Network Control and Signalling

Network Control and Signalling

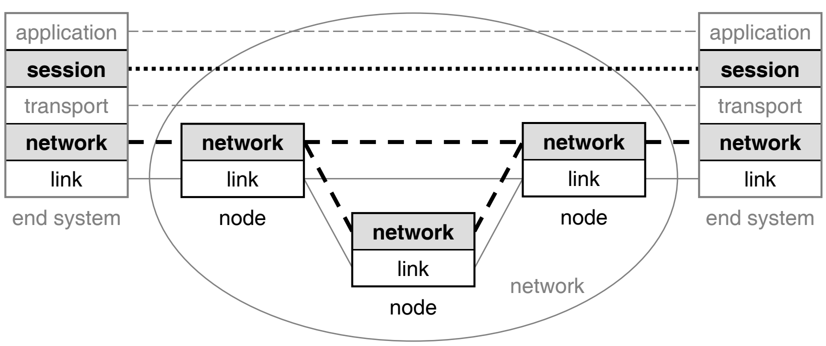

Network Control and Signalling

Signalling Paradigms

Network Control and Signalling © Sterbenz and Touch

Network Control and Signalling

Communication Flow Diagrams

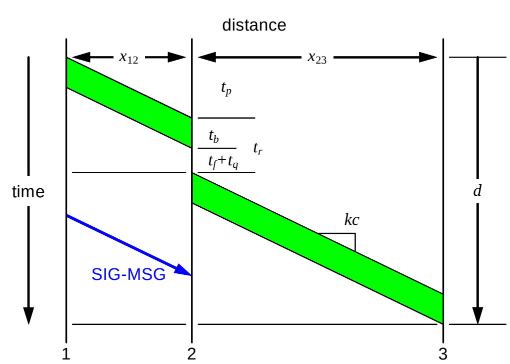

- Packets are parallelograms
- Messages are directed line segments

Signalling Paradigms

Circuit Switching

- First generation (<1970)
  - – PSTN switching technique
- Characteristics
  - – setup latency: RTT *before* data transfer
  - – no multiplexing efficiency
  - + negligible switch latency

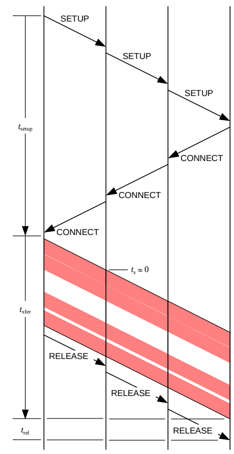

Signalling Paradigms

Message Switching

- First generation (<1970)
  - – data switching technique
- Characteristics
  - – no setup latency
  - + some multiplexing efficiency
  - – significant switch latency
    - store and forward
    - queueing delay: very bad behind large messages

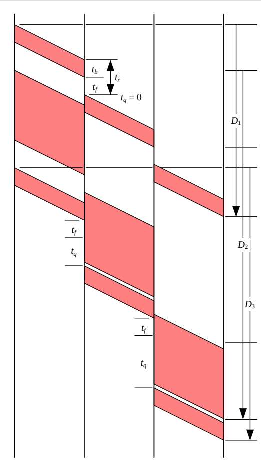

Network Control and Signalling © Sterbenz and Touch

Signalling Paradigms
Datagram Packet Switching

- Second generation (1970s)
- Messages packetised
  - + no setup latency
  - + significant multiplexing
  - + less queueing delay behind large packets
  - – store-and-forward per packet processing still required
  - – data rate limited by packet processing throughput

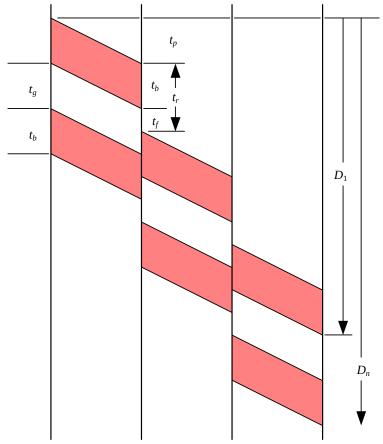

25 August 2006 **High-Speed Networking** 4-7

Network Control and Signalling © Sterbenz and Touch

Signalling Paradigms
Connection-Oriented Fast Packet Switching

- Motivation:
  - – eliminate store-and-forward
  - – high-performance switch design
  - – provision of QOS
- Connections
  - – establish state once to reduce per packet processing
    - + amortised for long flows
    - – expensive for transactions
    - – RTT delay before data transfer
    - + high throughput possible
    - + per hop messages reduce latency

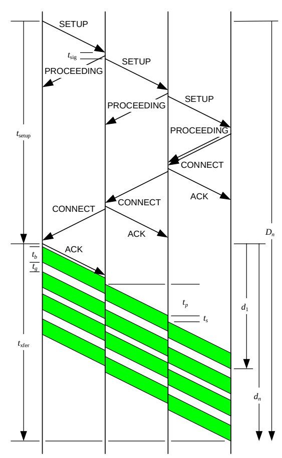

25 August 2006 **High-Speed Networking** 4-8

Signalling Paradigms
Connections vs. Connectionless

- Connections enabled fast switching in the 1980s
  - standardised as ATM, but with serious flaws
- Difficult tradeoffs between two paradigms
  - intermediate schemes are possible...

Connectionless vs. Connection Tradeoff N-5A

The latency of connection setup must be traded against the reduction in the end-to-end data transfer delay due to the elimination of store-and-forward delay, and faster transmission of due to nodes capable of increased bandwidth.

Signalling Paradigms
Signalling Efficiency

- Signalling complexity tradeoff between
  - needed functionality
  - simple state machines and message formats
- Signalling protocols must be robust to lost messages
  - state machine must account for this
  - messages should fit in single packet
    - reliable hop-by-hop protocols add significant complexity
    - (e.g. ATM S-AAL)

Efficiency of Signalling N-8B

Signalling messages should be simple in coding format and fit in a single packet to minimise the latency in processing. The signalling protocol should be robust to lost messages.

Signalling Paradigms

Round Trip Minimisation

- Techniques to minimise control round trips
  - hop-by-hop acknowledgements
  - parameter ranges
    - desired – minimum acceptable
    - desired, alternate
  - overlap of control and data

Minimise Round Trips N-6A

Structure control messages and the information they convey to minimise the number of round trips required to accomplish data transfer.

Signalling Paradigms

Round Trip Minimisation – HBH ACKs

- E2E handshake
  - requires RTT timeout
  - for *all* connect failures
- HBH ACKs
  - quicker local recovery

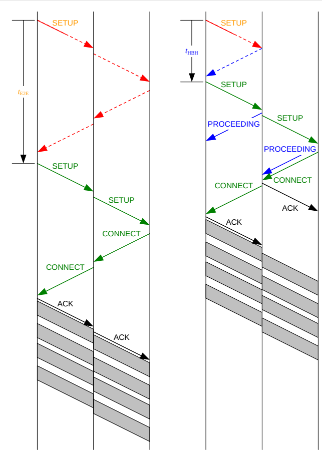

Signalling Paradigms

Round Trip Minimisation – Range Parameters

- Specific paramaters
  - exact match or...
  - another RTT attempt
- Incremental negotiation
  - initiator gives range or set
  - destination ACKs capabilities

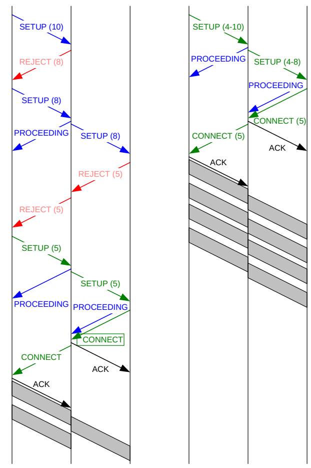

Signalling Paradigms

Round Trip Minimisation – Overlap Control/Data

- Separate phases
  - control to establish path
  - when successful
  - transfer data
- Overlap control/data
  - control to establish path
  - optimistically transfer data
  - when successful
  - continue transferring

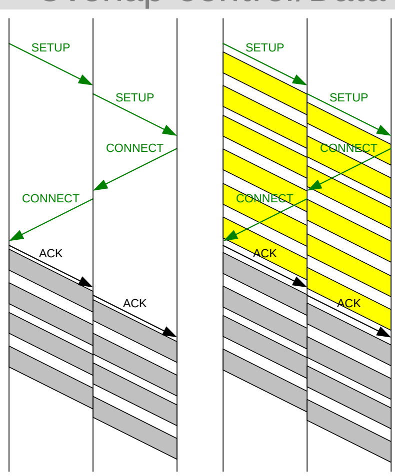

Signalling Paradigms

Intermediate Mechanisms

- Spectrum of signalling paradigms
  - per message forwarding (message switching)
  - per packet datagram forwarding
  - *data-driven soft state accumulation*
  - *control-driven soft state accumulation*
  - *optimistic connection establishment*
  - *fast reservation*
  - explicit virtual connection setup – fast packet switching
  - scheduled connection sharing – burst switching
  - explicit physical connection setup – circuit switching
- Intermediate mechanisms benefit from extremes

increasing state ↓

Signalling Paradigms

Intermediate Mechanisms

- Intermediate mechanisms benefit from extremes
  - benefits of state accumulation (from connections)
  - benefits of immediate transmission of data (from datagrams)

Network Path Establishment N-II.1

The routing algorithms and signalling mechanisms must be capable of forwarding datagrams or establishing connections on sufficiently high-performance paths and with low latency to meet application needs.

Intermediate Signalling Paradigms

Soft State Accumulation

- Features
  - allow immediate packet send
  - once state established: fast lookup and forwarding
- Variants
  - data driven
    - switch detects flow from headers
  - control driven
    - explicit inter-switch signalling
      - e.g. MPLS

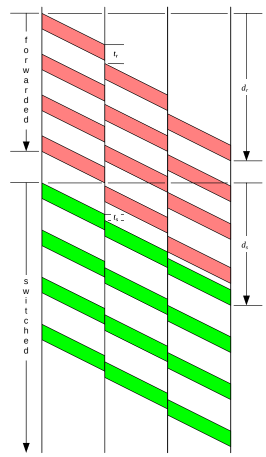

Intermediate Signalling Paradigms

Optimistic Connection Establishment

- Significant delay results from round trips
  - overlapping connection setup with data transfer to reduce

Overlap Signalling Messages with Data Transfer N-6A

To reduce unneeded end-to-end latency, signalling messages should be overlapped with data transfer.

Network Control and Signalling © Sterbenz and Touch

Intermediate Signalling Paradigms

Optimistic Connection Establishment

- Features
  - data sent with signalling
  - until COMMIT returned:
    - best effort or
    - per hop fast reservation

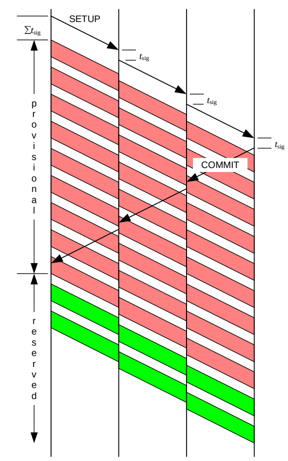

25 August 2006 **High-Speed Networking** 4-19

Network Control and Signalling © Sterbenz and Touch

Intermediate Signalling Paradigms

Burst Switching

- Features
  - packets group into bursts
  - bursts scheduled in circuit
- Role
  - optical networks:
    no photonic header processing
- Variants
  - release: signalled or timed
  - in- or out-of band signalling

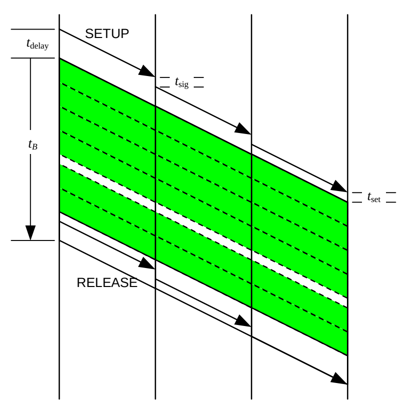

25 August 2006 **High-Speed Networking** 4-20

Multicast Signalling

Root vs. Leaf Join

- Root join
  - \+ simple
  - – doesn't scale

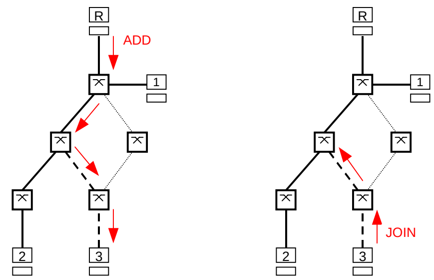

- Leaf join
  - – complex
  - \+ scales

Root vs. Leaf Multicast Control N-6B

Root multicast control is appropriate when a single entity has knowledge of the multicast group. Leaf-controlled multicast is needed for large multicast groups

Session Control

Signalling Efficiency

- Session (layer 5)
  - – set of transport layer associations (layer 4)
  - – point-to-point or multipoint (may be mixed)
  - – connection oriented or connectionless (may be mixed)

Session–Connection Interlayer Awareness N-4F

Session control awareness of network layer control parameters allows the overlap of session and connection control signalling to reduce overall latency

Session Control

Signalling Efficiency

- Latency reduction
  - minimise round trips
    - parallelise session and network signalling
  - properly locate resources

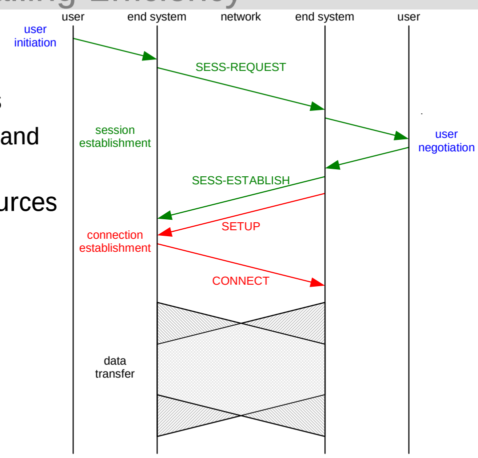

Network Control and Signalling

Traffic Management

Traffic Management
Goals

- Conflicting goals:
  - sufficient resources to deliver required QOS to users
  - minimise network resource use to keep costs low

Network Path Protection N-II.2

QOS mechanisms must be capable of guaranteeing bandwidth and latency bounds, when needed.

Traffic Management
Optimality

- Support for mixed traffic
  - distinct networks (1st and 2nd generations through 1980s)
  - virtual network partitioning
  - differentiated services: coarse grained service grades
  - integrated services: fine grained traffic classes
    - e.g. ATM-TM and intserv with RSVP

Optimal Resource Utilisation vs. Overengineering N-5A

Tradeoff *Balance the tradeoff of optimal resource utilisation and its associated complexity and cost against the cost of suboptimal resource utilisation resulting from overengineering the network.*

Network Congestion Control
Motivation

- Motivation
  - reservations are generally statistical
    - congestion can occur if there is any overbooking of resources
  - best-effort traffic must be limited to avoid congestion
  - not all end users will behave as they should

Congestion Control in the Network Improves Performance N-3A *Even though the end-to-end protocols must perform congestion control, there is substantial performance benefit in assistance from the network.*

Network Congestion Control
Performance Goal

- Goals
  - congestion control in the network reduces control loop delay

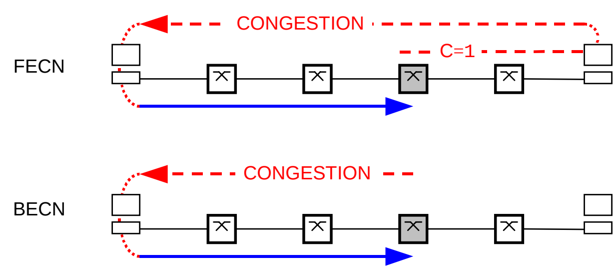

Network Congestion Control
Congestion Avoidance

- Congestion Avoidance
  - reacts to impending congestion *before* damage is done
    - e.g. RED (random early detection)
    - e.g. PPD/EPD (partial/early packet discard)

Avoid Congestion and Keep Queues Short N-II.4

Avoid congestion by network engineering, traffic management with resource reservation, and by dropping packets. Buffers should be kept as empty as possible, with queueing only for transient situations, to allow cut-through, and avoid the latency of FIFO queueing.

Network Congestion Control
Fairness vs. Complexity

- Fairness
  - desirable to allow fair sharing of network
  - difficult to discriminate well-behaved and misbehaving flows
  - fair mechanisms substantially more complex to implement

Congestion Control Fairness vs. Complexity N-2B*c*

The lack of fairness in simple congestion control and avoidance mechanisms must be traded against the complexity of fair implementations.

Network Control and Signalling

Path Routing Dynamics

Path Routing Dynamics

Dynamic Network Behaviour

- Dynamic network behaviour
  - leads to low performance paths over time
  - multipoint groups
    - prune and reroute to maintain optimal topology
  - node mobility
    - reconfigure topology to maintain performance

Dynamic Path Rerouting N-2B

Dynamic behaviour can require adjustments to topology to maintain a high-performance path. The overhead and frequency of topology maintenance must be traded against the lack of optimality.

Path Routing Dynamics

Node Mobility

- Node mobility changes path characteristics
- Example
  - $d_1 + d_2 < D_a$
  - intermediate node moves away
  - latency bound exceeded: $d_1 + d_2 > D_a$
  - reroute path
  - $d_3 < D_a$

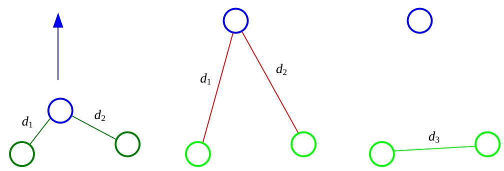

Path Routing Dynamics

Multipoint Groups$_1$

- Multipoint spanning tree
  - optimised to receiver group...
    ...at a particular point in time

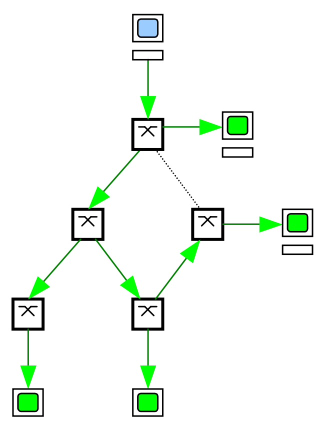

Path Routing Dynamics

Multipoint Groups$_2$

- Dynamic behavior
  - changes topology
  - latency bounds may be exceeded
  - bandwidth wasted
- Receivers leave group
  - inefficient path to remaining members
  - links carrying traffic to no receiver

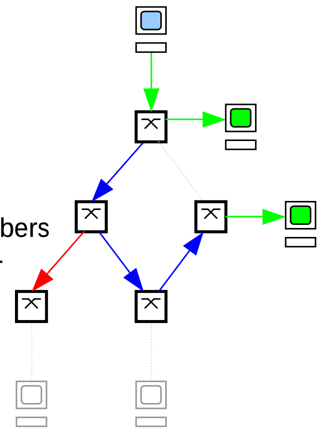

Path Routing Dynamics

Multipoint Groups$_3$

- Prune and reroute tree
  - optimise to remaining group
  - reroute to group members
  - prune unneeded leaf hops

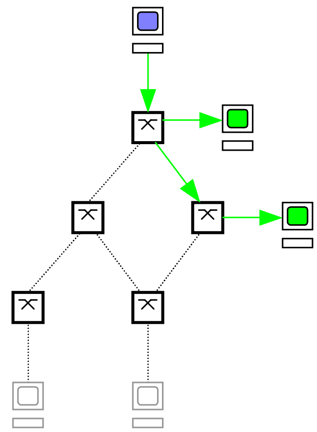

Network Control and Signalling

Monitoring and Management

Monitoring and Management

Issues

- Coarse granularity
  - management not high-speed, per se
- Issues
  - management must keep up with rapidly changing conditions
  - massive amounts of data must be filtered and reduced
  - monitoring should not interfere with high-speed flows

Network Monitoring Locality N-1B

Network monitoring functions must be built into the critical path to provide non-intrusive local filtering and aggregation of statistics.