
PNNI: Routing in ATM Networks

Raj Jain

Professor of CIS

The Ohio State University

Columbus, OH 43210

Jain@cis.ohio-state.edu

These slides are available at

<http://www.cis.ohio-state.edu/~jain/cis777-99/>



Distribution of topology information

Hierarchical groups



Source routing \Rightarrow Designated Transit Lists

Crankback and Alternate routing

Addressing

ref: "PNNI V1.0 Specification (Mar 1996)"

PNNI



Private Network-to-network Interface

Private Network Node Interface

Features of PNNI

Point-to-point and point-to-multipoint connections

Can treat a cloud as a single logical link

Multiple levels of hierarchy \Rightarrow Scalable for global networking

Reroutes around failed components at connection setup

Automatic topological discovery \Rightarrow No manual input required

Connection follows the same route as the setup message (associated signaling)

Uses: Cost, capacity, link constraints, propagation delay

Also uses: Cell delay, Cell delay variation, Current average load, Current peak load

Uses both link and node parameters

Supports transit carrier selection

Supports anycast

Addressing

Multiple formats.

All 20 Bytes long addresses.

Left-to-right hierarchical

Level boundaries can be put in any bit position

13-byte prefix \Rightarrow 104 levels of hierarchy possible




Link State Routing


Each node sends “Hello” packets periodically and on state changes.

The packet contains state of all its links


The packet is flooded to all nodes in the network





Very Large Networks



Hierarchical Layers



Hierarchical View



Terminology

Peer group: A group of nodes at the same hierarchy

Border node: one link crosses the boundary

Logical group node: Representation of a group as a single point

Logical node or Node: A physical node or a logical group node

Child node: Any node at the next lower hierarchy level

Parent node: Logical group node at the next higher hierarchy level

Logical links: links between logical nodes

Terminology (Cont)

Peer group leader (PGL):

Represents a group at the next higher level.

Node with the highest "leadership priority" and highest ATM address is elected as a leader.

Continuous process \Rightarrow Leader may change any time.

PGL acts as a logical group node.

Uses same ATM address with a different selector value.

Peer group ID: Address prefixes up to 13 bytes

Topology State Information

Metric: Added along the path, e.g., delay

Attribute: Considered individually on each element.

- q Performance, e.g., capacity or
- q Policy related, e.g., security

State parameter: Either metric or attribute

Link state parameter. Node state parameter.

Topology = Link + Nodes

Topology state parameter: Link or node state parameter

PNNI Topology state element (PTSE):

Routing information that is flooded in a peer group

PNNI Topology state packet (PTSP): Contains one PTSE

Topology State Parameters

Metrics:

- q Maximum Cell Transfer Delay (MCTD)
- q Maximum Cell Delay Variation (MCDV)
- q Maximum Cell Loss Ratio (MCLR)
- q Administrative weight

Attributes:

- q Available cell rate (ACR)
- q Cell rate margin (CRM) = Allocated - Actual
First order uncertainty. Optional.
- q Variation factor (VF) = CRM/Stdv(Actual)
Second order uncertainty. Optional.
- q Branching Flag: Can handle point-to-multipoint traffic
- q Restricted Transit Flag: Supports transit traffic or not

Database Synchronization and Flooding

Upon initialization, nodes exchange PTSE headers
(My topology database is dated 11-Sep-1995:11:59)

Node with older database requests more recent info

After synchronizing the routing database, they advertise the link between them

The ad (PTSP) is *flooded* throughout the peer group

Nodes ack each PTSP to the sending neighbors, update their database (if new) and forward the PTSP to all *other* neighbors

All PTSEs have a life time and are aged out unless renewed

Only the node that originated a PTSE can reissue it.

PTSEs are issued periodically and also event driven.

Information Flow in the Hierarchy


Information = Reachability and topology aggregation

Peer group leaders *summarize* and circulate info in the parent group

A raw PTSE never flows upward.

PTSEs flow horizontally through the peer group and downward through children.

Border nodes do not exchange databases (different peer groups)






Topology Aggregation


Get a simple representation of a group

Alternatives: Symmetric star (n links) or mesh ($n^2/2$ links)

Compromise: Star with **exceptions**



Address Summarization



Summary = All nodes with prefix xxx, yyy, ...
+ foreign addresses

Native addresses = All nodes with prefix xxx, yyy, ...

Example:

q A.2.1 = XX1*, Y2*, W111 A.2.2 = Y1*, Z2*

q A.2.3 = XX2*

q A.2 = XX*, Y*, Z2*, W111. W111 is a foreign address


Address Scope

Upward distribution of an address can be inhibited, if desired.

E.g., Don't tell the competition B that W111 is reachable via A.

Each group has a level (length of the shortest prefix).

Each address has a scope (level up to which it is visible).



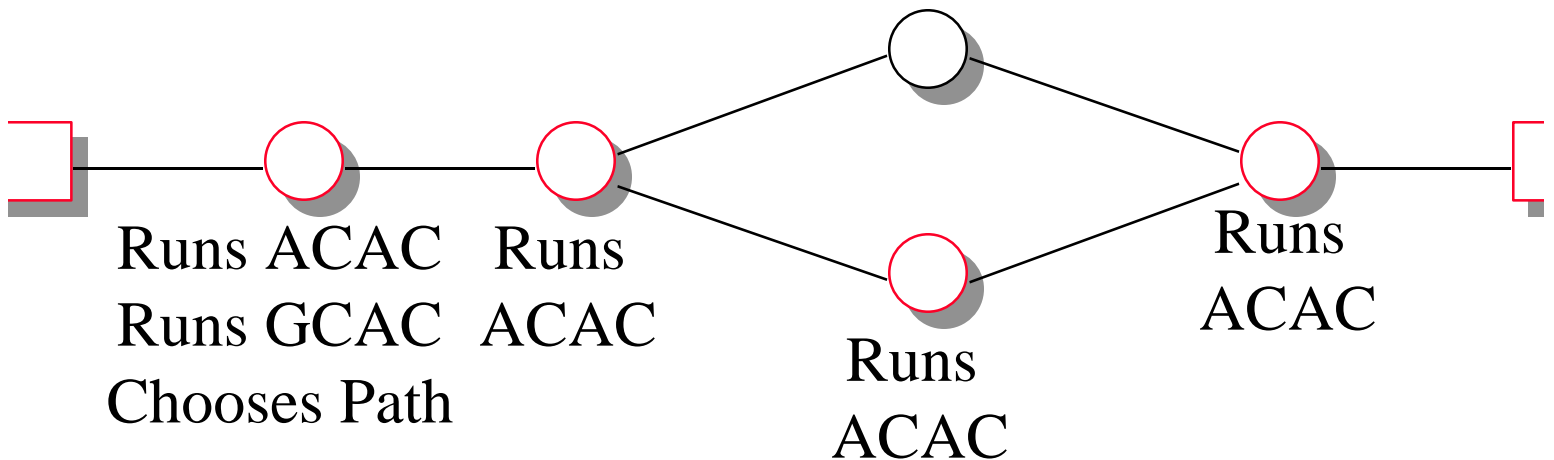
Call Admission Control

Generic Call Admission Control (GCAC)

- q Run by a switch in choosing a source route
- q Determines which path can probably support the call

Actual Call Admission Control (ACAC)



- q Run by each switch
- q Determines if it can support the call



Source Routing

Used in IEEE 802.5 token ring networks

Source specifies all intermediate systems (bridges) for the packet



Designated Transit Lists

DTL: Source route across each level of hierarchy

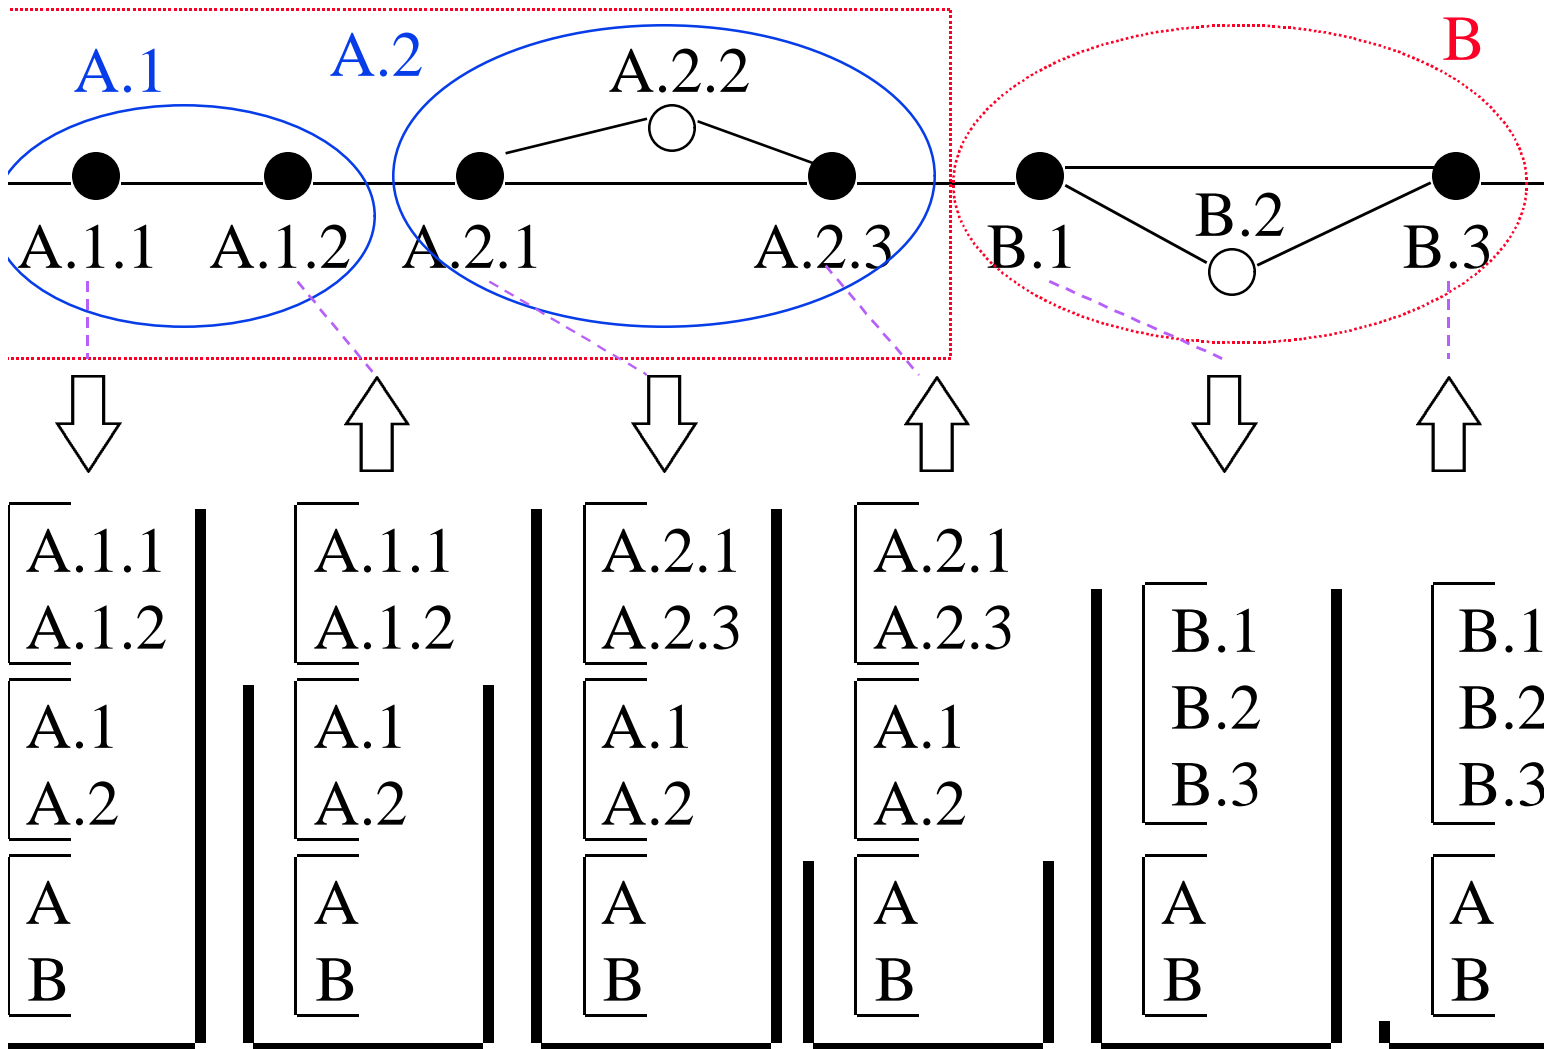
Entry switch of each peer group specifies complete route through that group

Entry switch may or may not be the peer group leader

Multiple levels \Rightarrow Multiple DTLs

Implemented as a stack


DTL: Example



Cranks back and Alternate Path Routing

If a call fails along a particular route:

- q It is *cranked back* to the originator of the top DTL
- q The originator finds another route *or*
- q Cranks back to the generator of the higher level source route



Summary



Database synchronization and flooding

Hierarchical grouping: Peer groups, group leaders

Topology aggregation and address summarization

Designated transit lists

Crankback

Homework

Read chapter 10 of Sackett and Metz