QoS-Based Routing - Some Issues

(June, 28, 1996)

Bala Rajagopalan
Bellcore
(braja@bellcore.com)

Raj Nair
Ascom/Nexion
(nair@nexen.com)

Motivation for QoS-Based Routing

- *Assumption: Internet will be resource-constrained during the foresseable future.*

- **Benefits of QoS-based routing:**

  —On-demand QoS-accommodaing path computation

  —Economic network engineering

  —Compensation for imprecise network engineering

  —Efficient routing

  —Graceful performance degradation during overload conditions.

High-Level Issues

- Requirements on IP QoS-based Routing

  —determines the complexity

- QoS-based routing architecture
- Technical and policy issues in QoS-based routing

Some Generic Requirements

- Support for multiple path metrics (bandwidth, delay, etc.). Dynamic determination of paths satisfying multiple constraints (e.g., bandwidth and delay).
- QoS support for multiple types of flows, i.e., unicast and many-to-many multicast.
- Robustness.
- The ability to route short-lived data flows with minimal overhead.
- Efficient utilization of network resources, through load spreading, global admission control techniques and the like.
- Minimization of routing overheads.
- Support for resource control to limit resource consumption by different classes of traffic.

Generic Requirements (Continued)

- Scalability w.r.t. nodes, links, and network diameter.
- Support for prioritizing flows, and to permit high priority flows to be established with precedence over lower priority flows.
- Interdomain and Policy?
- Other?

Architectural Considerations

- Evolution of the traditional intra and interdomain organization

  —Add QoS-related features to existing intra and interdomain routing schemes, if possible.

- New routing architectures

  —Scalable QoS-based routing based on new architectures such as Nimrod or ATMF PNNI, or an entirely new approach.

In either case, there are many technical issues.

Some Technical Issues

- Unicast and Multicast Path Computation
- Routing efficiency
- Traffic prioritization
- Short-lived flows
- Resource control
- Integration with the IP signalling model

Unicast Path Computation

- Algorithmic intractability
  - —Path computation subject to two or more additive constraints (e.g., delay and jitter) is NP-Complete (Wang and Crowcroft, 1996).
  - —Implies that some combination of link metrics cannot be used meaningfully in path computation.
  - —Bandwidth and delay combinatin can be used.
- QoS Determination
  - —Finding the QoS available on some links may be tricky.
  - —Delay estimation?
- Aggregation?

Providing a Delay QoS (Continued)

how do we admit nodes satisfying the delay constraint

Lets say it turns out that the next link chosen must contribute no greater than **y**

The question is how do we know that the delay seen by this flow after routing it through the link will not be greater than **y**

To answer that we need more than current delay on the link; we to know the currently available bandwidth on the link

Providing a Delay QoS (Continued)

To calculate the expected delay, we need the available

current delay. This will tell us where we are on the delay curve.

For example,

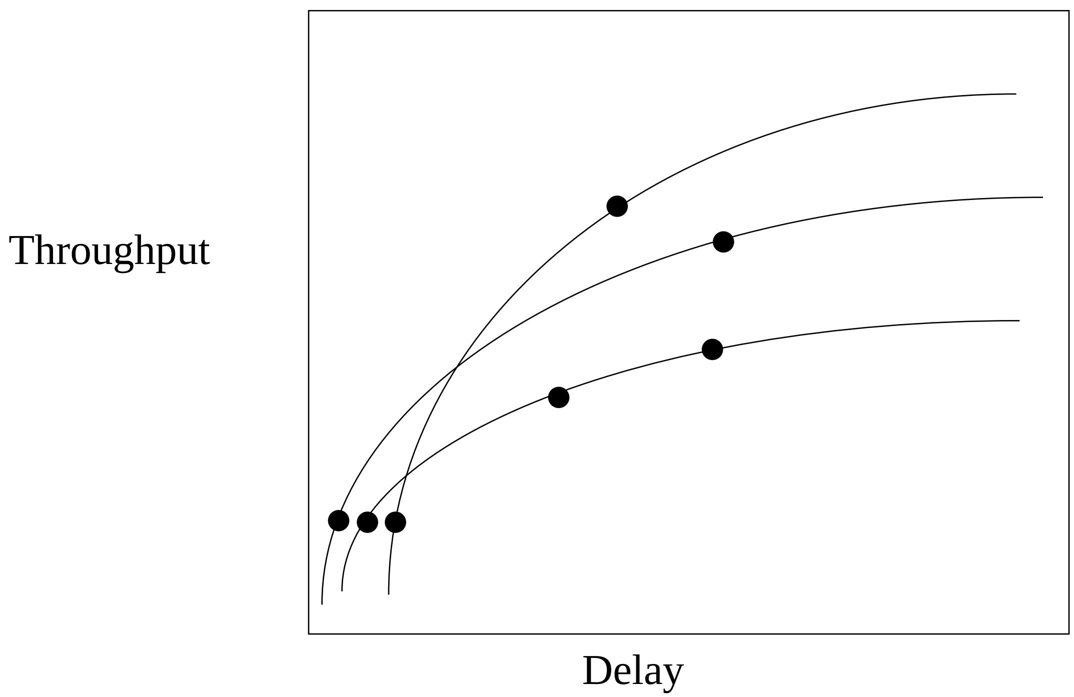

Information Loss Due to Aggregation

Area information is subject to loss due to aggregation

Each routing domain may have its own proprietary scheme for encoding metrics

The question is whether we can have a uniform encoding to allow destination to have the same view as the sender, even after the information has undergone multiple encoding/decodings through transit domains

Assume that each area advertises aggregated metrics

Issues with Aggregation

Consider a configuration of heterogenous networks

in the figure below.

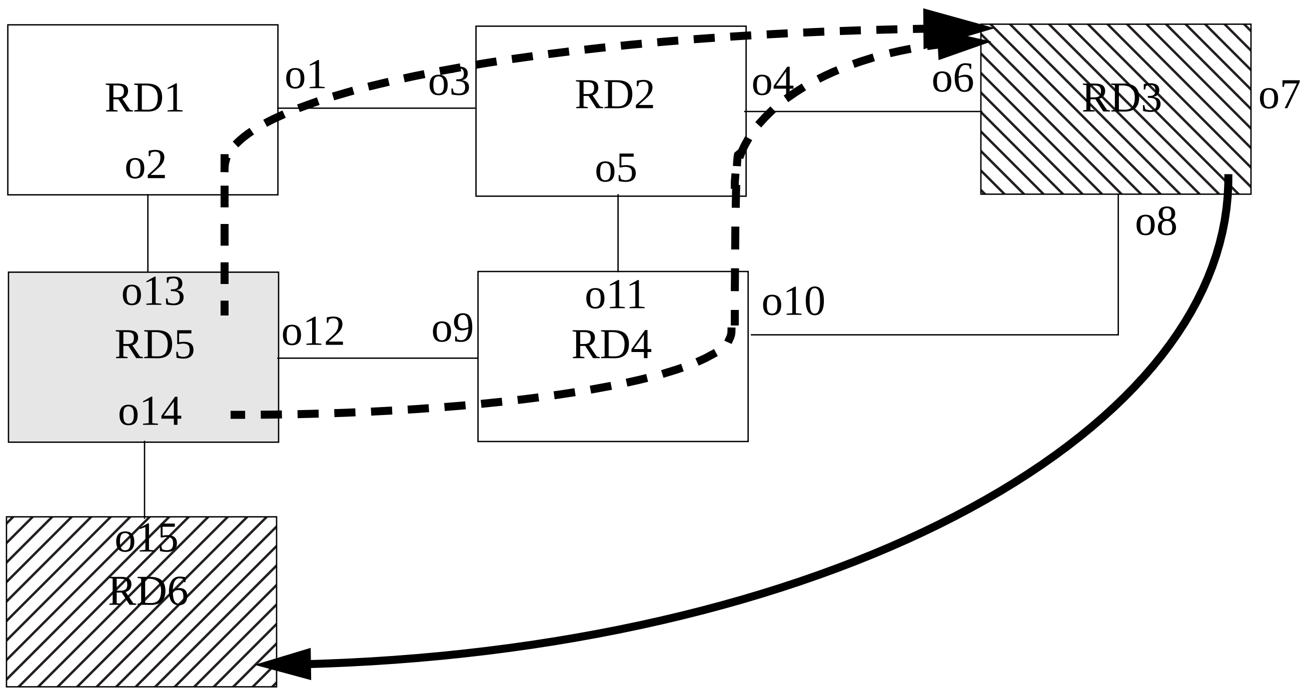

Issues with Aggregation (Continued)

For example, in passing the metric of logical links in RD5 to nodes in RD6, the information maintained by the routing algorithm in RD5 about the metric must be first converted to a common format, and then converted to the encoding used in RD6.

Say, a source in routing domain RD6 wants to route a flow to a destination in routing domain RD3. Which of many alternate paths from RD6 to RD3 is selected by the source depends on the metrics of these paths, as injected into RD6.

Issues with Aggregation (Continued)

The source may select the path RD5-RD4-RD2-RD3, which looks feasible to the source, but in reality infeasible. Also, the path RD5-RD1-RD2-RD3 may be feasible, but look infeasible to the source.

Worse, the call may be accepted on an infeasible path, because the QOS checks during flow set up are not done by a proprietary scheme for some link metric.

For example, if a flow needs a delay and bandwidth guarantee. The flow may be setup through the area but may actually support the two guarantees on two different physical paths

Multicast Path Computation

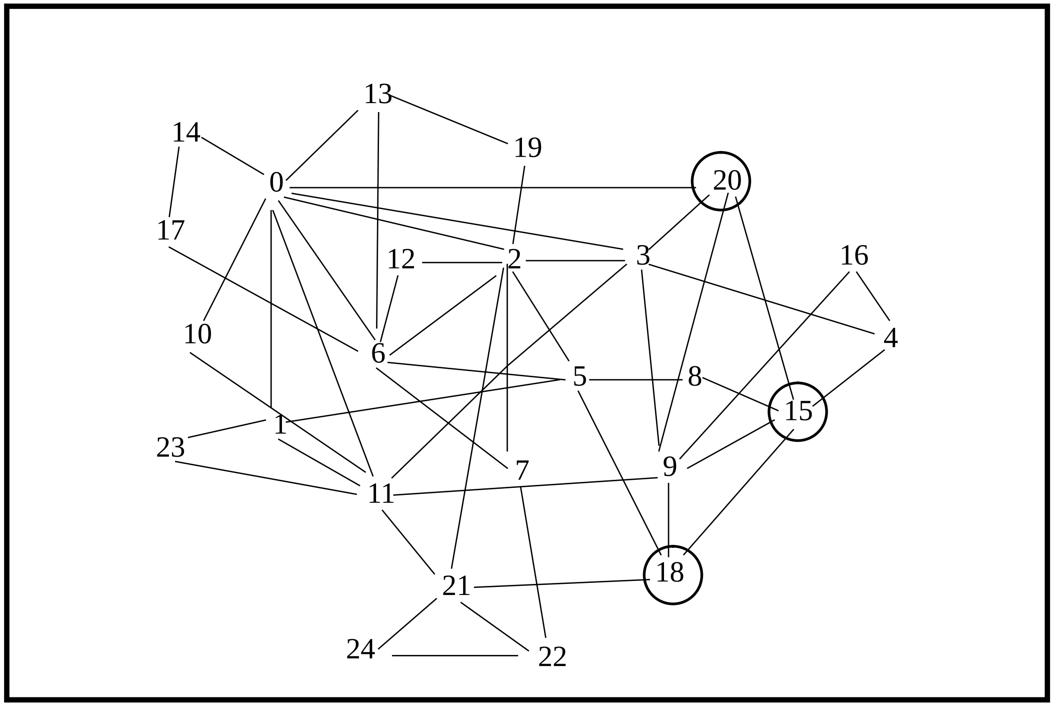

Heuristic algorithms for incremental path computation require each node to know the locations of all members plus their QoS requirements. Scalability implications?

Computing shared trees based on shared reservation style: Both sources and receivers must be known!

Routing Efficiency

- Connection admission control

—Admissibility of connections must be determined based on some notion of their overall resource utilization.

—Fairness in admitting connections

C = 12 for all links

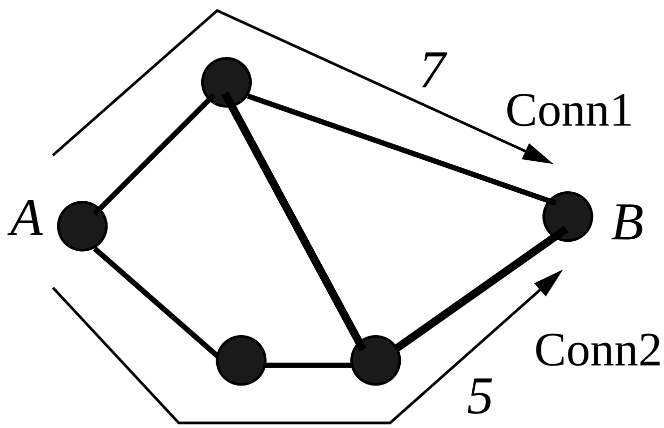

Connection 3, with request = 6 may be rejected admission.
Larger connections should be treated fairly by admission control.

Routing Efficiency (Contd.)

- Routing overhead reduction

—Tree-based forwarding of frequent state updates

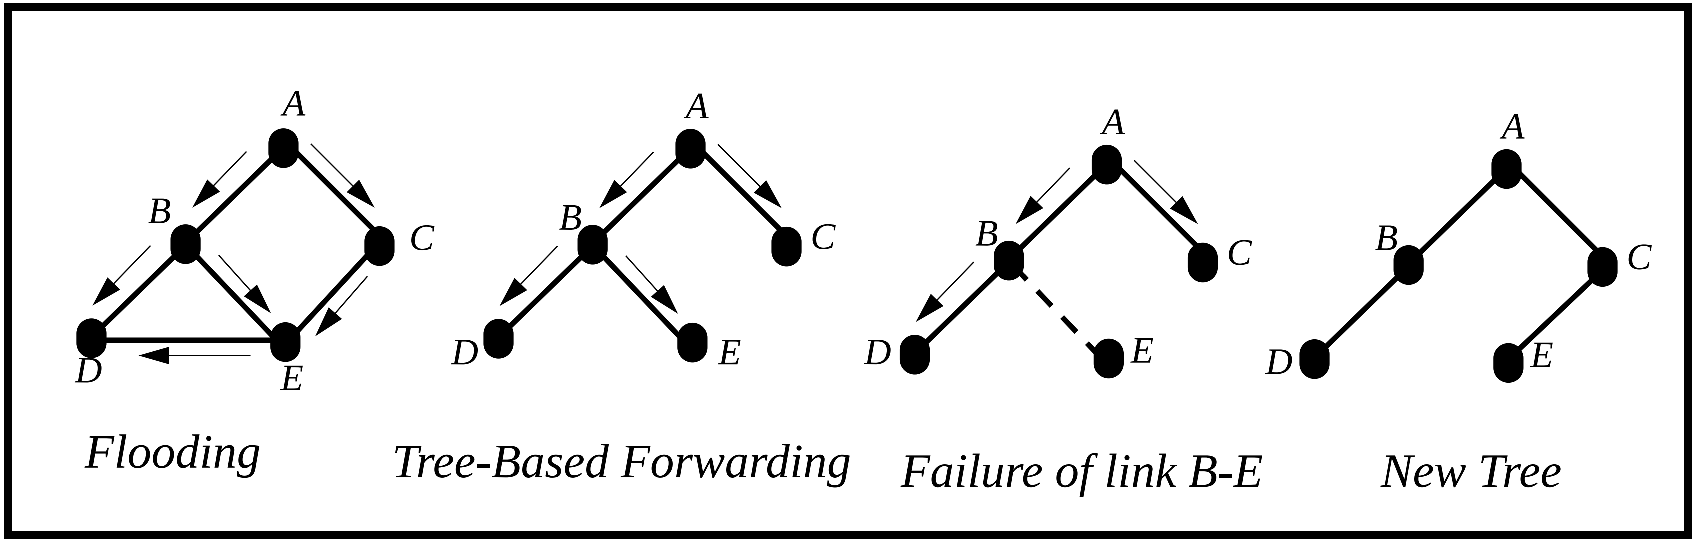

—Provable convergence for such procedures (Rajagopalan, 95)

—Metric quantization and effects on routing accuracy.

Flow Prioritization

- Problem: Prioritize flows and permit high priority flows to be established with precedence over lower priority flows.
- Example:

$C = 12$ for all links

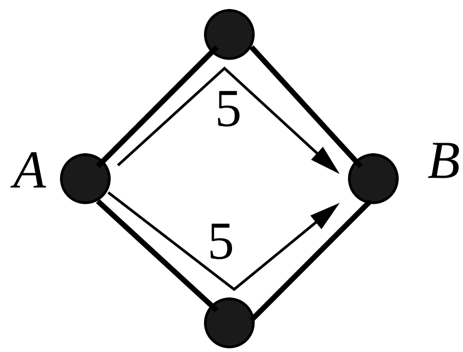

- A new higher priority flow of 8 units between A and B arrives. Preempt one of the flows and establish the new one?
- Prioritization may require additional information to be kept at each router. Scalability implications?

Short-Lived Flows

- Minimize routing overheads for short-lived flows
  - —Route cacheing?
  - —Route short-lived flows as best effort traffic?

Resource Control

- Control allocation of resources to different traffic classes

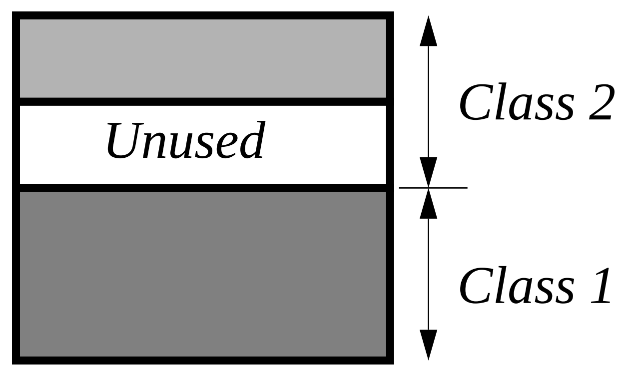

Allocations on a single trunk

- What is the policy in allocating unused resources?
- How to reclaim allocated resources from one class and give it to another?

Interaction with RSVP

- Unicast flow

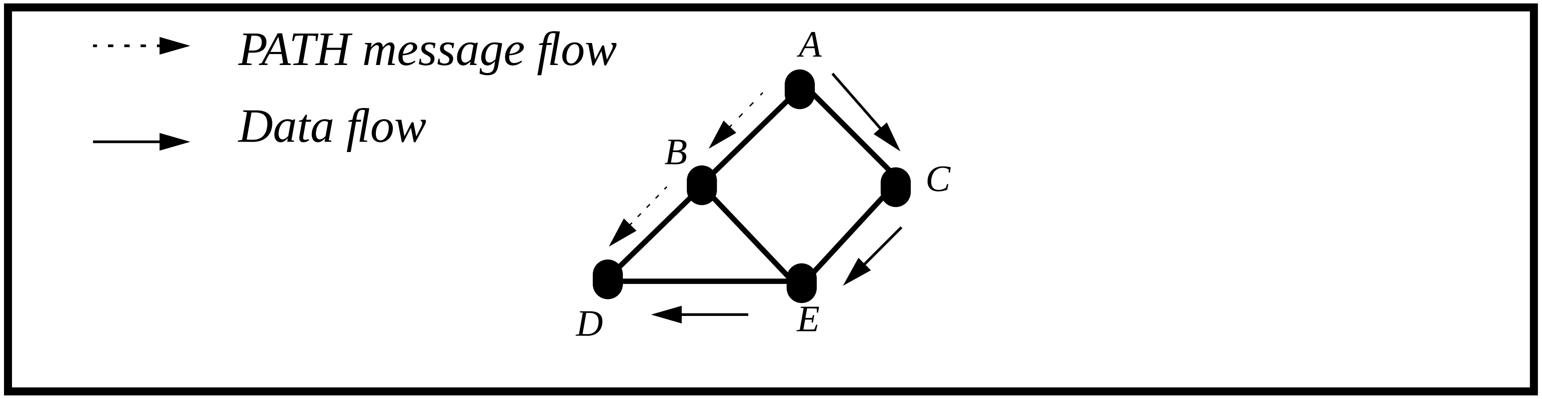

- Multicast flow