Who Talks to Whom: Using BGP Data for Scaling Interdomain Resource Reservation

Ping Pan and *Henning Schulzrinne*

Columbia University

ISMA Workshop – Leiden, Oct. 2002

Columbia Computer Science

Overview

- Reservation scaling
  - CW: “per-flow reservations don’t scale”
  - → true only if every flow were to reserve
  - may be true for sub-optimal implementations...
- Based on traffic measurements with BGP-based prefix and AS mapping
- looked at all protocols, since too little UDP to be representative

Reservation scaling

- Reserve for sink tree, not source-destination pairs
  - all traffic towards a certain network destination
  - provider-level reservations
    - within backbone
    - high-bandwidth and static trunks (but not necessarily MPLS...)
  - application-level reservations
    - managed among end hosts
    - small bandwidth and very dynamic flows
- Separate intra- and inter-domain reservations
- Example protocol design: BGRP

Different growth curves

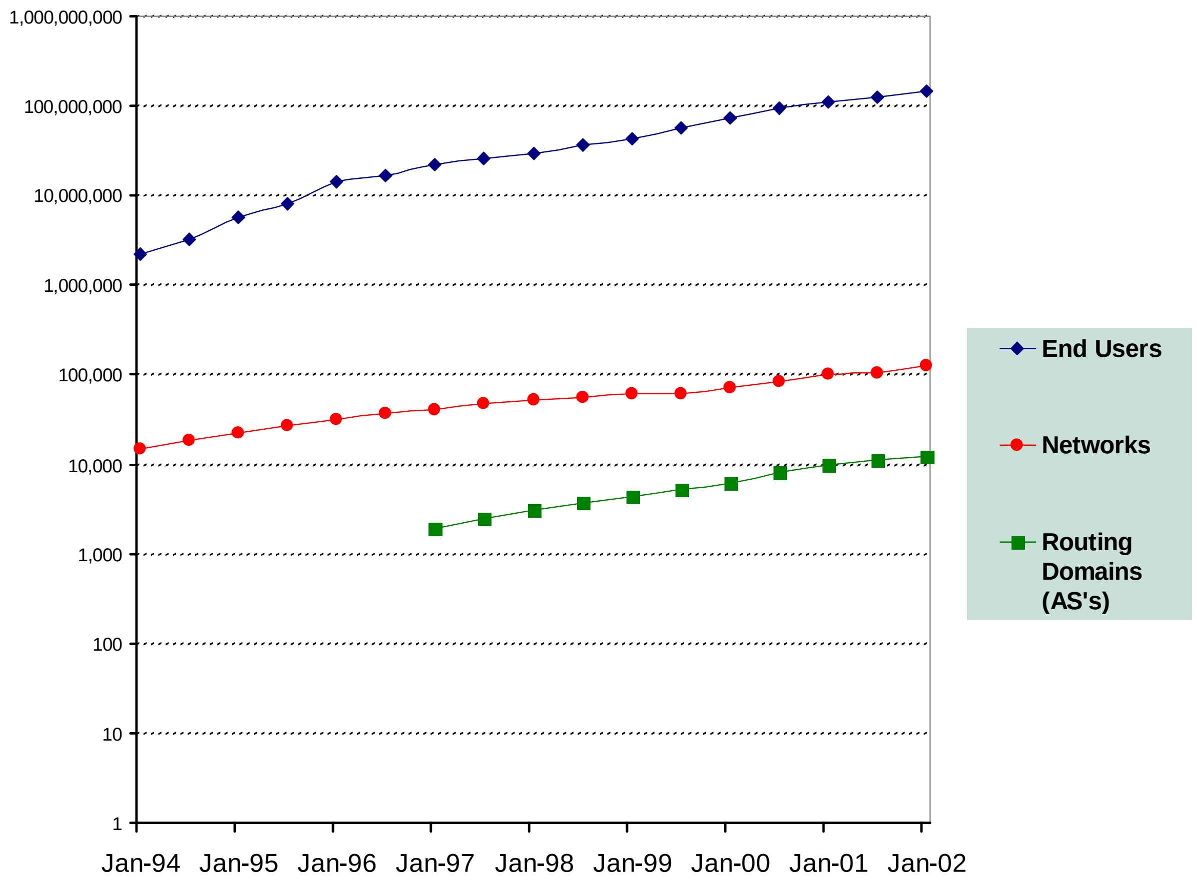

Columbia Computer Science

Estimating the max. number of reservations

- Collected 90-second traffic traces
  - June 1, 1999
- MAE West NAP
- 3 million IP packet headers
- AS count is low due to short window:
  - were about 5,000 AS, 60 network prefixes then
  - May 1999:
    - 4,908 unique source AS's
    - 5,001 unique destination AS's and
    - 7,900,362 pairs (out of 25 million)

Columbia Computer Science

A traffic snap shot on a backbone link

Granularity	flow discriminators	potential flows
application	source address, port	143,243
	dest. address, port, proto	208,559
	5-tuple	339,245
IP host	source address	56,935
	destination address	40,538
	source/destination pairs	131,009
network	source network	13,917
	destination network	20,887
	source-destination pairs	79,786
AS	source AS	2,244
	destination AS	2,891
	source-destination pair	20,857

How many flows need reservation?

- Thin flows are unlikely to need resource reservations
- Try to compute upper bound on likely reservation candidates in one backbone router
- Eight packet header traces at MAE-West
  - three hours apart on June 1, 1999
  - 90 seconds each, 33 million packets
  - bytes for each
    - pair of source/destination route prefix
    - destination route prefix

Distribution of connection by bandwidth

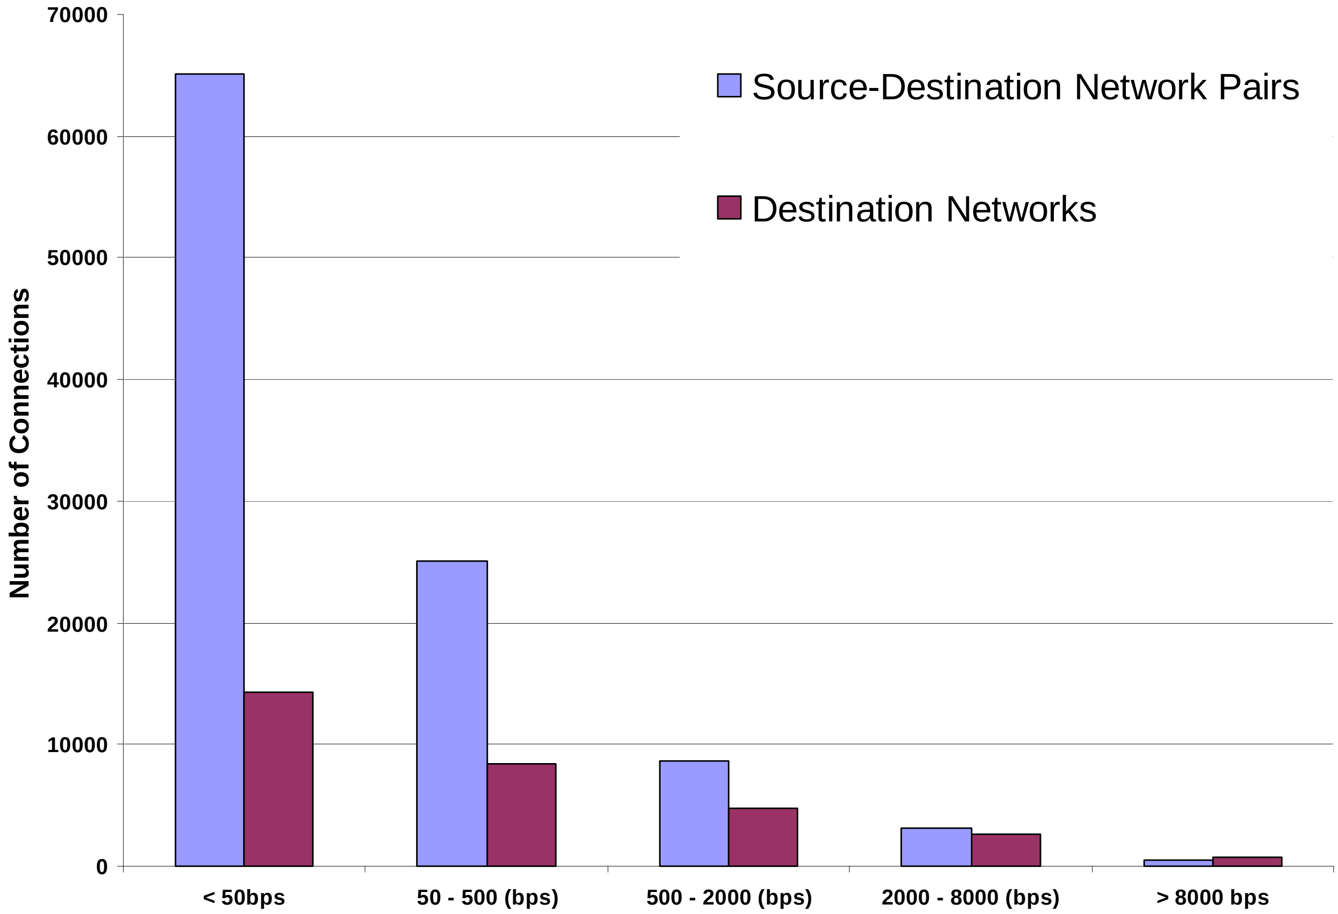

Columbia Computer Science

The (src-dest / destination) ratio

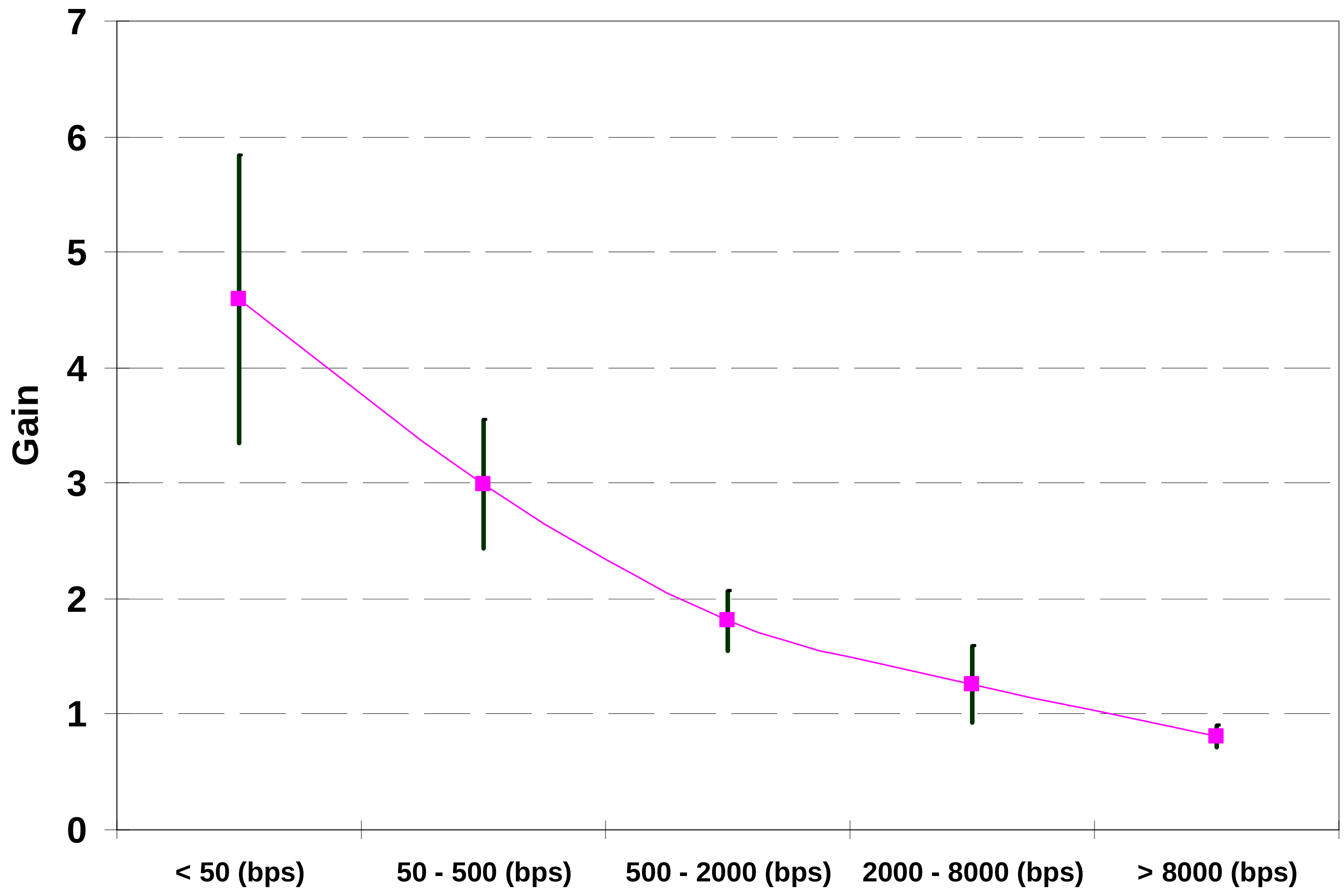

Results

- Most packets belong to small flows:
  - 63.5% for source-destination pairs
  - 46.2% for destination-only
- only 3.5% (3,261) of the source-destination pairs and 10.9% (1,296) of destinations have average bit rate over 2000 b/s
  - thus, easily handled by per-flow reservation
- more above-8000 b/s destination-only flows than source-destination flows
  - large web servers?

Columbia Computer Science

Aside: Estimating the number of flows

- In 2000,
  - 4,998 bio. minutes ~ 500 bio calls/year
    - local (80%), intrastate/interstate toll
  - 15,848 calls/second
    - not correct → assumes equal distribution
- AT&T 1999: 328 mio calls/day
  - 3,800/second

The Hierarchical Reservation Model

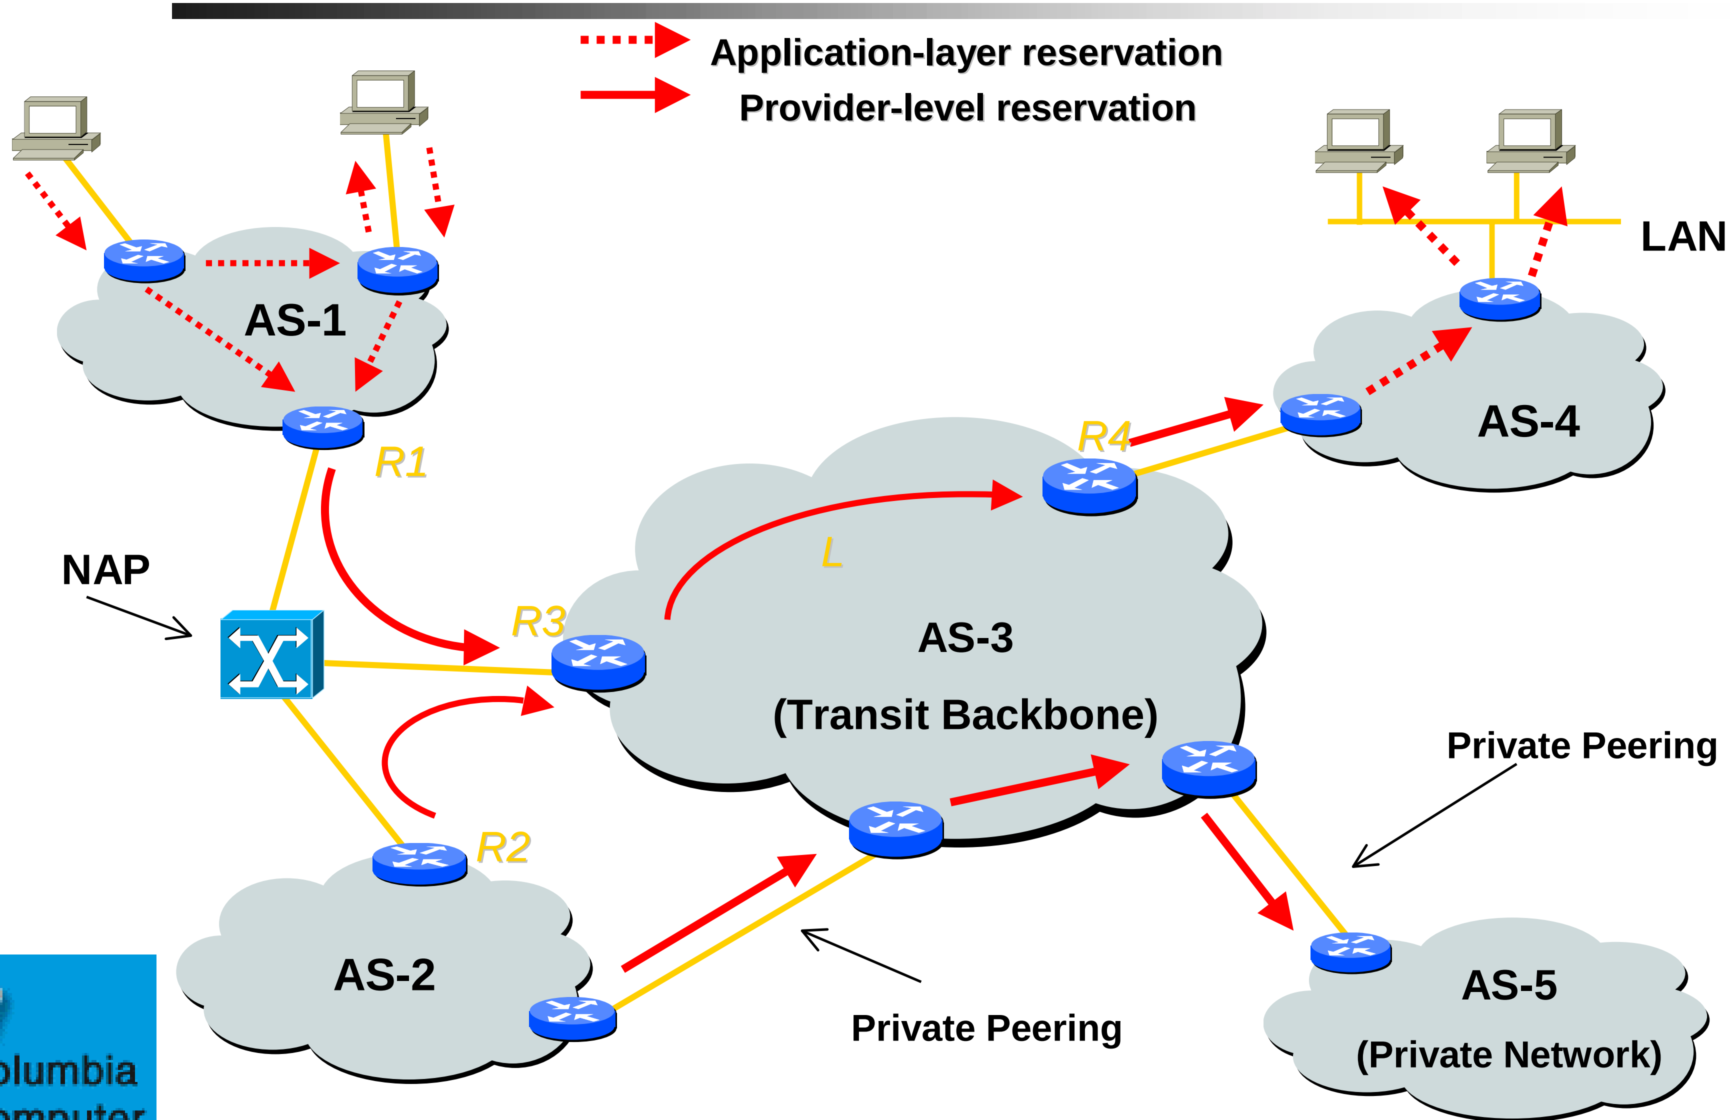

Columbia Computer Science

Conclusion

- Communications relationships
  - granularity and “completeness”
  - flow distribution
- Questions:
  - traffic seems to have changed qualitatively
    - more consumer broadband, P2P
    - see “Understanding Internet Traffic Streams”
  - protocol behavior
    - funnel-behavior may differ for QoS candidates
    - e.g., large PSTN gateways
    - but no funnel for (e.g.) media servers