Internet Traffic Engineering

Jeng-Long Chiang

Nov 21, 2002

References

[1] Brownlee N. and Claffy K.C., “Understanding Internet traffic streams: dragonflies and tortoises.”

[2] Fortz B., Rexford J., and Thorup M., “Traffic engineering with traditional IP routing protocols.”

[3] Mortier R.M., “Multi-timescale Internet traffic engineering.”

- *IEEE Communications Magazine, Oct. 2002, Vol. 40, No. 10, pp 110-131.*

Outline

- Introduction
- Internet traffic streams
- Internet traffic engineering
  - Data timescale resource allocation
  - Control timescale resource allocation
- A simple example
- Conclusion

Introduction

- Traffic engineering involves adapting the routing of traffic to network conditions, with the joint goals of good user performance and efficient use of network resources.

- Traffic analysis helps the operators on arranging network resources and applying pricing policies.

Internet Traffic Streams

- Internet traffic streams can be classified by size and lifetime
  - By size:
    - mice v.s. elephants
    - (small v.s. high-volume stream)
  - By lifetime:
    - dragonflies v.s. tortoises
    - (short v.s. long-run stream)

Observations in [1]

- University of Auckland (UA)
- University of California at San Diego (UCSD)
- By school:
  - 9Mb/s at UA (charged per megabyte)
  - 622Mb/s at UCSD (utilization = 15%)
- By type:
  - Web TCP = 50%(UCSD)-80%(UA)
  - Non-Web TCP = 20%-50%
  - UDP = around 1%

Observations in [1]

- By size:(short streams)
  - Web TCP =
    - 20B, 50B, 100B, 200B (<300B, failed request)
    - 300B-800B (<40KB)
  - Non-Web TCP =
    - 30B, 50B, 800B, 1500B (<3KB)
  - UDP =
    - 30B, 80B (<300B)

Observations in [1]

- By lifetime:
  - Lasting 2s or less = 40%(UCSD)-70%(UA)
  - Lasting less than 15min = 98%
  - LR streams (>15min) = 1.5 % (NTP, SSH, DNS)
- By rate:(LR Streams):
  - TCP = variable rate (congestion)
  - UDP = constant rate
- By link utilization:
  - 5%(UA)-50%(UCSD) LR streams

Internet Traffic Engineering

- Traffic engineering is concerned with the performance optimization of networks.

- It addresses the problem of efficiently allocating resources in the network so that user constraints are met and operator benefit is maximized.

Resource Allocation

- Resources are allocated over three timescales:
  - Data timescales
  - Control timescales
  - Management timescales
- Network service provisioning
  - Service-oriented
  - Technology-oriented

RA in the Three Timescales

- Data timescales:
  - Congestion control
  - Per-packet payment
  - Router scheduling and marking
- Control timescales:
  - Admission control in ATM and MPLS
- Management timescales:
  - Load balancing routing (rfc891)

A Data Timescales Scheme [2]

- Recent standards activity has proposed TE extensions to OSPF and IS-IS to incorporate traffic load in the LSA and path selection decisions. However, modifications on routers are required.
- It is possible to select static link weights that are resilient to traffic fluctuations and link failures, allowing the use of the traditional incarnations of OSPF and IS-IS.

Tuning of IGP Weights

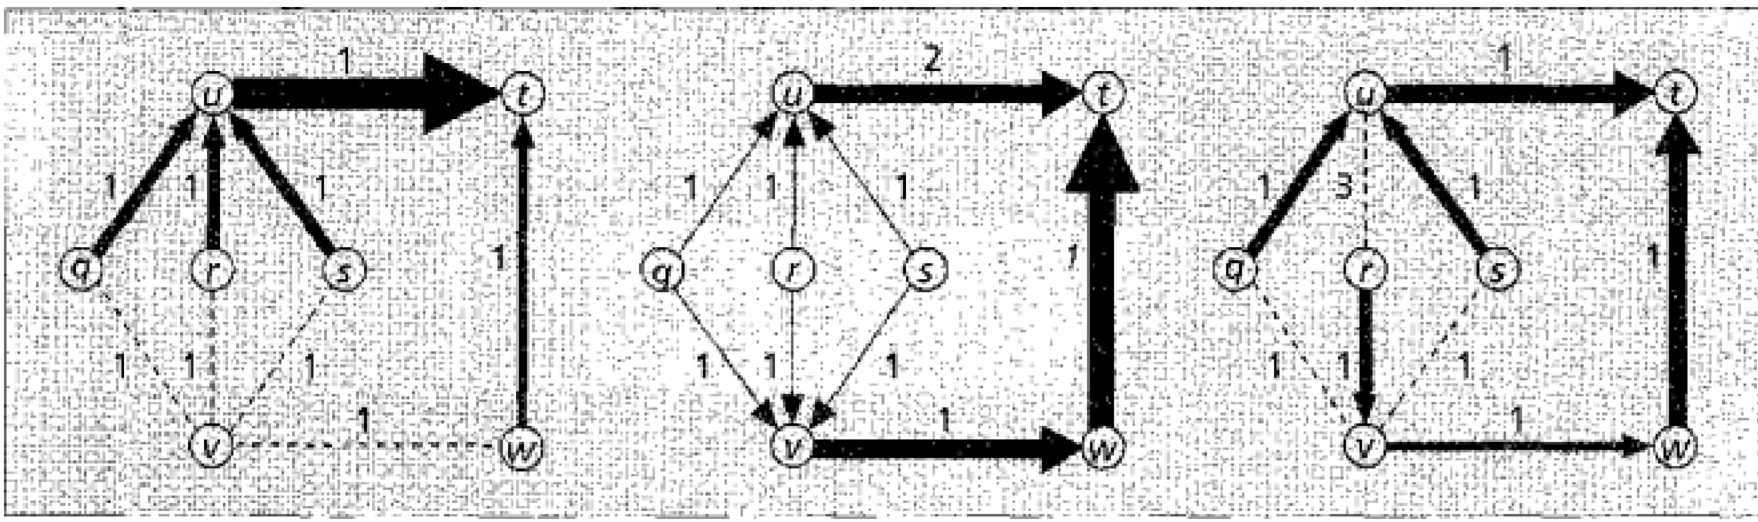

Figure 2. *Routing the same demands with different weight settings: each link has an integer weight, arrows illustrate the flow of traffic, the thickness of the arrows indicates the volume of traffic traversing the link, and a dashed line represents a link that carries no traffic.*

Advantages of Using Traditional OSPF/IS-IS

- Centralized setting of routing parameters:
  - Protocol stability
  - Low protocol overhead
  - Diverse performance constraints
- Path selection based on link weights:
  - Compatibility with traditional shortest path IGPs
  - Concise representation
  - Default weights and backup routes

Traffic Engineering Framework

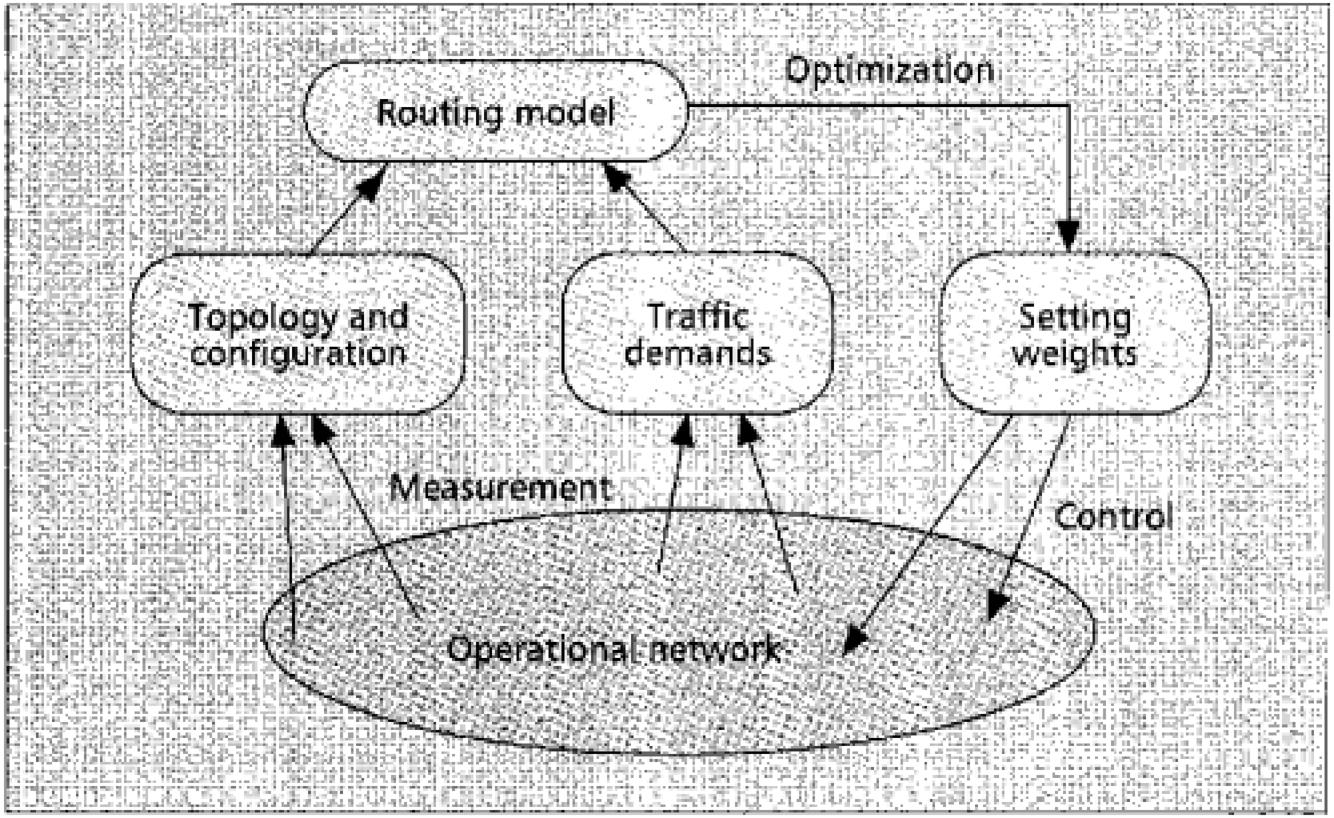

Figure 3. *Key components of the traffic engineering framework.*

Evaluation

- AT&T WorldNet backbone (90, 274)
- UnitOSPF(Wi=1, for all I)
- Cisco’s InvCapOSPF (Wi=1/Ci)
- AdvancedOSPF
- OPT direct traffic along any path in any proportions

Cost Function

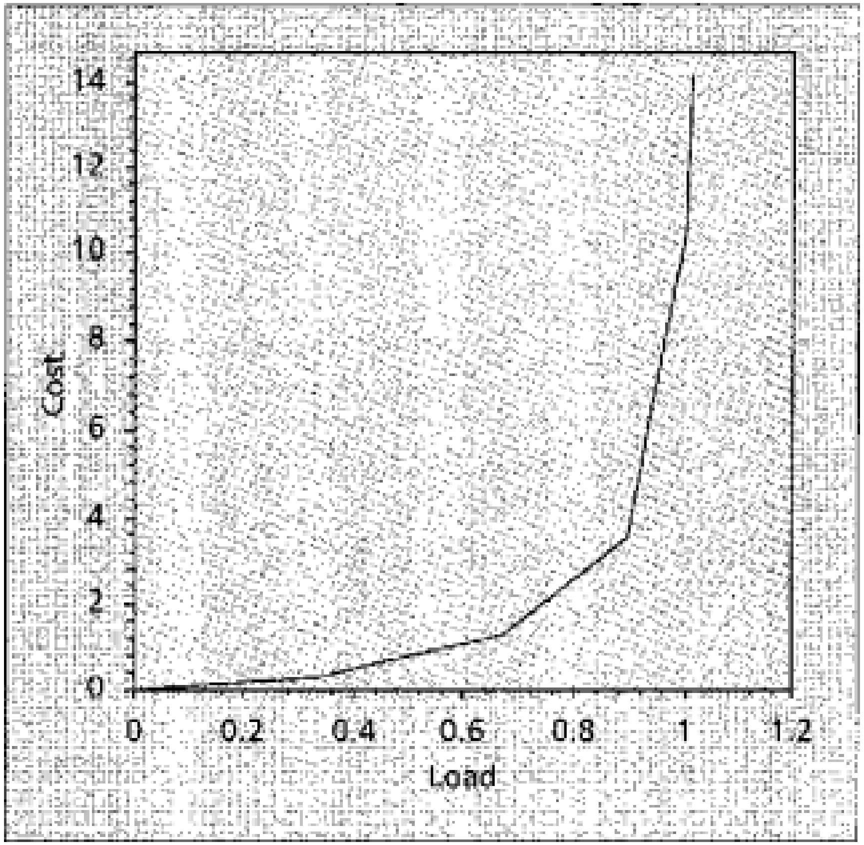

Figure 4. *Link cost as a function of the load for a link capacity 1.*

Performance Comparison

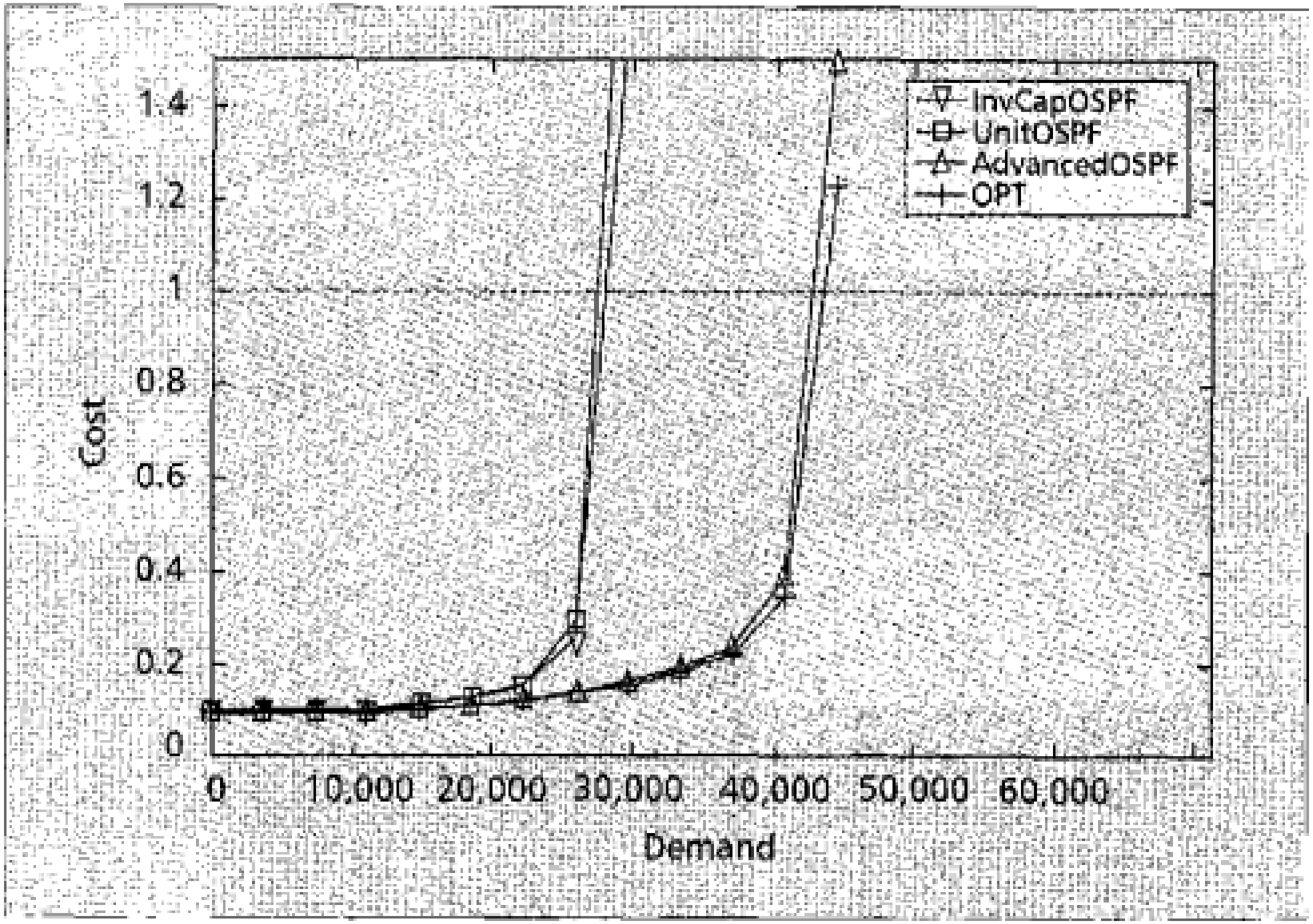

■ **Figure 5.** *Networkwide cost vs. demand for a proposed AT&T backbone.*

Observations in [2]

- Changing a single link weight is quite effective.
- The existing IGP weights continued to perform relatively well after a single link failure.
- A single weight change is sufficient to alleviate congestion that would arise after a link failure.
- Necessary weight changes could be pre-computed for any link failure.

Two Control Timescales schemes

- Implicit admission control scheme for TCP
  - TCP v.s. ATM
  - Measurement-based admission control (MBAC)
- ECN proxy for RTP
  - ECN=Explicit Congestion Notification
  - RTP=Real-Time Transport Protocol
    - RTP, RTCP(receiver report, RR), RTSP

Implicit Admission Control

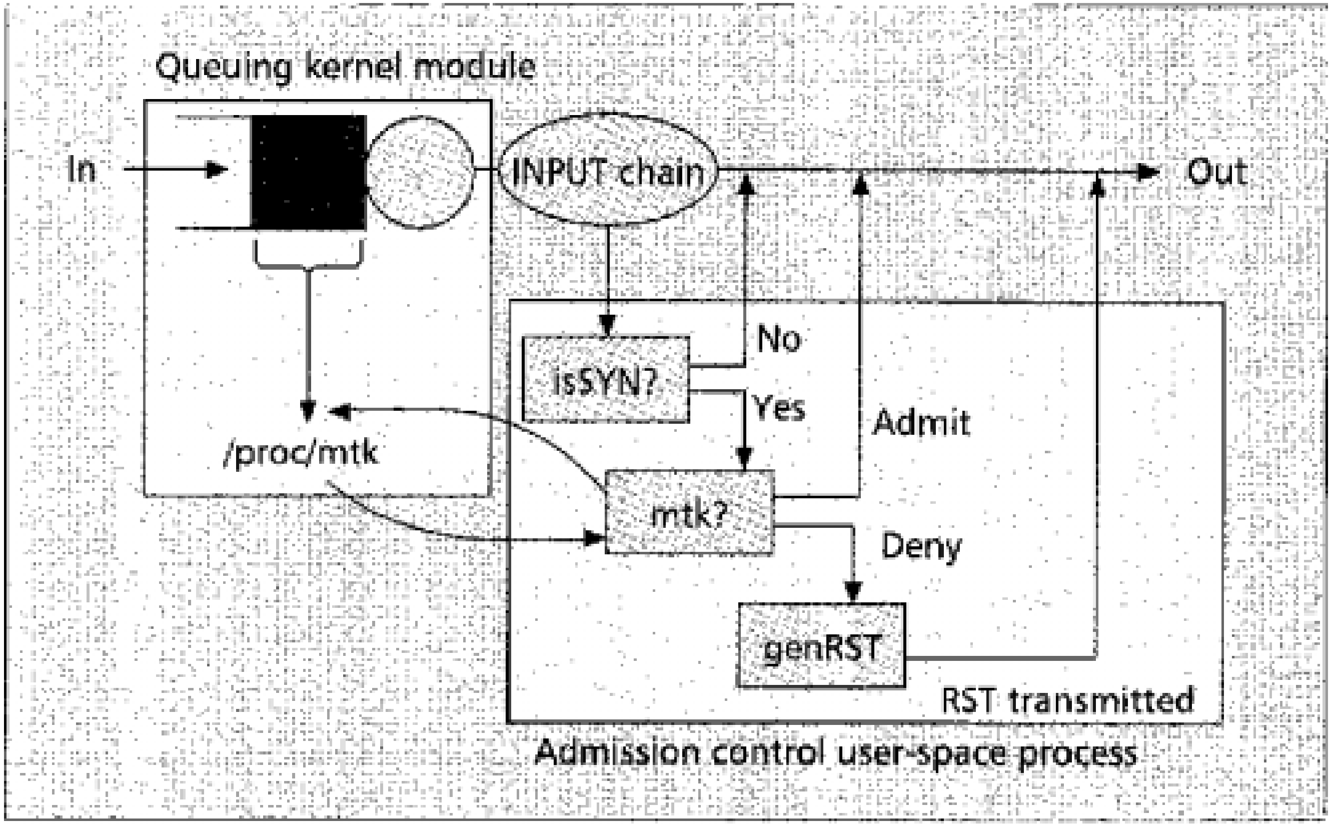

Figure 1. *Linux implicit admission control implementation.*

ECN Proxy for RTP

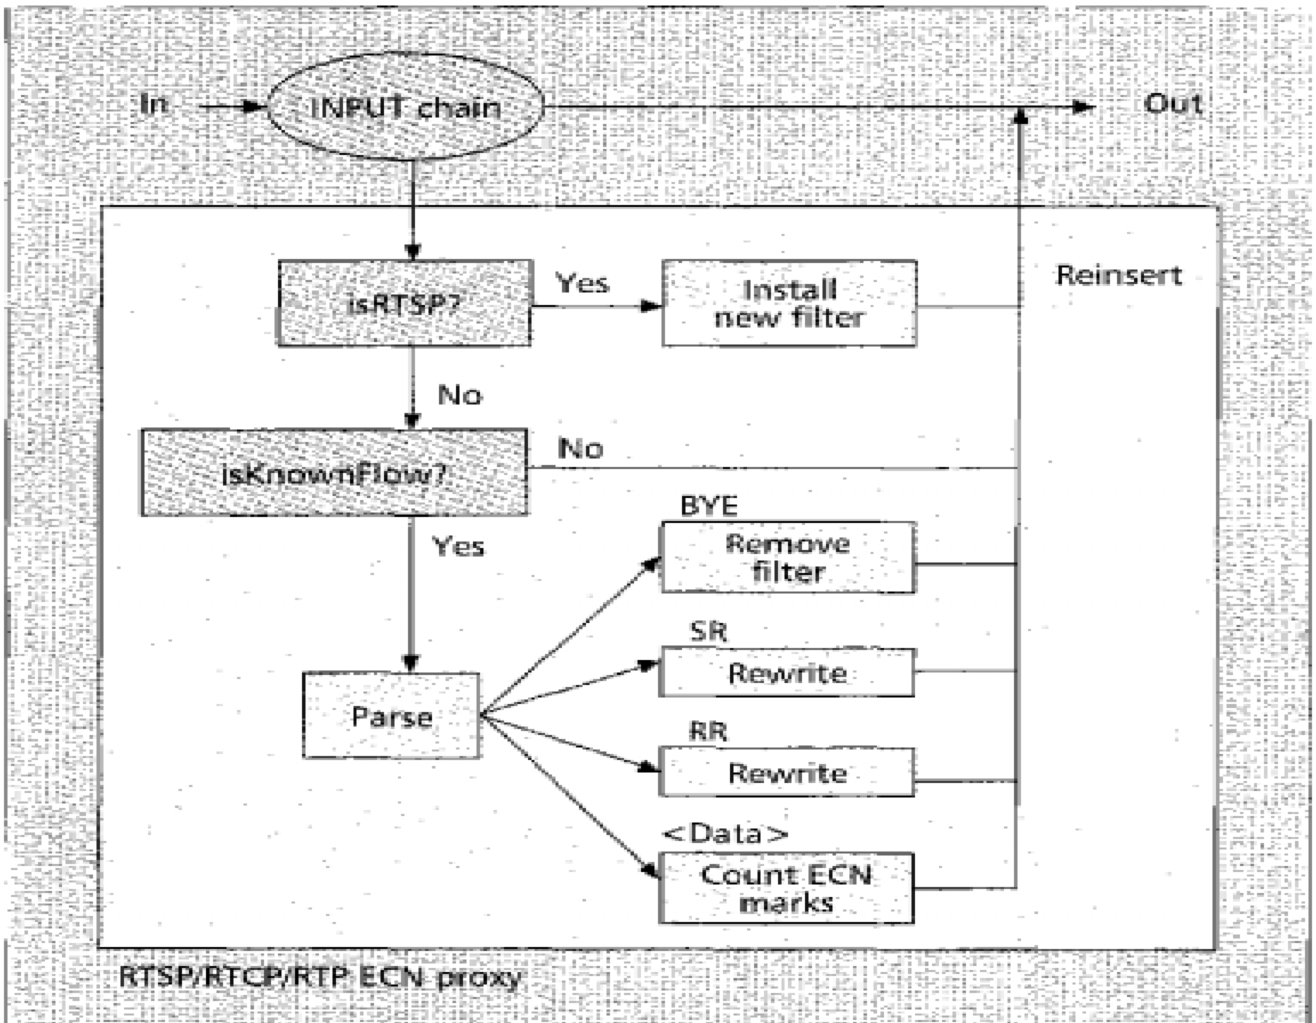

■ **Figure 2.** *Linux RTP-ECN-proxy implementation.*

A Simple Example

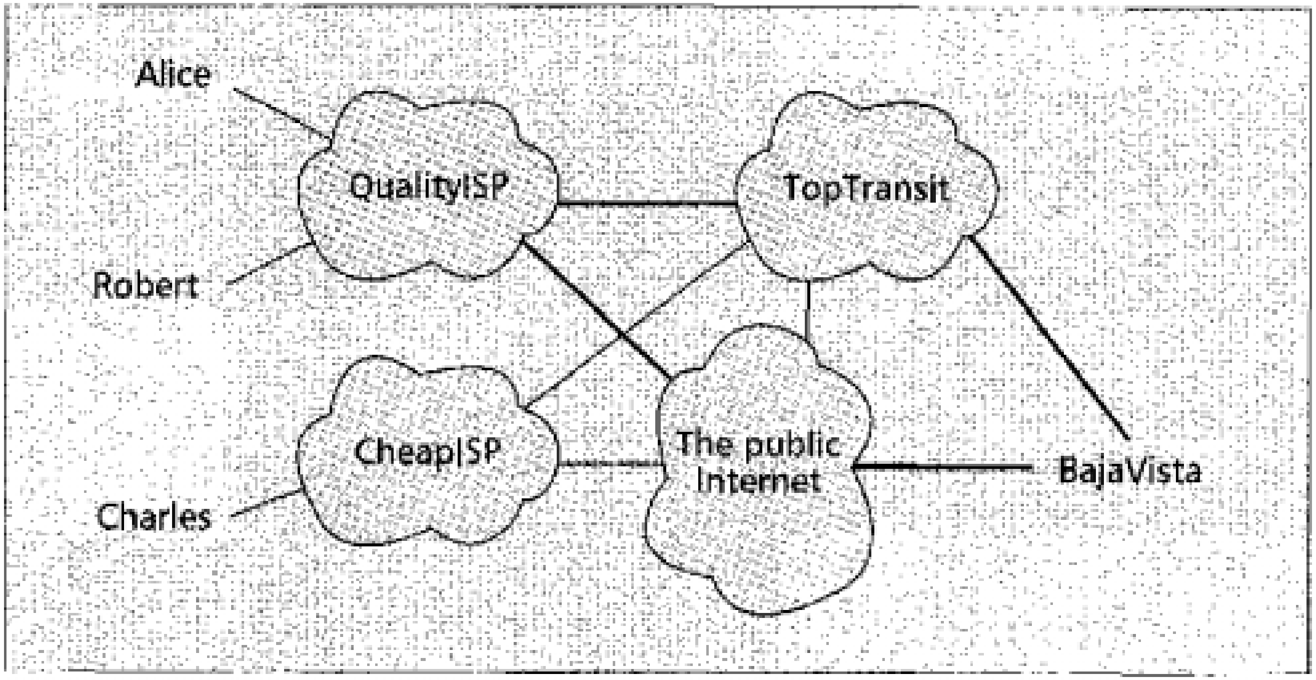

■ **Figure 6.** *A simple example.*

Conclusion(1/3)

- Streams in the Internet continue to become larger
  - Cable moden, xDSL, PC, P2P, streaming
- Short v.s LR streams:
  - Caching, queueing, scheduling, explicit route
  - TCP-regulated LR stream v.s. short stream

Conclusion(2/3)

- TE can be treated as a network operations task, rather than the responsibility of the underlying routing protocol.

- Traditional shortest path routing protocols are surprisingly effective for engineering the flow of traffic in large IP networks.

Conclusion(3/3)

- Multi-timescale TE techniques allow dealing with the competing desires of the operators to simplify the services they offer while still providing sufficient flexibility for users to express their individual requirements.