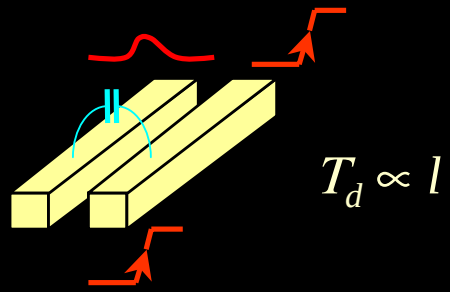

WC FO1≅167*L_{gate} ps
(R. Ho, IEEE Proc. vol.89 no.4, 2001)

Maximum synchronous resource size

Wire Type	Parameter	0.18-µm	0.13-µm	0.10-µm	0.07-µm	0.05-µm
Semi-Global	R (ohm/mm)	107	185	317	611	1196
	c (fF/mm)	331	268	208	170	155

Wire parameters

	0.18-µm	).13-µn	).10-µn	).07-µn	.05-µ
Cost Perf. (GHz)	0.56	0.77	1.0	1.4	2.0
High Perf. (GHz)	1.1	1.5	2.0	2.9	4.0

Clock frequency

$$T_d = 2.13\sqrt{rcFO1} \Rightarrow 4.26L\sqrt{rcFO1} < 1\ clock\_cycle$$

	Technology	.18-µm	.13-µm	.10-µm	.07-µm	.05-µm
	Chip Size (mm)	20	21	23	25	28
High Performance	ax Resource Size	6.5	4.7	3.5	2.4	1.5
	Nr of Resources	9	20	42	112	350
Cost Performance	ax Resource Size	13	9.3	7.1	4.7	3.0
	Nr of Resources	2	5	10	28	87

Optimal repeater insertion with crosstalk

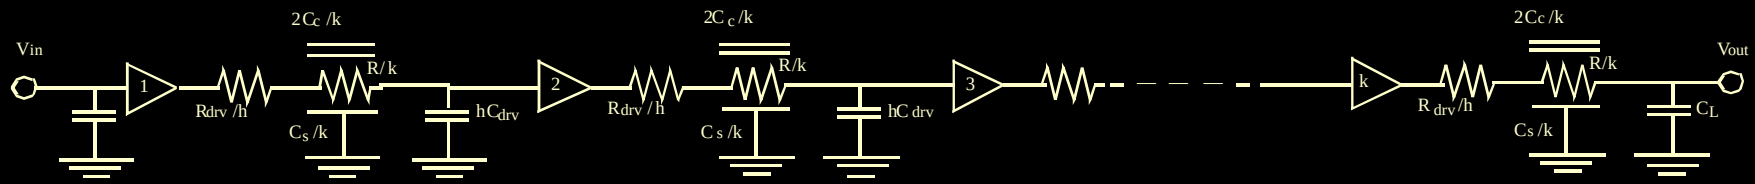

$$t_{0.5} = k\left[0.7\frac{R_{drv}}{h}\left(\frac{C_s}{k} + hC_{drv} + 4.4\frac{C_c}{k}\right) + \frac{R}{k}\left(0.4\frac{C_s}{k} + 1.5\frac{C_c}{k} + 0.7hC_{drv}\right)\right]$$

C_c-- crosstalk capacitance, k- number of repeaters, h- repeater size

$$\frac{\partial t_{0.5}}{\partial k} = 0 \Rightarrow k_{opt} = \sqrt{\frac{0.4RC_s + 1.5RC_c}{0.7R_{drv}C_{drv}}}$$

$$\frac{\partial t_{0.5}}{\partial h} = 0 \Rightarrow h_{opt} = \sqrt{\frac{0.7R_{drv}C_s + 3.1R_{drv}C_c}{0.7RC_{drv}}}$$

(D. Pamunuwa & H. Tenhunen, *Proc. of 14th Int. Conference on VLSI Design*, January 2001)

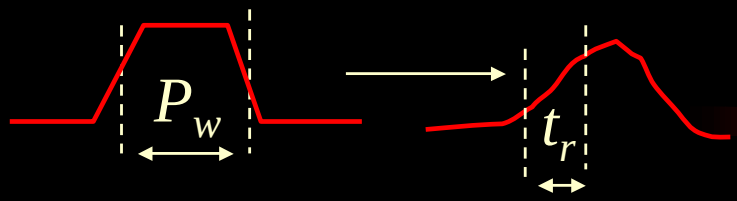

$P_w > 3t_r$

$t_{r(0\text{-}90\%)} = 3.18 t_{d(0\text{-}50\%)}$

Limit: Inter-symbol interference

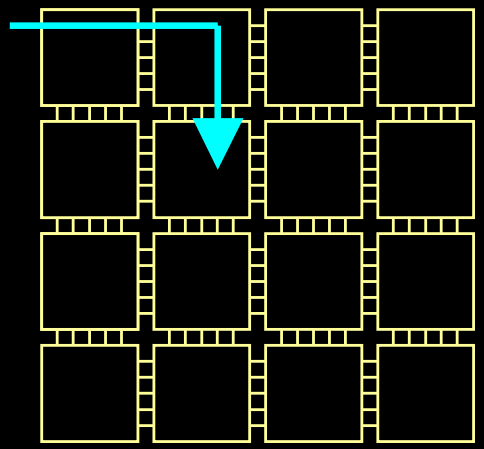

(a) no refresh; (b) refreshed;
@ same total signal delay,
(a) 1.3mm (b) 3 mm.

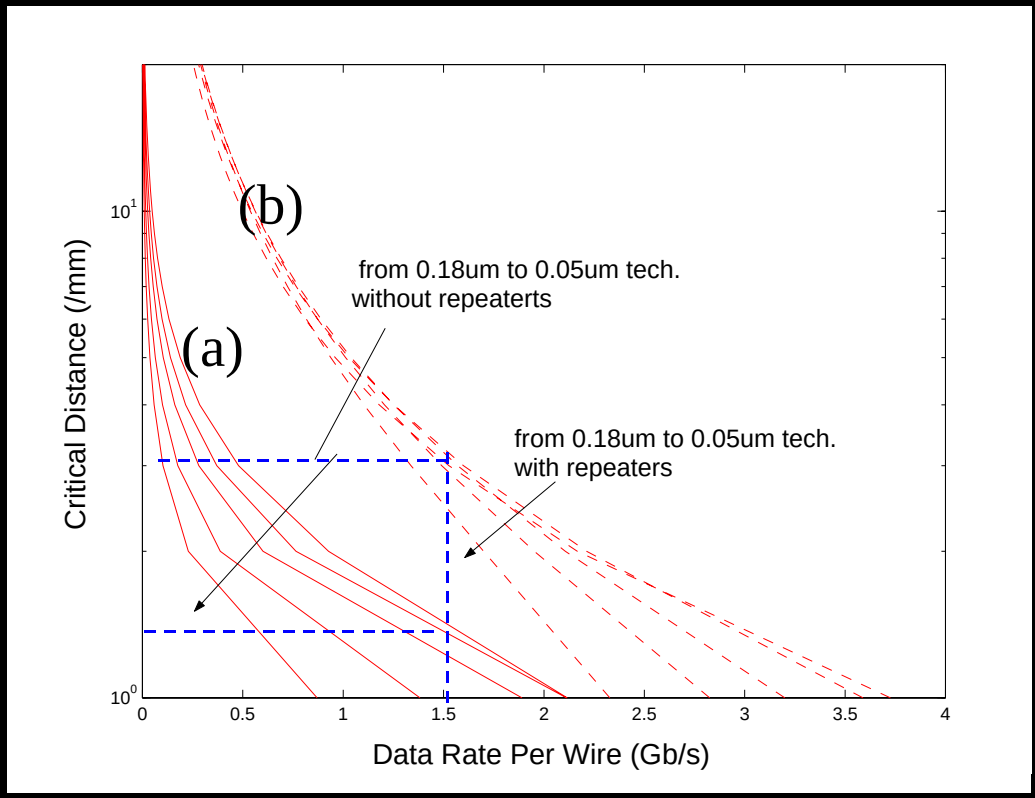

Inter-resource bandwidth optimization

- PROBLEM: Given a fixed metal area for communication link: what is the optimal configuration?

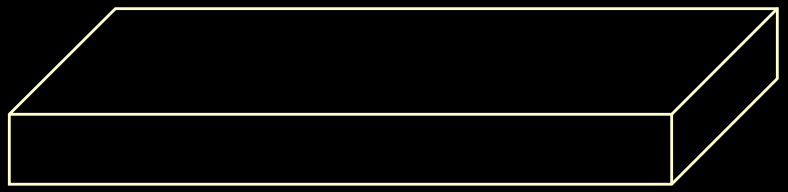

1 huge wire?

Signal speed close to speed of light, but only one wire and that's it!

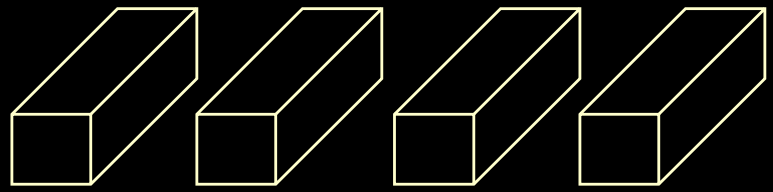

A few fat wires?

Now cross-talk comes into play!

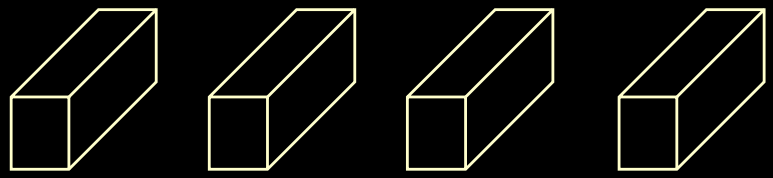

Make them thinner & Space them further apart?

Now wire resistance is higher!

Large number to benefit from parallelism?

Now wire resistance and coupling capacitance are both much higher!

How does repeater insertion affect all this?

Fixed width W_T

$$W_T = NW + (N-1)S$$
$$(= 15\ \mu m)$$

D. Pamunuwa, ISCAS 2002

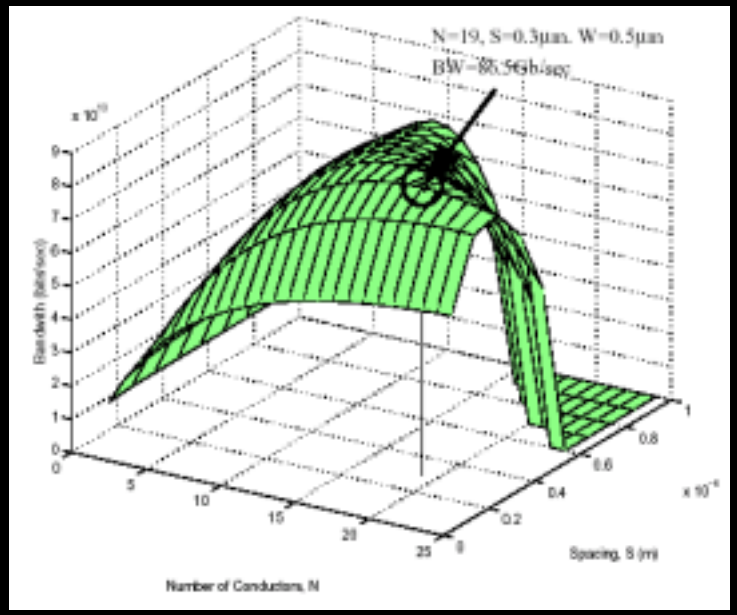

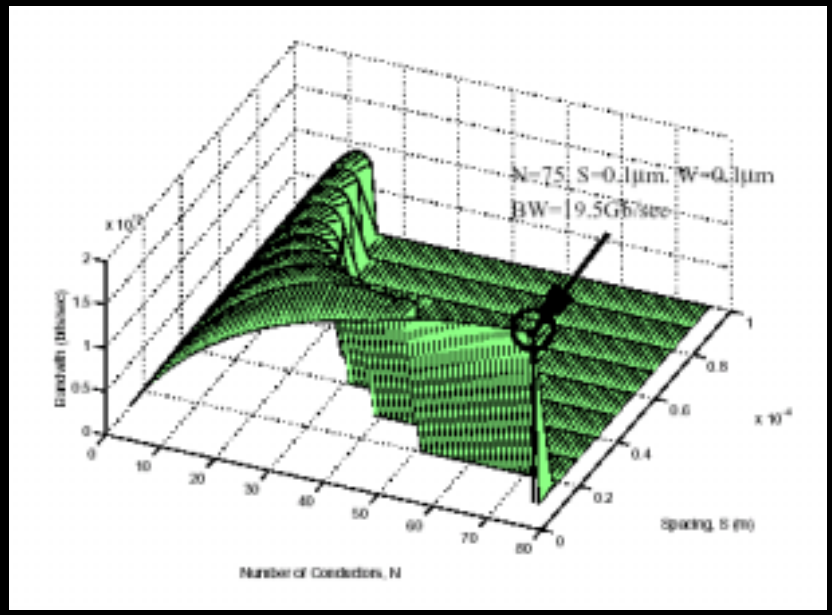

Without buffers

With optimal buffering

Components in interconnect IP

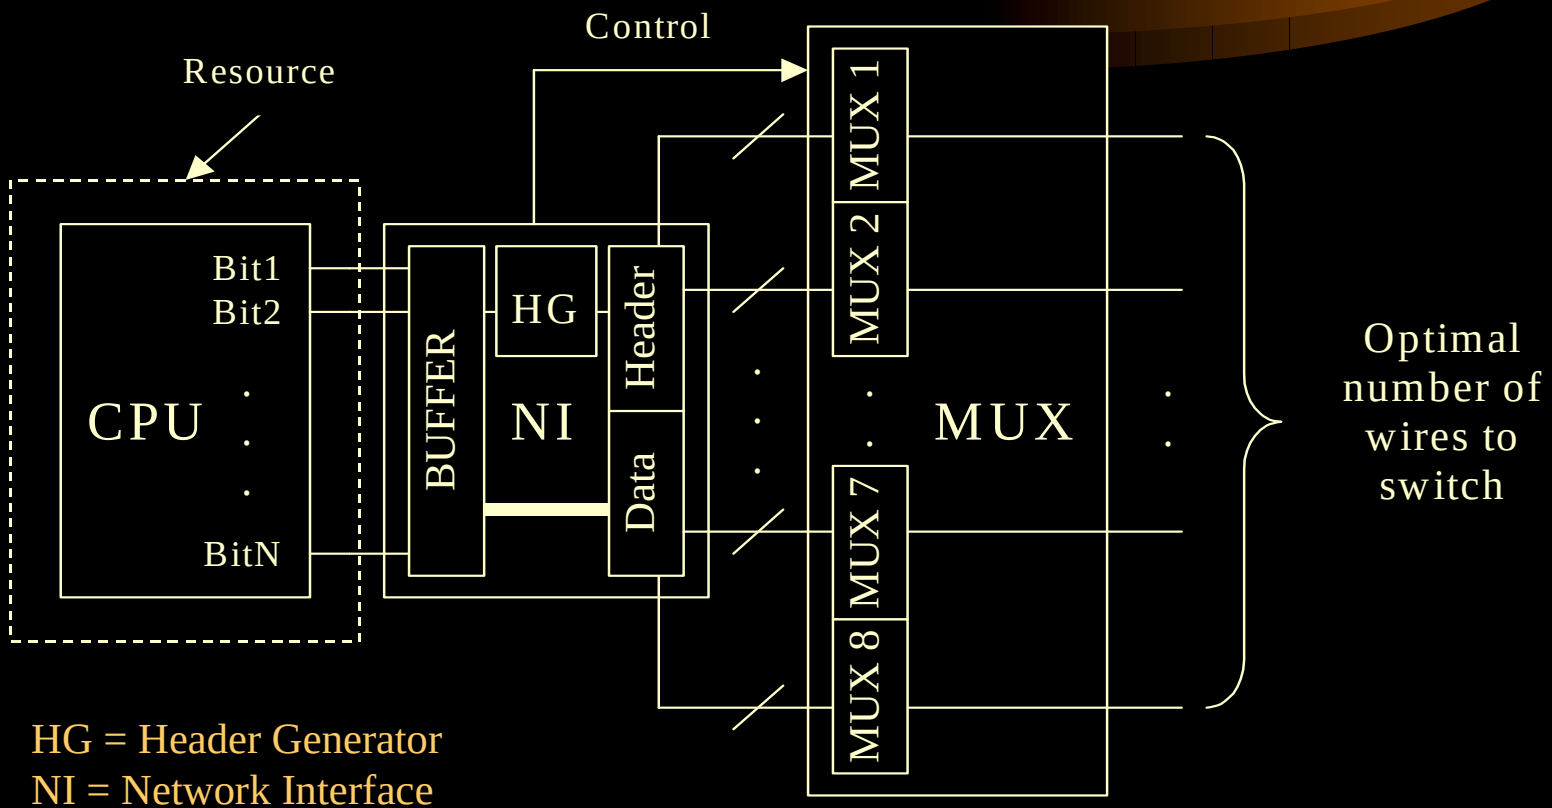

HG = Header Generator
NI = Network Interface

Network taxonomy

NoC network core similar to traditional telecom network

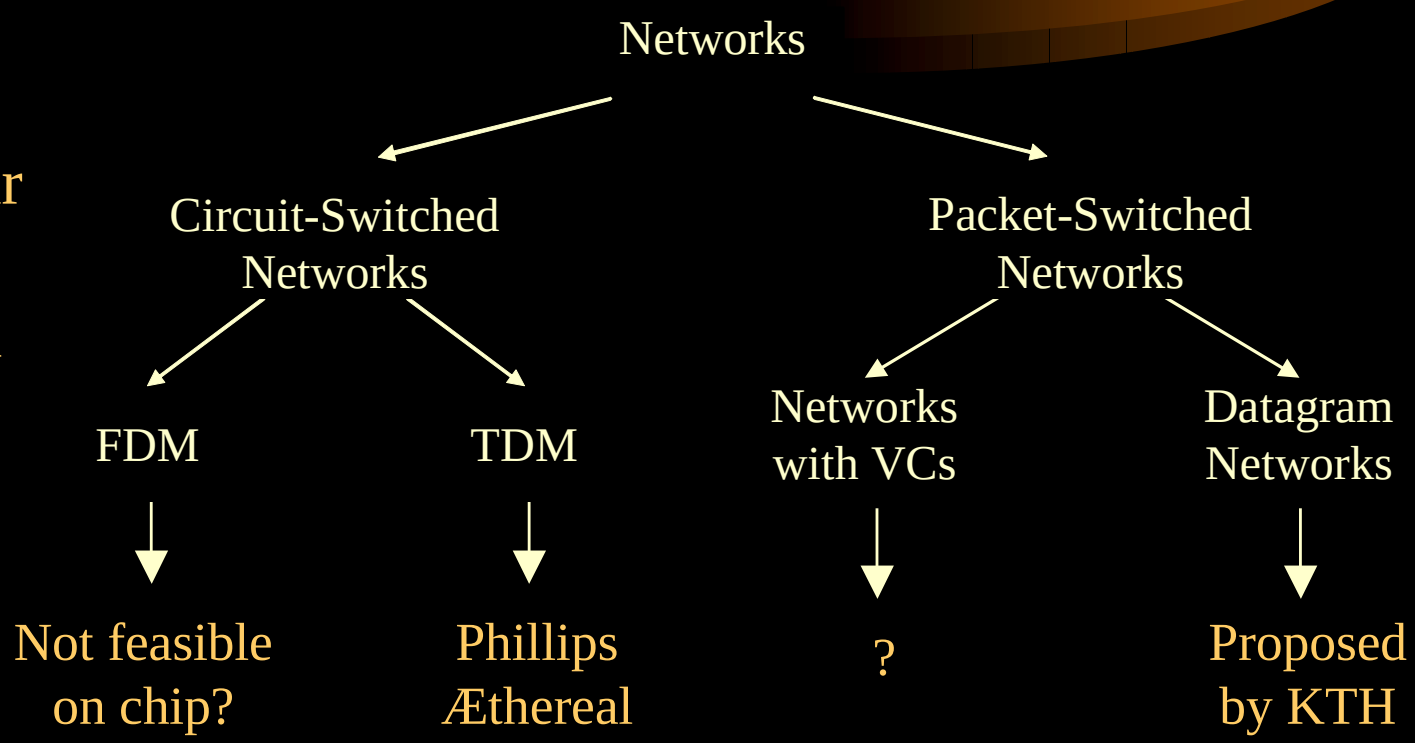

Switch for datagram network

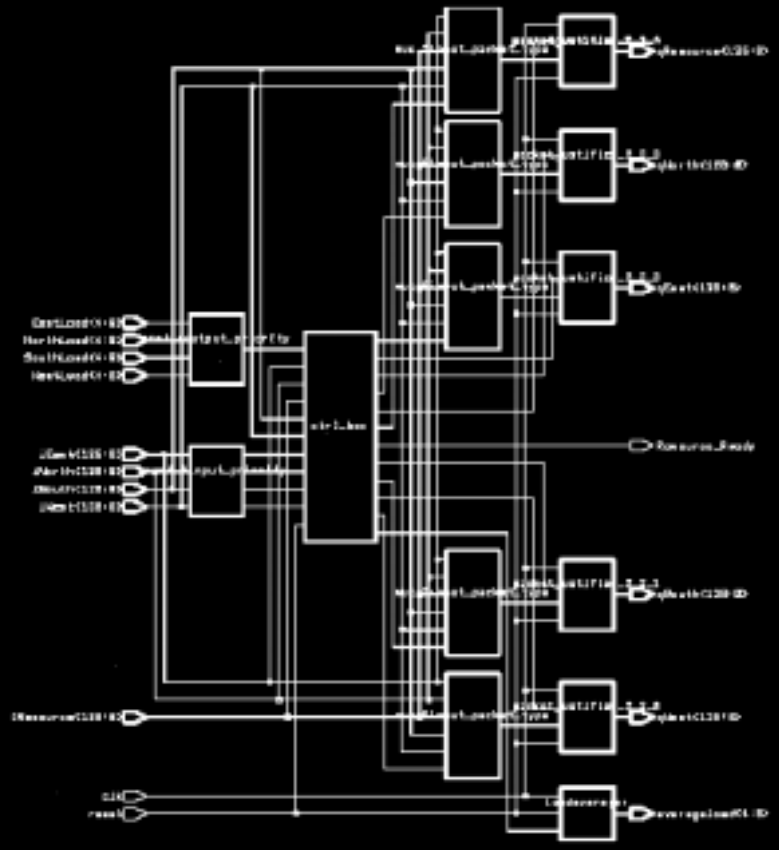

Buffer-less switch

Synthesis with Synopsys

constraint	*total combined logic*	*critical path gate depth*
optimised for area	13 964	79
optimised for speed	21 029	48

E. Nilsson, Design and Implementation of a hot-potato Switch in a Network on Chip, Master thesis, KTH.

Why (not) datagram network

- Advantages
  - State- and memoryless switch gives simple design
  - Easily adapts to changes in the network
- Disadvantages
  - Delay variation
  - Out-of-order delivery

Switch for TDM network

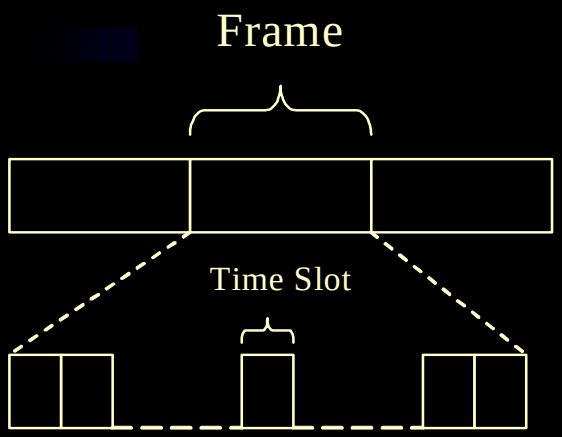

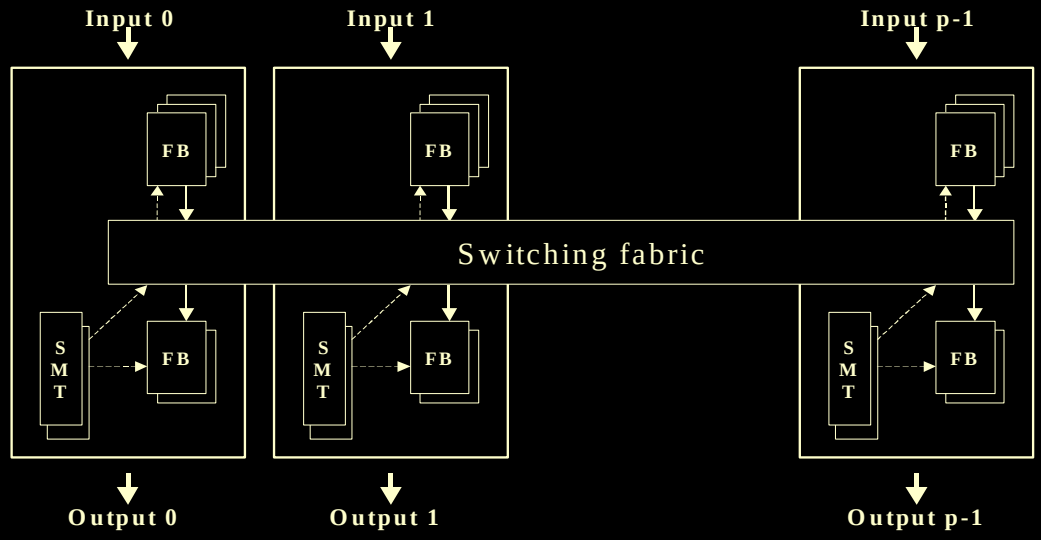

- An end-to-end “channel” is needed, compare to bus control.
- Three RAM frame buffers for each input
- Two RAM frame buffers for each output
- Two slot map tables (SMTs) for each input/output pair
- Fully connected switching fabric

Phase compensation scheme

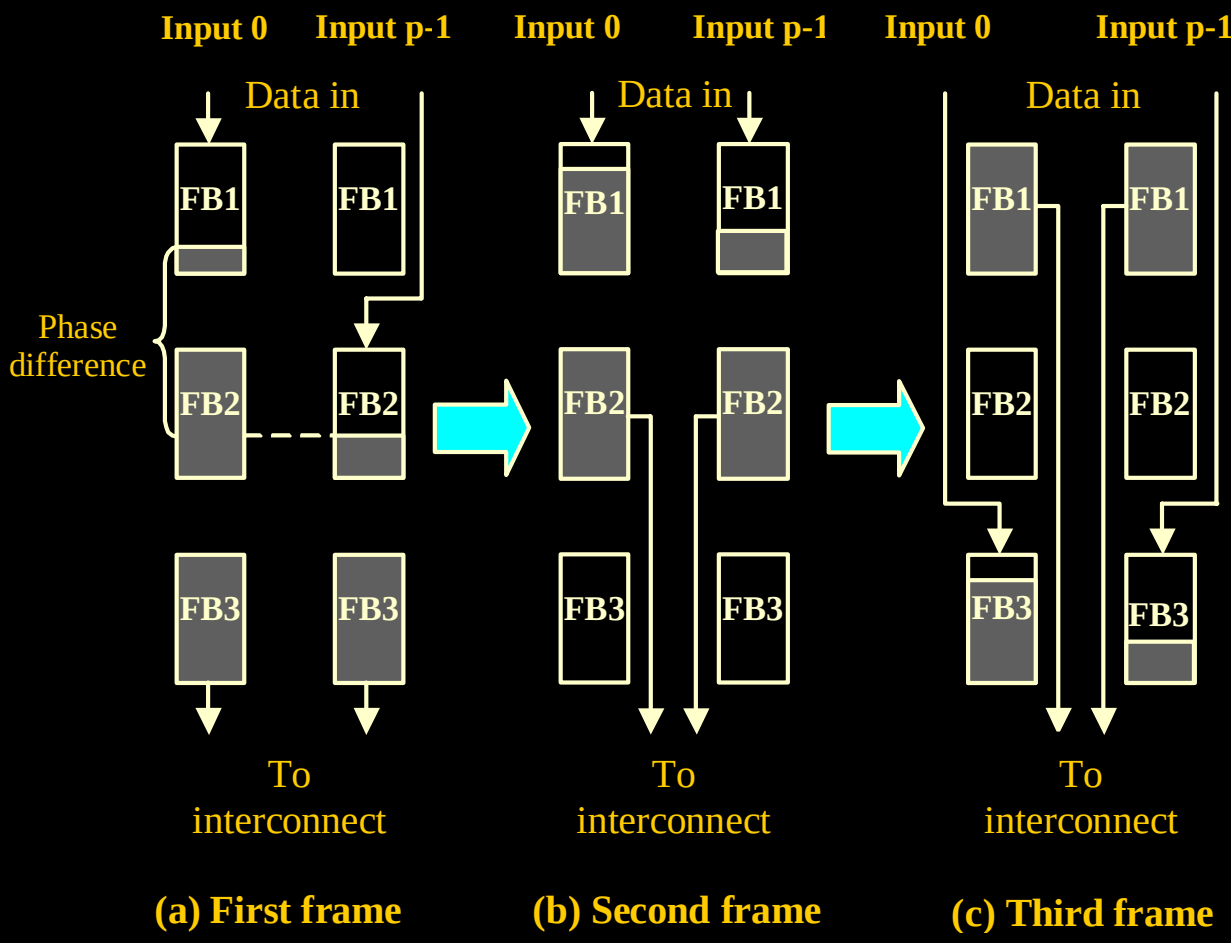

Equivalent time-space-time (TST) switch

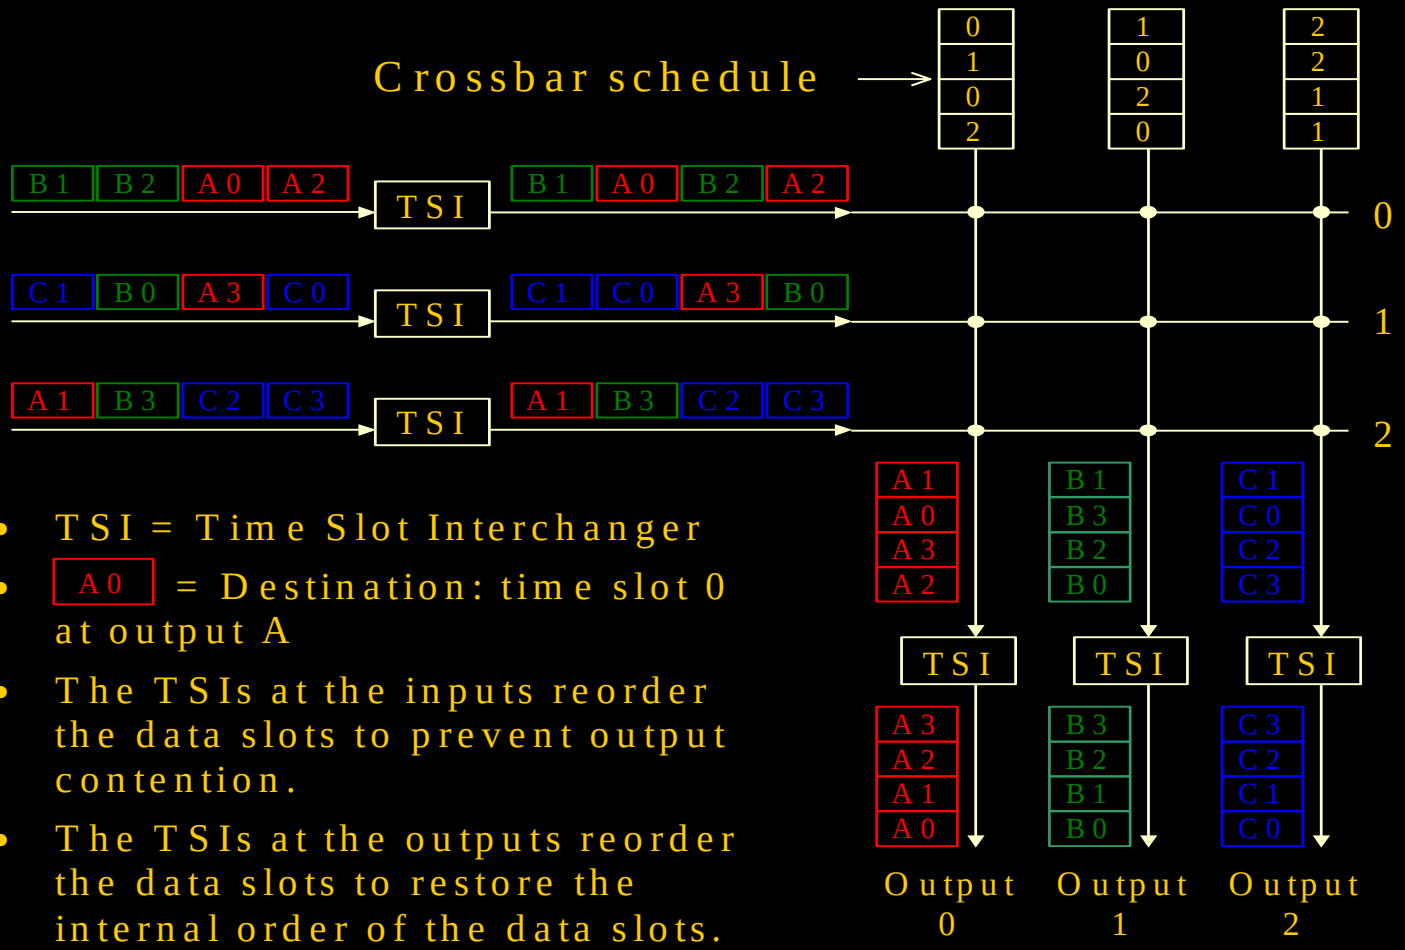

- TSI = Time Slot Interchanger
- A0 = Destination: time slot 0 at output A
- The TSIs at the inputs reorder the data slots to prevent output contention.
- The TSIs at the outputs reorder the data slots to restore the internal order of the data slots.

Slot Map Table (SMP)

- Has same function as the TSIs and crossbar schedule together in the TST switch.
- Controls the read/write operations using three fields: input buffer number (ib), input slot number (is) and output slot number (os).

ib	is	os
0	0	2
1	1	3
0	1	0
2	3	1

(a) SMT for output buffer 0

ib	is	os
1	2	0
0	2	2
2	2	3
0	3	1

(b) SMT for output buffer 1

ib	is	os
2	0	3
2	1	2
1	0	0
1	3	1

(c) SMT for output buffer 2

SMT properties

- A row in the SMTs corresponds to a time period (cycle) long enough to perform a read/write operation pair
- Two different outputs are not allowed to access different memory locations of an input in same cycle (assume one read port)
- Due to possible multicasts, it may require more time cycles than the number of time slots in a frame

Problem with multicast

Red entries may not be in same row since different outputs read from same input.

Green and blue entries may not be in same row since they are reads from diffent memory locations of same input

ib	is	os
0	0	0
0	1	1
1	0	2

ib	is	os
0	0	0
0	1	1
1	1	2

ib	is	os
0	0	0
0	1	1
1	2	2

ib	is	os
1	0	2
~~0~~	~~0~~	~~0~~
0	1	1

(a) SMT foroutput buffer 0

ib	is	os
0	0	0
1	1	2
0	1	1

(b) SMT for output buffer 1

ib	is	os
0	0	0
~~0~~	~~1~~	~~1~~
1	2	2

(c) SMT for output buffer 2

Tasks

- How many rows do the SMTs need?
- Effective scheduling algorithms? Here, scheduling means arbitrary column permutations of the SMTs. It’s OK to permute a SMT since the RAM buffers can be read in any order.

SMT length

- Let the number of inputs and outputs be *p*
- Let the number of time slots in a frame be *n*
- The required SMT length is *n* for unicast
- Best known algorithm is O(*nplogn*)
- Run-time rescheduling of an entry O(*p*)

Multicast scheduling

- The required SMT length is $\min\{np, n^2\}$
- Finding the optimal schedule is NP-hard
- Approximative greedy algorithms exit
- SMT length = $qn + p^{1/q}(n-1)$, here q is integer

  (Teofilo F. Gonzalez, "Multimessage Multicasting: Complexity and Approximations", 1996)
- For 2D mesh, $p = 5$. Multiport RAM keeps SMT length = n. No scheduling needed.

Conclusions

- From computation to communication and interconnect centric design
- NoC platform uses interconnect IPs to provide well-controled global wire delay and efficient global commnication
- Wire planning and repeater insertion maximizes the inter-resource bandwidth
- Different network cores give different switch architecture