

Logistics

- Overdue for class representatives. Please we are looking for nominations!
- Next lecture on 08/03, 4-6PM at the Anthony Low Bldg 124 Innovations Theatre (replaces 11/03 lecture, Canberra day)
- Assignment 1 will be released on 11/03 (Mon, next week)

Overview: Routing and Communication Costs

- routing mechanisms and communication costs
- routing techniques: store-and-forward (SF), cut-through (CT)
- communication algorithms and cost modelling. Special focus on broadcast:
 - in particular, two broadcast variants: one-all and all-all
 - impact of store-and-forward (SF) and cut-through (CT) on communication cost
 - we consider rings, meshes, torus, trees and hypercubes in the analysis
- improving the speed of (some) communication operations

References. Grama et. al.:

- Section 2.5 and 2.6
- Ch 4 (includes algorithms/cost analysis beyond one-all/all-all broadcast)

Routing Mechanisms

Efficient algorithms for routing messages through the network are essential for the performance of parallel systems

Given a source and destination node in the network, determine which path(s) a message takes through the network to reach its destination. Can be classified into:

- Minimal versus non-minimal routing. Minimal takes (one of the) shortest paths, each link brings you closer to destination. Non-minimal may route the message along a longer path to avoid network congestion.
- Deterministic versus adaptive routing. Deterministic determines a unique path solely based on the source and destination node. Adaptive routing uses information on the state of the network to determine the path of the message.

Example of a minimal deterministic routing is the so-called dimension-order routing in 2D meshes (e.g., XY-routing) and hypercubes (E-cube routing)

Communication Costs

Communication cost model of a point-to-point message parameterized by:

- start-up time (t_s): time required to handle a message at the sending node and at the receiver. This includes the time to prepare the message for transmission (adding header, trailer, error correction info, etc.), execute the routing algorithm, and the time to establish an interface between the local node and routing switch.
- per-hop time (t_h): time taken for the header of the message to travel between two directly-connected nodes. Related to the latency in the routing switch
- per-word transfer time (t_w): time taken by each word to traverse the link.

Store-and-Forward (SF) Routing Technique

- common on networks for early parallel systems
- each intermediate node in the path from the sending to receiving node forwards the message **ONLY AFTER** the entire message has been received and stored
- assume a message of size m traversing l links, point-to-point comm. cost is:

$$t_{\text{comm}} = t_s + l(t_h + mt_w)$$

- usually the per-hop time $t_h \ll mt_w$ (even for “small” m), and therefore:

$$t_{\text{comm}} \approx t_s + lmt_w$$

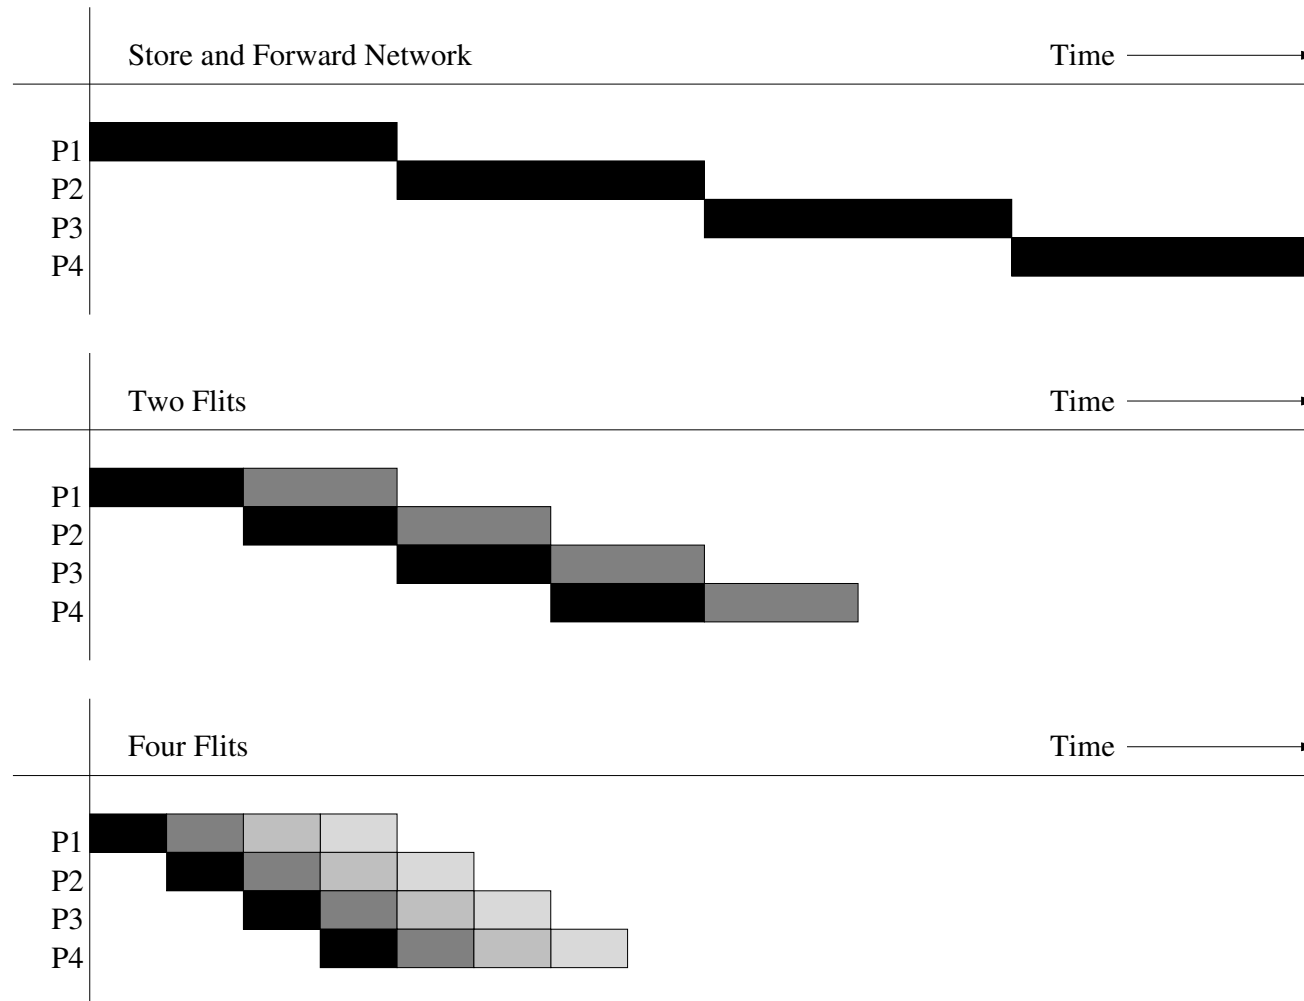
- poor utilization of communication resources as only one link is active at a time for the communication at hand

Cut-Through (CT) Routing Technique (aka wormhole routing)

- the CT term is more broadly used in long-haul networks (e.g. the Internet); for high-performance networks (this course) CT is aka wormhole routing
- message split into **small-sized** fixed-length segments called flow-control digits (flits)
- message header takes lt_h to arrive (acts as a tracer that sets up the connection)
- flits are forwarded ASAP at each intermediate node, in a pipelined fashion, increasing communication resources usage (all links asymptotically used at a time)
- communication cost can be (asymptotically) modelled as (why?):
$$t_{\text{comm}} = t_s + lt_h + mt_w$$
- compared to store-and-forward routing, l does not multiply mt_w . Besides, its implementation requires less memory at the intermediate nodes (why?)
- more sophisticated versions extend the concept to multi-lane CT routing (aka virtual channels) to handle more efficiently variable message sizes
- pioneered (mostly) by Bill Dally for the Torus Routing chip (1986), with milestone papers on deadlock avoidance (1987) and virtual channels (1992)

Pipelining of Flow Control Digits (aka flits)


node P1 sends a message to node P4 that passes through P2/P3 nodes
(shaded regions represent the time the message is in transit)




Communication Costs on Static Networks

- for a single point-to-point message transfer, communication costs can be modelled as (recap):
 - store-and-forward: $t_{\text{comm}} = t_s + l(mt_w + t_h) \approx t_s + lmt_w$
 - cut-through: $t_{\text{comm}} = t_s + lt_h + mt_w$
- how can we extend this to derive cost models for communication patterns?
- for simplicity, let us focus on broadcast communication pattern:

One-All Broadcast




All-All Broadcast



Store-and-Forward/One-All/Ring

- (from now on) let us assume that:
 - node 0 is the source of the broadcast
 - each node can only send a single message at a time
 - the network uses deterministic minimal routing (shortest path)
- two possible algorithms to perform one-all broadcast on a ring:



(left: naive; right: broadcast 2-way split)


- naive communication cost modelled as (assuming that p is an odd number):

$$t_{\text{one-all}} = (p - 1)t_s + \sum_{l=1}^{\lfloor \frac{p}{2} \rfloor} 2l(t_h + mt_w)$$

- broadcast 2-way split communication cost can be modelled as:

$$t_{\text{one-all}} = \lceil \frac{p}{2} \rceil (t_s + t_h + mt_w)$$

Cut-Through/One-All/Balanced Binary Tree



- processing nodes at the leaves, intermediate nodes are just switching units
- assume t_h between switches and nodes is equivalent to t_h between switches
- broadcast completes in $\log_2(p)$ stages; at each stage the number of parallel message transfers doubles (as with the hypercube)
- communication cost can be modelled as:

$$t_{\text{one-all}} = \log_2(p)(t_s + mt_w) + (\sum_{i=1}^{\log_2(p)} 2^i)t_h$$

Store-and-Forward/All-All/Torus

Three steps:

1. each process in an arbitrarily selected row performs an all-all broadcast within the row using the algorithm SF/All-All/Ring
 2. each node collects the \sqrt{p} messages they have received from the other nodes in that row, and consolidates them into a single message of length $m\sqrt{p}$
 3. each column performs an all-all broadcast within the column
- communication cost can be modelled as:

$$\begin{aligned} t_{\text{all-all}} &= \underbrace{(\sqrt{p} - 1)(t_s + t_h + mt_w)}_{\text{step 1}} + \underbrace{(\sqrt{p} - 1)(t_s + t_h + m\sqrt{p}t_w)}_{\text{step 3}} \\ &= 2(\sqrt{p} - 1)(t_s + t_h) + (p - 1)mt_w \end{aligned}$$

Improving the speed of (some) communication operations

- the algorithms and communication costs modelled so far assume that the message (or its **flips**) always follow the same path among the sender and receiver
- if we split messages into parts, and let each part follow different paths among the sender/receiver, we may sometimes improve network utilization (see Grama et. al, Ch 4.7.1 for examples)
- they also assume that a node cannot be transmitting to several neighbours at a time
- some parallel computers have hardware able to send messages in multiple directions at the same time (aka all-port communication) or even specialized networks for certain kind of communication primitives (e.g., IBM BG/L supercomputer, see Lin & Snyder, Ch 2, "Supercomputers")