BB State Aggregation and Requirements for Wireless QoS Networks

Presented by

Olov Schelén

Luleå University/CDT, Operax AB

www.cdt.luth.se/~olov

Thanks to Telia Research, Ericsson, SICS

Today's presentation

- Trends for new applications
- Service requirements for adopting these applications into the Internet
  - VLLs
  - Advance reservations
- State aggregation in Bandwidth Brokers
- Actions for going wireless

Trend: A multi-service network

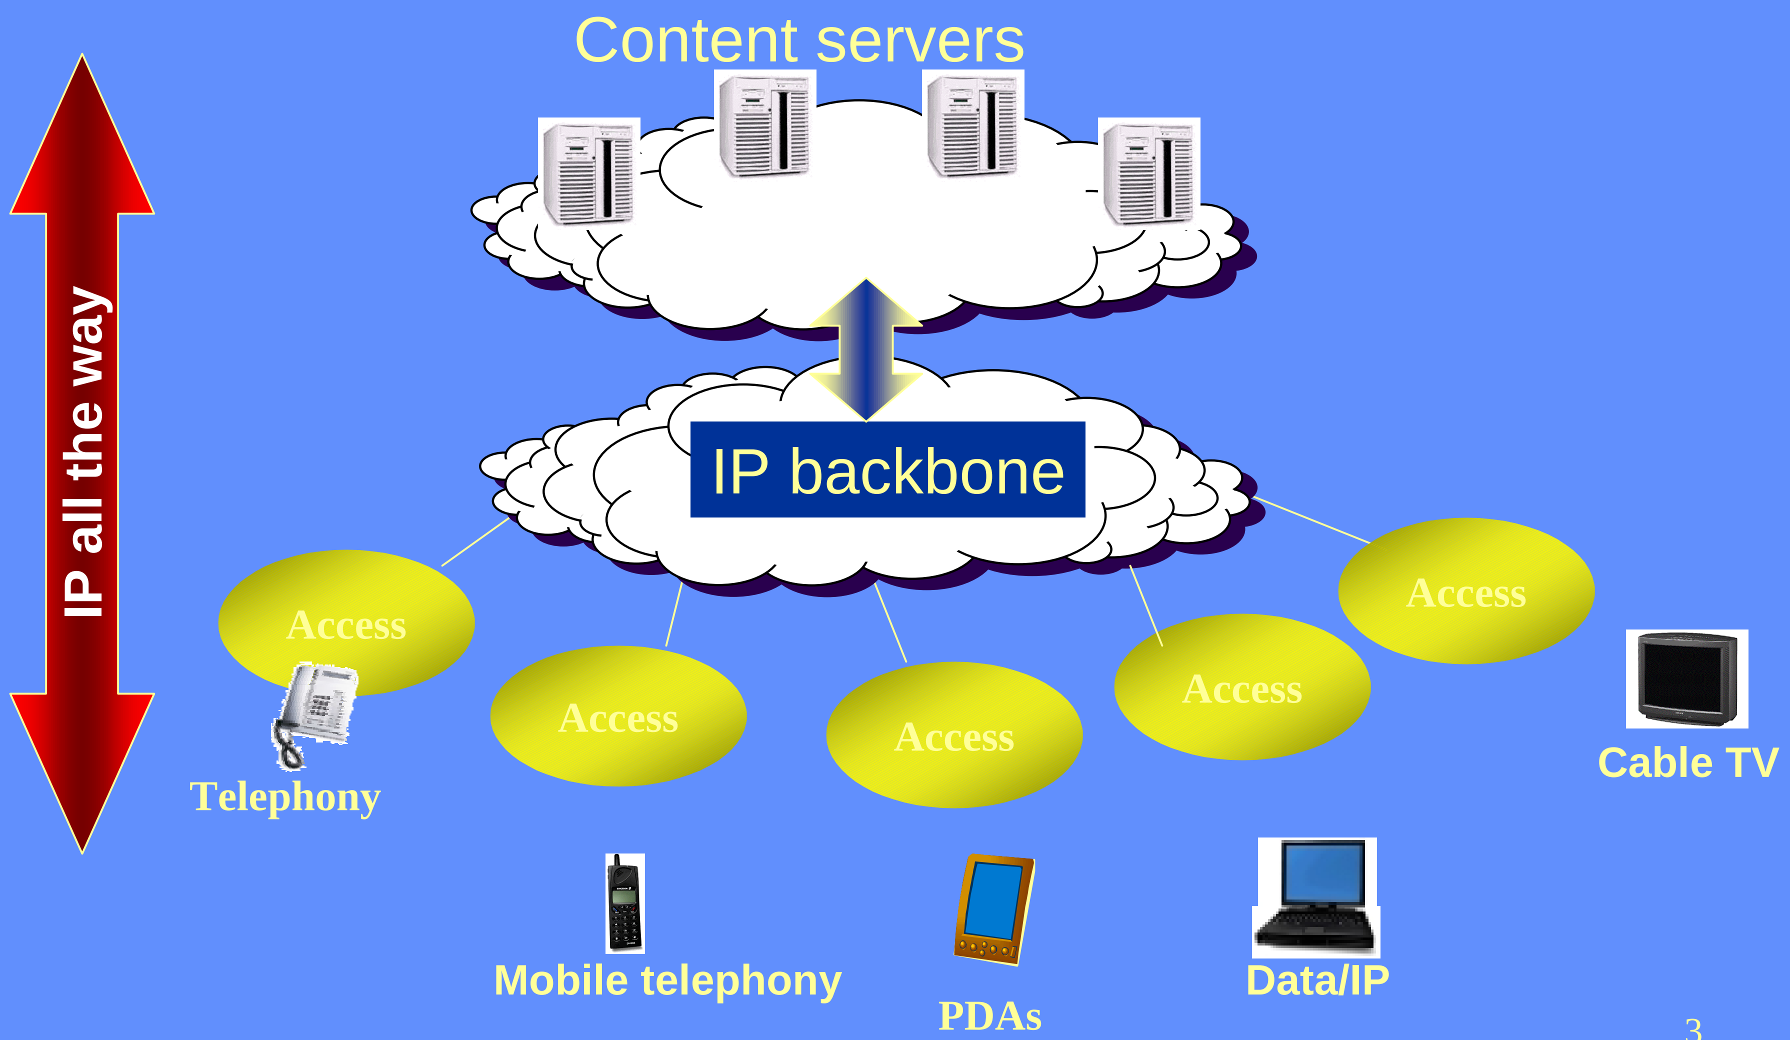

Mobile telephony

- Mobile telephony (GSM, TDMA, …) growth is comparable to the Internet
  - Projections on over 1 billion users 2004
- Focus
  - Mobile Internet!

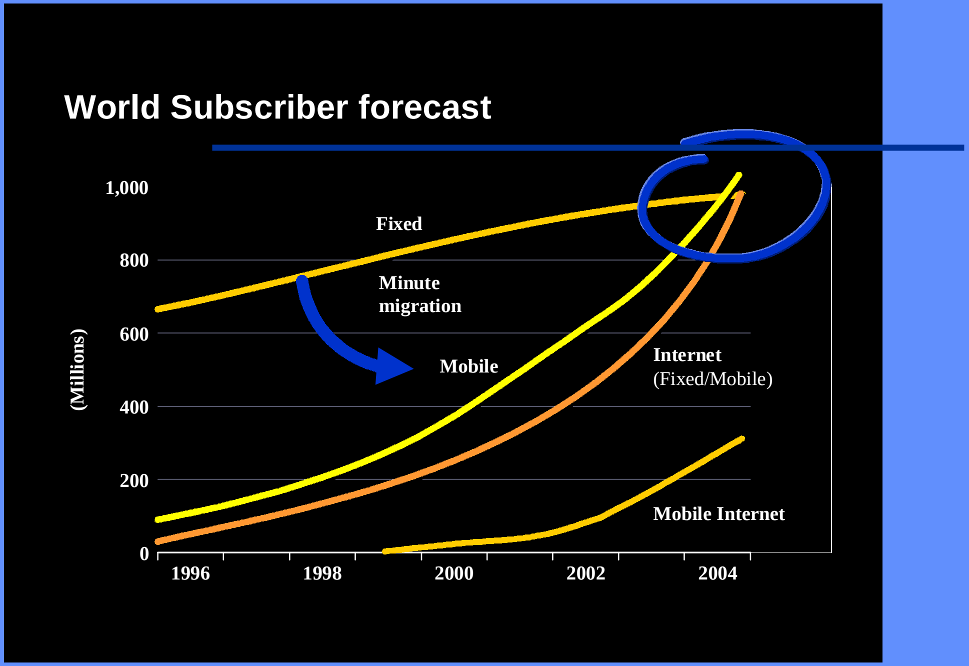

Some Internet2 QoS objectives

- Provide predictable QoS end-to-end
- Keep architectural simplicity in the backbone
  - Class of service
  - Aggregation
- Allow for more elaborate admission control where needed
  - e.g., over wireless links

Example: Internet resource provisioning for telephony

- A virtual telephone network within the Internet (VLL, VPN)
- Well provisioned
- Renegotiated periodically, in advance
- Time dependent (e.g., time of day, days of week)
- Per-call QoS signaling not appropriate in the backbone, but mandatory in wireless access.

Our BB Architecture

- unidirectional virtual leased lines, VLLs
- automatic setup of DiffServ traffic conditioners
- one broker per routing domain
- brokers know the static resources of their domain (through OSPF and SNMP)

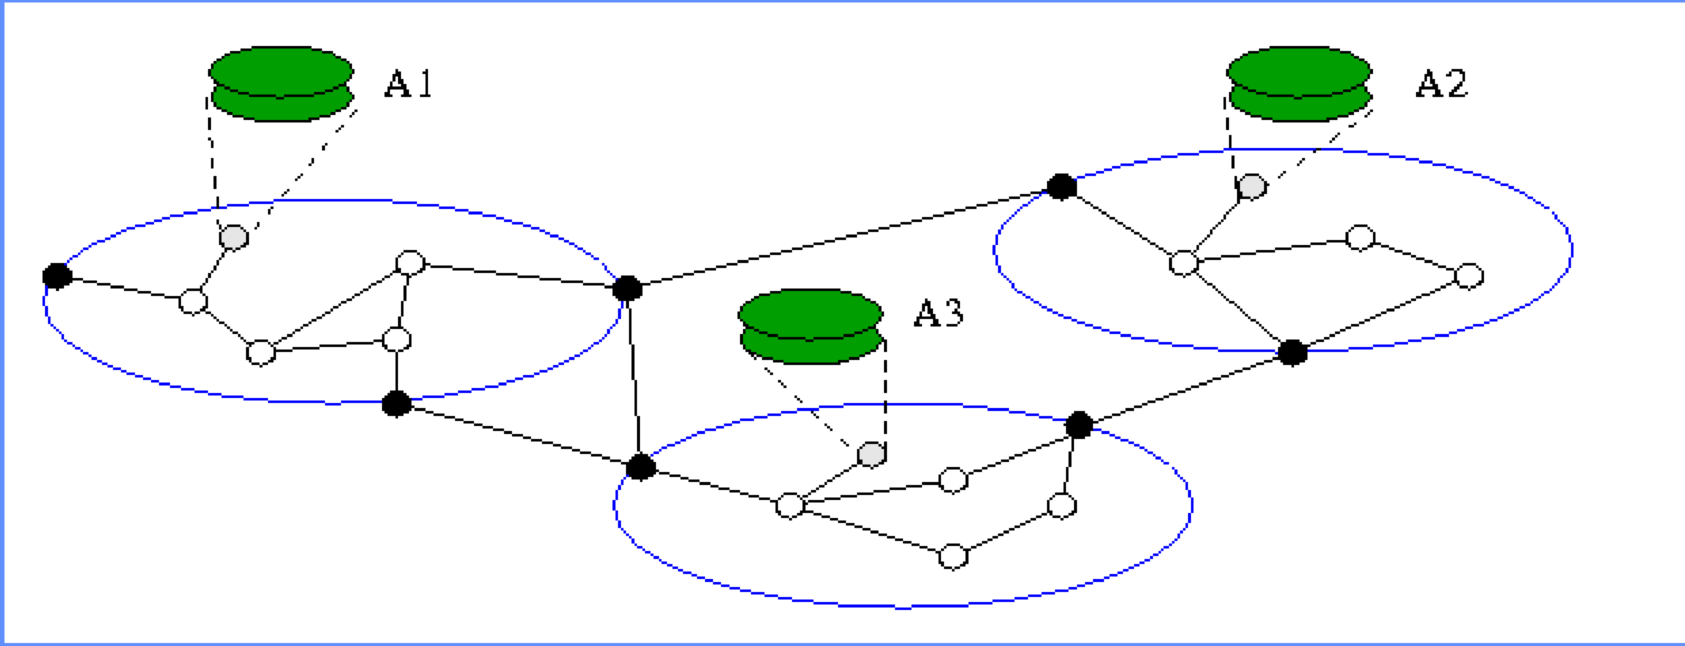

Admission Control for VLLs

- Objective: reserve resources over links that are actually going to be used
  - to obtain better utilization
- Follow intra-domain unicast routing
  - to predict the path
  - to find the border point to next domain
- Independent packet filter at the ingress

Admission Requests for VLLs

Essential content

- The requested bandwidth
- An ingress address (prefix)
  – determines where the reservation starts
- A destination address (prefix)
  – determines where the reservation ends
- Duration (in case of advance admission)
  – start and stop time

Advance admission decision

- Fitting boxes into a two-dimensional time-bandwidth diagram

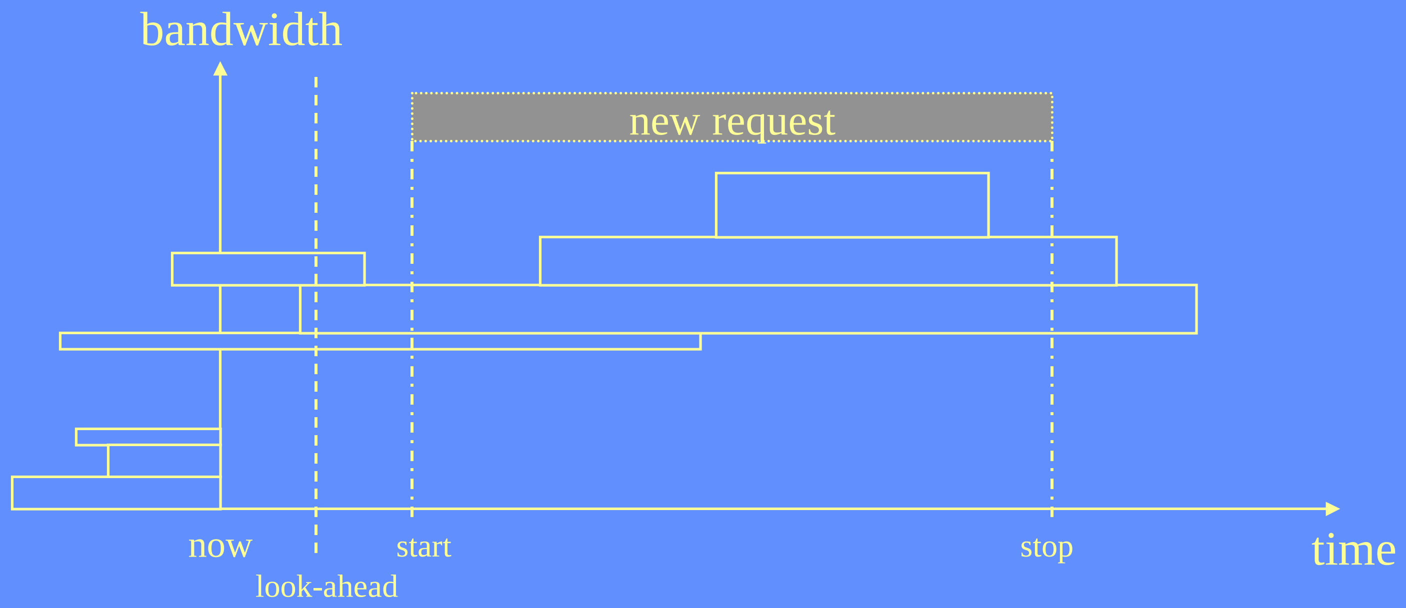

We provide a web-interface to make admission requests to our BB.

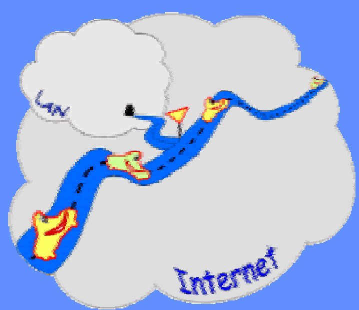

In addition, applications can use a TCP-based protocol to make admission requests.

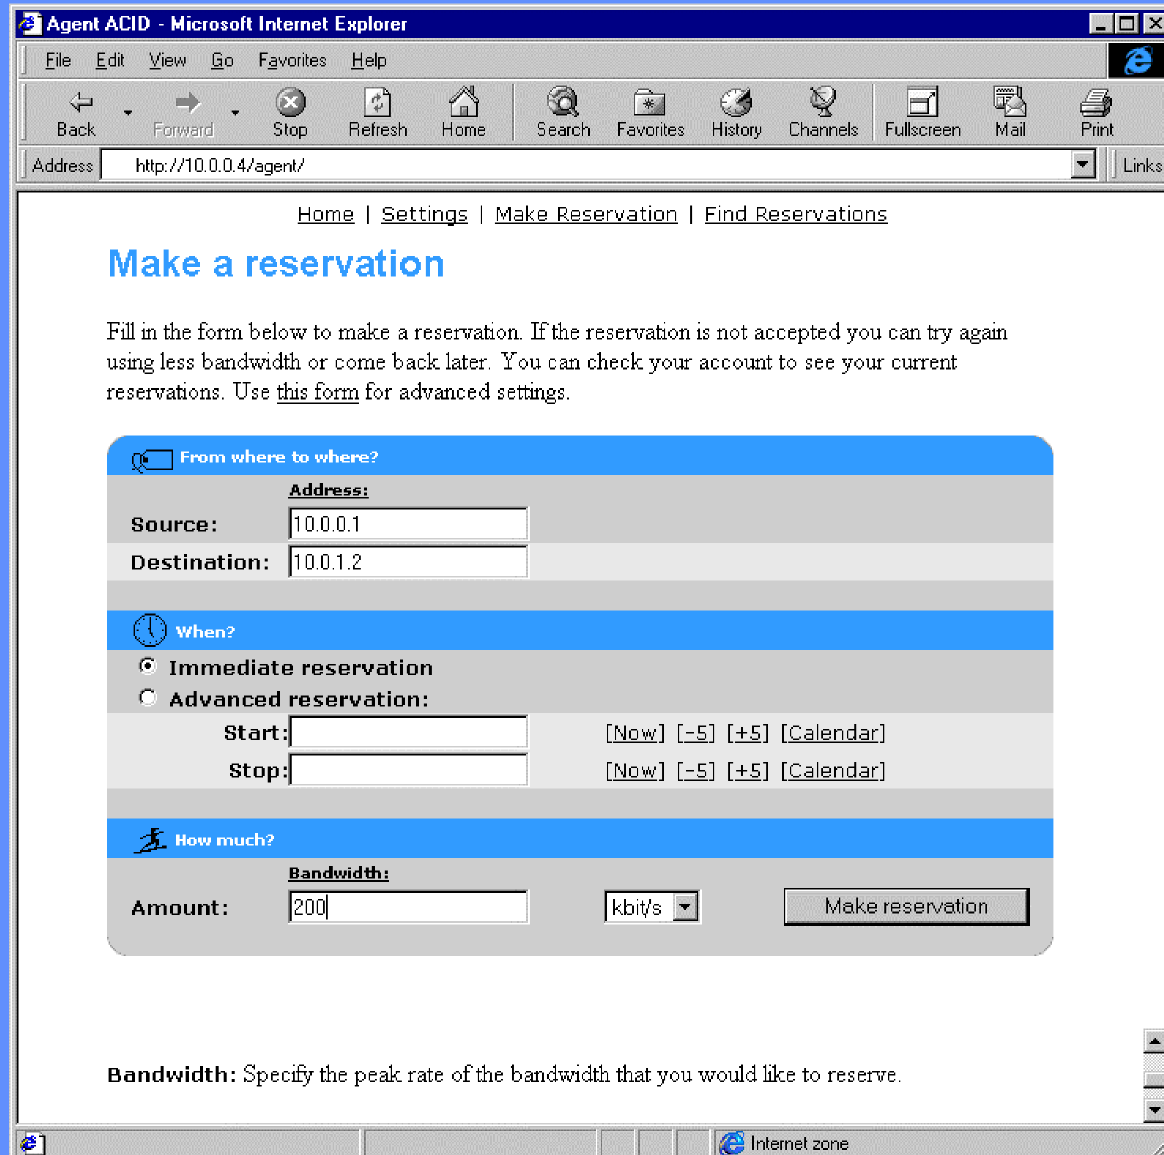

Sending admission requests to a BB

- General Interface: TCP-based protocol
  - e.g., COPS
- Other signaling will be translated to this protocol, e.g.,
  - Web interface (VLL/trunk reservations)
  - RSVP, at boundary routers (micro-flows)
  - Other BBs

How did we evaluate our BB?

A . By building a complete prototype

- operating in small domain of real DiffServ capable routers (a lab network)
- with real traffic

B . By evaluating admission control cost in a real network domain

- for a medium-sized link-state routing domain (the LTU campus network)

C . By evaluating admission control cost over multiple routing domains (work in progress)

State aggregation for Advance Reservations

- The amount of admission control state depends on static topology only, (independently of the number of reservations).
- less than 100 KB per control point (e.g., per hop) to support one month, slotted in intervals of five minutes.

- Advance admission decisions are fast
  - e.g., less than 7 microseconds per hop

Multi-hop admission performance (path lookup, ~6 hops)

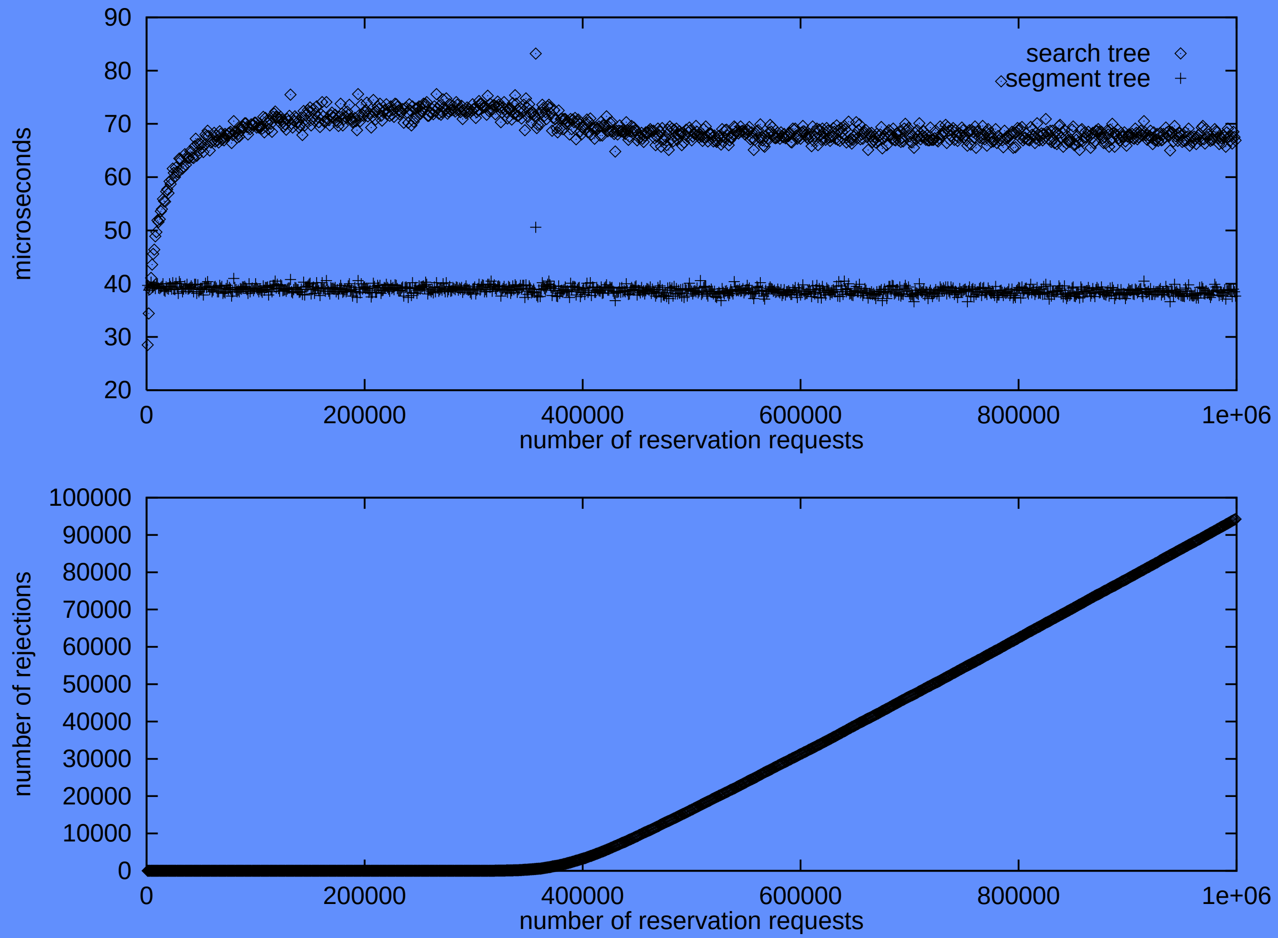

Operability event test-bed

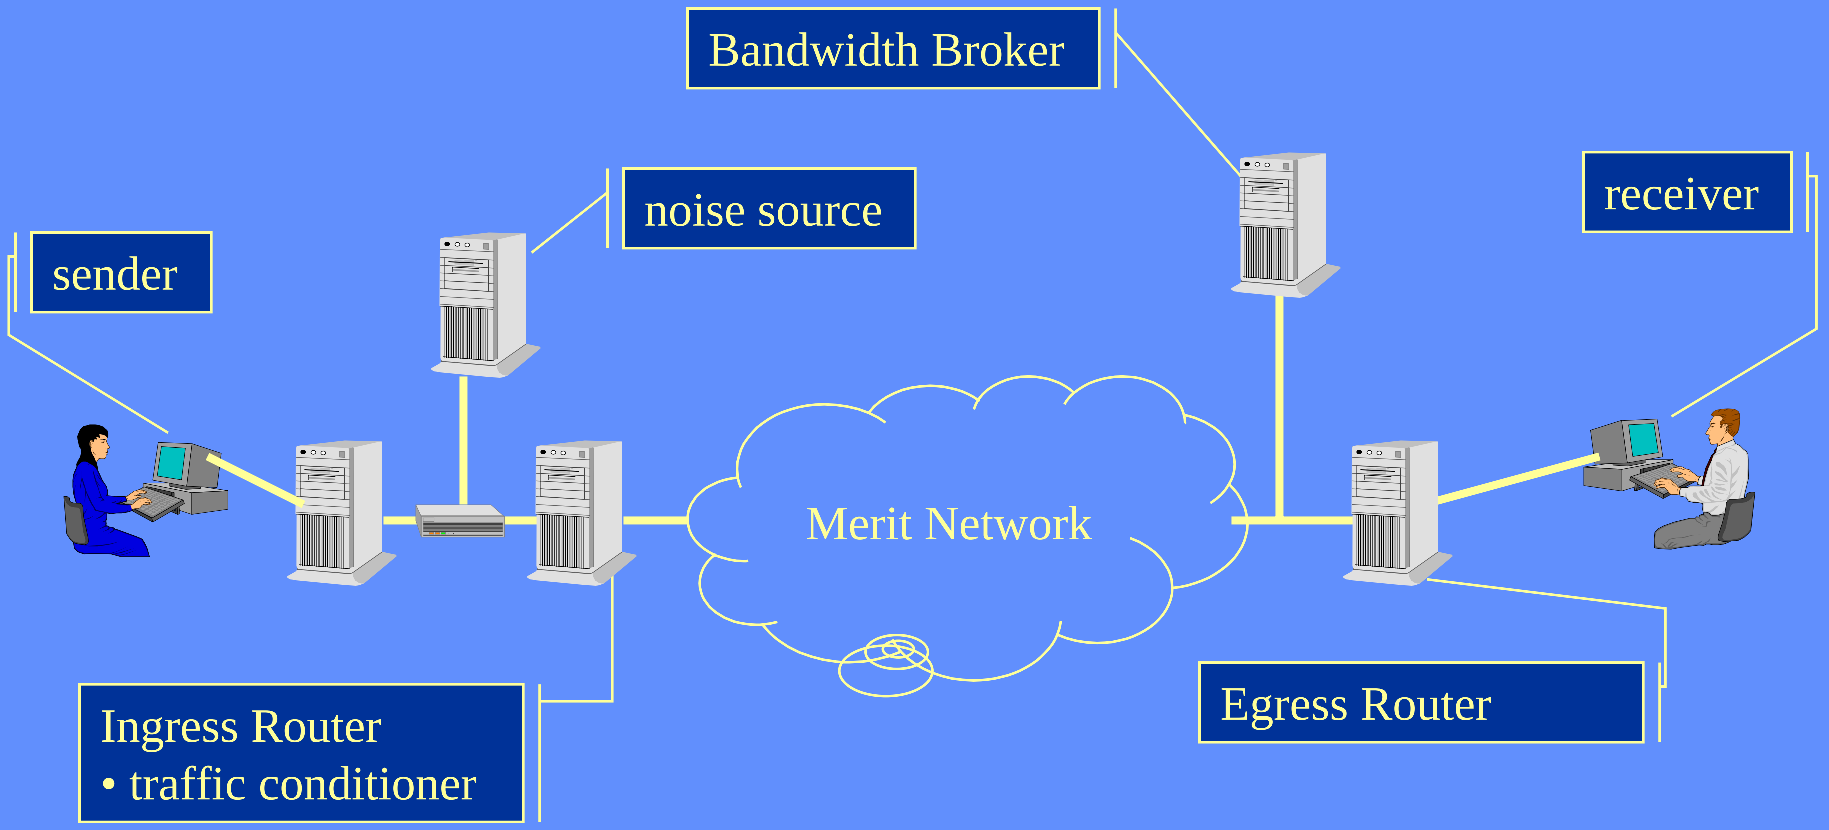

Edge routers

- PC-based boundary routers
  - FreeBSD
  - Traffic conditioner
  - ALTQ
- Configured with SNMP/COPS
- Monitored by lots of visualization tools

Inter-domain VLL Aggregation

- BGP for figuring out next AS and ingress point.
- When reserving with a neighboring agent, one reservation can be made per destination
  - independently of from where traffic originated,
  - for traffic using the same downstream ingress,
- Thus, agents can aggregate several reservations into one reservation with their neighboring agent.
- Trade-off: Signaling vs. over-allocation

Funnels

- Each reservation can be seen as a funnel covering a specific destination.
- An end-to-end reservation can be constituted from several consecutive funnels
- Funnels form a **sink tree** towards the destination

Aggregation towards a destination domain

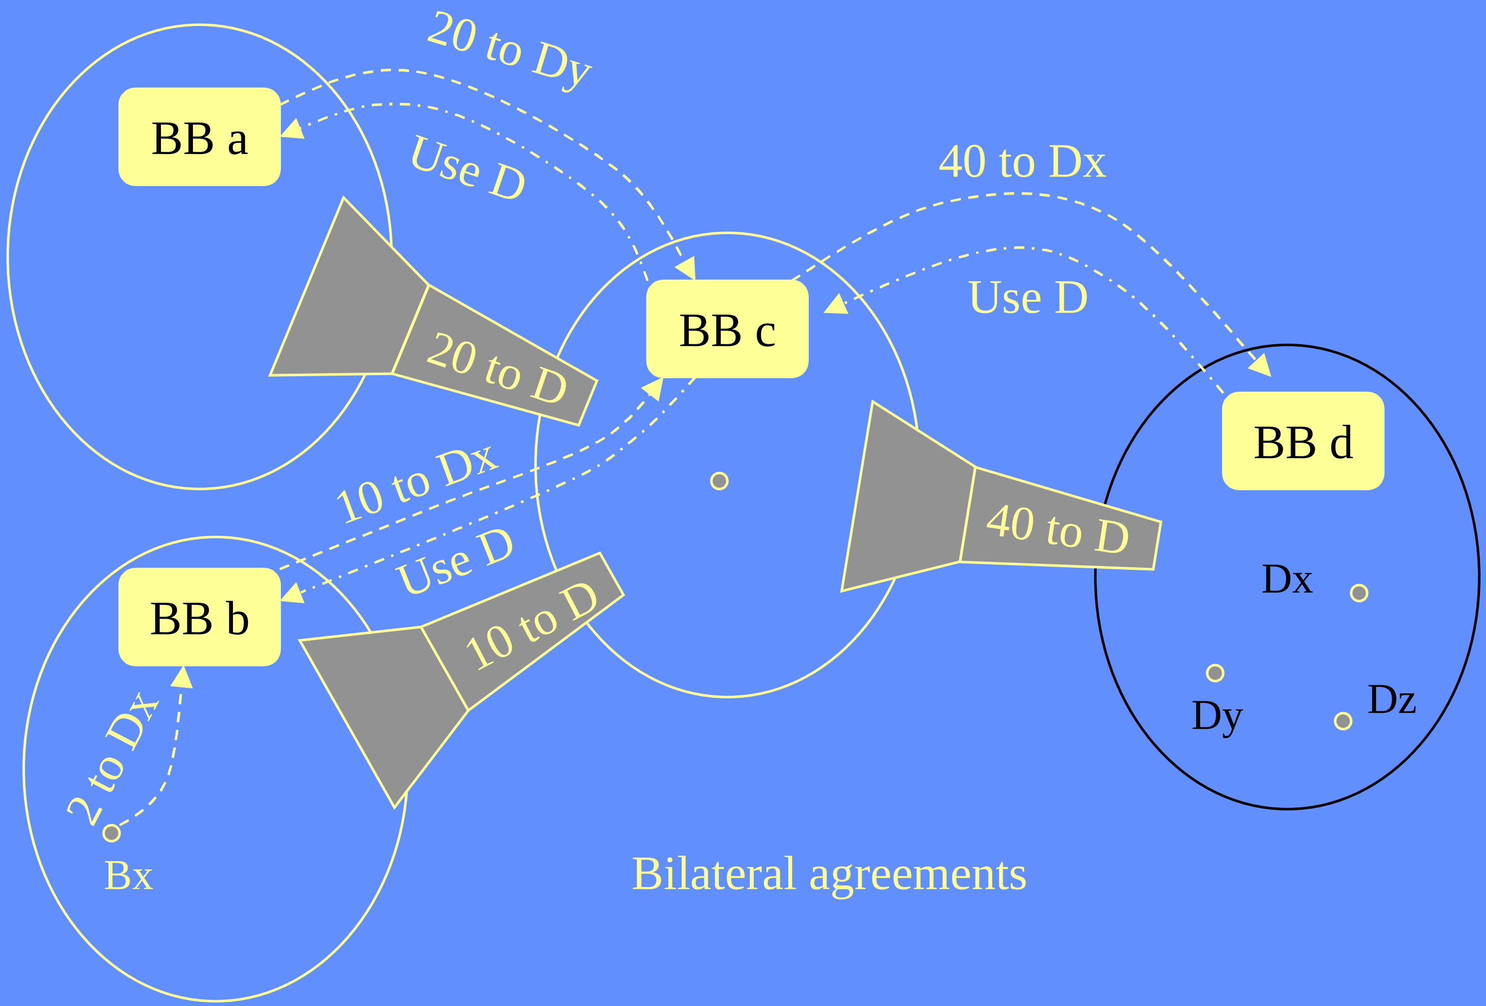

Bilateral agreements

Prefix generalization

- One aggregate reservation can be used for all destinations within a well provisioned domain (e.g., a LAN).
- The destinations can be represented by an address prefix (subnet prefix)
- Only the agent in charge of the destination domain can judge when such aggregation is appropriate.

Trade offs

- Pre-reservations vs on-demand signaling
  – admission requests stop propagating where sufficient resources are reserved further on
- No fan-out (pure sink trees) enables aggregation
  – fan-out in destination domain (if desirable) should be provided by separate mechanism
- Scales as O(n) in worst case, where n is the number of destination prefixes.

Important Notes

- The same inter-domain protocol as for simple end-to-end forwarding can be used.
  - Aggregation goes on internally in the BB
- The inter-domain protocol should support returning modified admission requests
  - to allow for prefix generalization, to suggest a different time or bw that might be available, etc.
- The degree of gambling (over-allocation) is a policy, e.g.,
  - can be peak-rate, or can be measurement based

QoS going wireless

- wide-area wireless (e.g., GSM, UMTS)
  - bandwidth very expensive and scarce
- local-area wireless (e.g., WLAN, Bluetooth)
  - bandwidth cheap but still somewhat scarce

Wide-area wireless

- Insufficient parameter set for admission control over radio bearer
  - 3GPP defines a large number of attributes (>10), e.g. bit rate,
    - delay, BER, delivery of erroneous packets
- Classes
  - conversational: delay ~0, BER >0
  - streaming: delay >0, BER >0
  - interactive: delay >0 or ~0, BER = 0
  - background: delay >0, BER = 0

Bit error rate

- For spectral efficiency, the wireless link will run with high BER
  - e.g., as high as the voice encoders can tolerate
- Lower BER can be provided to some traffic
  - by link layer FEC and/or ARQ
- This should be hinted in admission requests

Prerequisites

- A revision of the IP stack
  - Application: sensitive and insensitive data
  - Transport: UDP lite - Transport layer that checksums transport and application headers, but allows bit-errors in the payload
  - Network: Carry sensitivity information to the links
  - Header compression for wireless links
    - IETF, robhc

Admission requests

- Intserv RSPECs
  - Do not contain a sufficiently large parameter set for wireless
  - A new kind of rspec could be defined that is suitable for wireless links
- Local USE of RSVP?
  - Carrying admission requests to local gateway and then to radio resource manager

Subnet Bandwidth Managers (SBM)

- link-layer specific
- already designed for wide area radio bearers
  - using extended parameter set
- will be provided for WLANs
- How should BBs and SBMs interact?
  - Clients → SBM → BB for resources?
  - Clients using RSVP (locally)?

Conclusion

- Provide advance planning and enable for aggregated resource management in the backbone (learn from BGP).
- Revise the IP stack for wireless
- Provide extended fine granular admission control for wireless

Oh, that's a piece of cake!

Our BB work: http://www.cdt.luth.se/~olov/publications