

DAT575: Interconnection Networks

Exam - January 2025

Date: January 13th, 2025

Time: **8:30-12:30**

Examiner: Ioannis Sourdis

Department: Computer Science and Engineering

Inquiries: contact through phone, phone extension 1744

Duration: 4 hours

Grading scale: 100 points in total

0: 0%-49%, **3:** 50%-64%, **4:** 65%-84%, **5:** 85%-100%

Available references:

- two A4 pages (e.g. one A4 paper used in both sides) of Lecture notes (handwritten and/or computer produced) are allowed.
- Chalmers approved calculators are allowed.
- Textbooks are NOT allowed.

General: Submit your solutions, in English, on blank paper sheets. Write legibly; feel free to use figures to get your point across.

The order of answering the questions does not matter (start with the easiest ones).

Please start the solutions for each problem on a new sheet. Please number the sheets so that the solutions are in numerical order.

Note that it is possible to receive partial credit for an answer even if it is not 100% correct.

Your personal identity code is required on each submitted sheet!

Good luck!

Note: These are example answers to the exam questions for one possible set of randomized variables.

Question 1 Topologies: (15 points)

Consider a 4x4 2D mesh network.

- Draw the topology and find the bisection bandwidth of the network considering that all channels have equal bandwidth $b = 1\text{Gbps}$.
- Calculate the maximum channel load ($\gamma_{\max}$) and the ideal throughput (Θ_{ideal}) of the network (per injection port) for (i) uniform random and (ii) bit-complement traffic pattern.

Consider also that each injection and ejection port have bandwidth “b”.

Answer

- $B_b = 8 \cdot b$ (because of channels in both directions) so $B_b = 8\text{ Gbps}$

UR:

The $N=16$ injection ports inject $16 \cdot b$ traffic if injection rate is 100%. in UR 50% of this traffic ($8 \cdot b$) goes (evenly) through the 8 bisection channels which support a bandwidth of $8 \cdot b$, and that is sufficient to allow each injection port to inject maximum traffic, i.e. $b = 1\text{Gbps}$.

Or

$$\gamma_{\max} = N/2 \cdot B_c = 1, \Theta_{\text{ideal}} = b / \gamma_{\max} = 1\text{ Gbps (per injection port)}$$

Bit complement:

For bit complement 100% of the traffic ($16 \cdot b$) goes through the bisection which supports however only $8 \cdot b$ bandwidth. That is 50% of the maximum traffic that could be injected. So each injection port will be constraint to 50% of its maximum injection rate, that is $b/2 = 0.5\text{ Gbps}$.

Or

$$\gamma_{\max} = N/B_c = 2, \Theta_{\text{ideal}} = b / \gamma_{\max} = 0.5\text{ Gbps (per injection port)}$$

Question 2 Routing: (10 points)

- Explain how the choice of a routing algorithm can (i) improve network throughput, and/or (ii) reduce network latency.
- Describe a good routing algorithm for a torus that can achieve well both (i) and (ii).

Answer

- Throughput is improved if routing algorithm offers path diversity
- Latency is improved if routing algorithm offers the minimal (or close to minimal) path.
- An algorithm with good balance is Channel Queue Routing (CQR) – lecture on routing slide 58. (other justified algorithms can be considered correct too)

Question 3 Flow Control: (15 points)

- Give an example of head of line blocking (HoL).
- What type of flow control suffers from HoL and why?
- Show how the HoL blocking problem of your example can be solved/alleviated?

Answer

- Lecture on flow control, slide 40
- Buffered, Flit-based (wormhole) because flits of different packets may share the same buffer.
- Virtual channels. Lecture on flow control, slide 43

Question 4 Deadlocks: (10 points)

- Explain the necessary conditions for a network to have the risk of deadlock
- which of the following networks have a probability to deadlock and why?
 1. 2D mesh with dimension order routing
 2. A butterfly network
 3. A torus network with minimal adaptive routing
 4. A torus network with fully adaptive routing.

Consider that all networks use wormhole routing and a single channel (there are no parallel virtual channels)

- Suggest a way to avoid deadlocks for the above networks that may suffer from deadlocks.

Answer

- Lecture on deadlocks, slide 2
1. No deadlocks because routing is constrained and has no circles
 2. No deadlocks because topology has no circles
 3. Yes because there are circles and require an extra VC0 to avoid deadlocks (slide 55)
 4. Yes, because there are circles and require an extra VC0 and additional constraints to prevent a packet to transition from VC0 back to VC1 or otherwise enforce some channel order (slide 61).

Question 5 Router Microarchitecture (15 points)

Consider a network composed of routers which the following characteristics:

- Lookahead routing (Next Route Computation – NRC),
- Speculative switch allocation
- Switch traversal and Link traversal performed in separate stages, one clock cycle each.

Calculate the average zero-load latency of a 5-flit packet that traverses 3 hops in the network, considering that speculation has a success rate of 2/3.

Answer

$$T = T_h + T_s$$

$$T_h = (3 + 1 \cdot 1/3) \cdot \text{hops} = 10$$

$$T_s = 5 - 1 = 4$$

$$T = 14$$

Question 6 Arbiters & Allocators: (15 points)


Consider the 4x4 matrix arbiter depicted bellow

With the following configuration:

X	0	0	1
1	X	0	1
1	1	X	1
0	0	0	X

Consider that all four requesters have one active request.

Show the order of the requests been granted and the state of the matrix after each grant.



Answer

Sequence of grants: R2, R1, R0, R3

X	0	0	1
1	X	0	1
1	1	X	1
0	0	0	X

R2, R1, R0, R3

X	0	1	1
1	X	1	1
0	0	X	0
0	0	1	X

R1, R0, R3, R2

X	1	1	1
0	X	0	0
0	1	X	0
0	1	1	X

R0, R3, R2, R1

X	0	0	0
1	X	0	0
1	1	X	0
1	1	1	X

R3, R2, R1, R0,

X	0	0	1
1	X	0	1
1	1	X	1
0	0	0	X

Question 7 Reliability: (5 points)

Explain the tradeoff of detecting transient faults in transmitted packets (and resending faulty packets) when it is performed hop-by-hop versus when it is performed end-to-end.

Answer

Hop-by-hop

- (+) Shorter retransmission delay
- (-) More delay, area, and energy cost to detect at every hop
- (-) More area cost (buffer and logic) to support retransmission at every hop


End-to-end:

- (-) longer retransmission delay
- (-) larger buffers at source node to keep pending packets for longer.
- (-) requires network to support end-to-end acknowledgements
- (+) lower delay, area, energy costs at the intermediate routers (no need for detection and retransmission there)

Question 8 Quality of Service: (5 points)

Calculate the fair throughput allocated to flow-A, flow-B, and flow-C according to the Max-Min algorithm, considering that:

- the incoming bandwidth of flow A is 0.3,
- the incoming bandwidth of flow B is 0.5,
- the incoming bandwidth of flow C is 0.6, and
- the total bandwidth of the share channel is 0.75.



Answer

0.25, 0.25, 0.25

Question 9 Network Interfaces: (5 points)

- Which are the two main programming paradigms that can be used by a multi-processor system? How does each paradigm generate network traffic?
- How does the choice of programming paradigm affect in general the design of a network interface? Describe the general principles a network interface needs to support in each case (no need to describe the details a network interface)?


Answer

Message passing: processes running in processors send explicit messages to each other. Network traffic is generated by the explicit messages exchanged by the processors. The network interconnects only processors with each other so the network interface is symmetric (processor interface) and supports transforming the processor messages to network packets. This is either control by software or performed by a coprocessor.

Shared memory. Processors issue load and store (memory) requests. The network traffic is then implicit, as each load and store request generates messages between processor and memory. The network interconnects processor and memory modules, so the network interface is asymmetric, i.e. different network interfaces are required to connect a processor or memory. Status holding registers are often used to keep track of pending requests. When a cache coherent protocol is supported then a large number of message types need to be supported. E.g. for exclusive read requests and cache line invalidations.

Question 10 Busses: (5 points)

Consider write and a read requests in a bus pass through the stages shown in the following figure, i.e., write: (a), read: (b):




where:

- AR: Access Request,
- ARB: Arbitration,
- AG: Access Granted,
- RQ: Read/Write Request,
- P: Processing,
- RPLY: Reply
- ACK: Acknowledgement

Show how the following requests will be scheduled in a pipelined bus:

ReadRQ1, WriteRQ1, WriteRQ2, if they get a grant in the above order.

How many cycles did it take to complete all three requests and for how many of these cycles was the bus utilized?

Answer

10 cycles out of which bus is used in 6 cycles (60% bus utilization).