An Experimental Study of the Skype Peer-to-Peer VoIP System

Saikat Guha (Cornell)
Neil Daswani (Google)
Ravi Jain (Google)

IPTPS 2006

About Skype

- Voice over IP (VoIP)
- 50 million users
- Valued at $2.6 billion
- Proprietary

About Skype

"Internet Telephony that Just Works"

- Adaptive voice quality

  modem, broadband, T1, ...
- Works in any network topology

  one or more NATs, Firewalls, ...

Outrageous Opinion[1] #1

- NAT Traversal in Skype:
  - Level 0: Initiator NAT'ed
  - Level 1: Recipient NAT'ed
  - Level 2: Both NAT'ed (well-behaved NATs)
  - Level 3: Both NAT'ed (broken NATs)

Outrageous Opinion

NAT Traversal is essential for P2P

- How to pick NAT Traversal complexity, and make it scale?

[1] does not, necessarily, reflect the views of all co-authors or employers

NAT Traversal in Skype

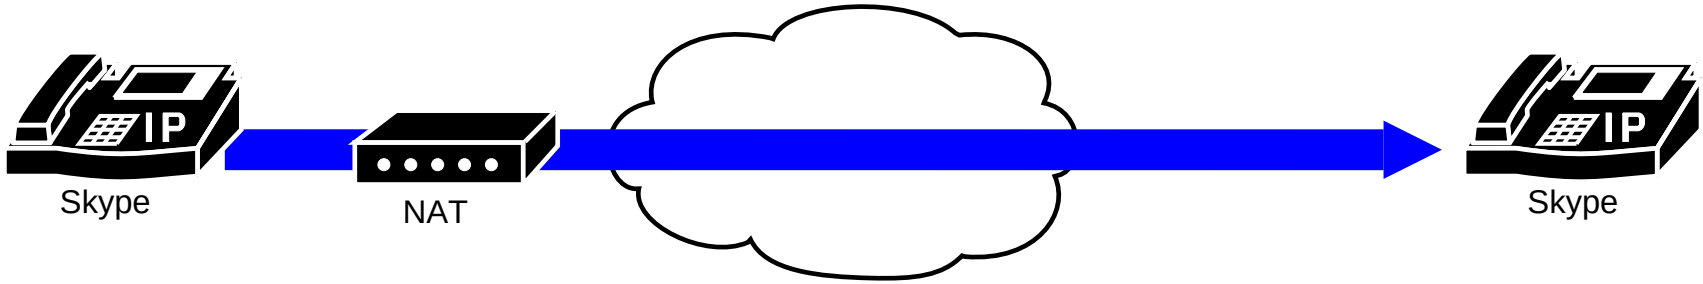

Level 0: Initiator NAT'ed

Solution: Don't embed IP address in payload

Apps: Most old software, almost all new software

NAT Traversal in Skype

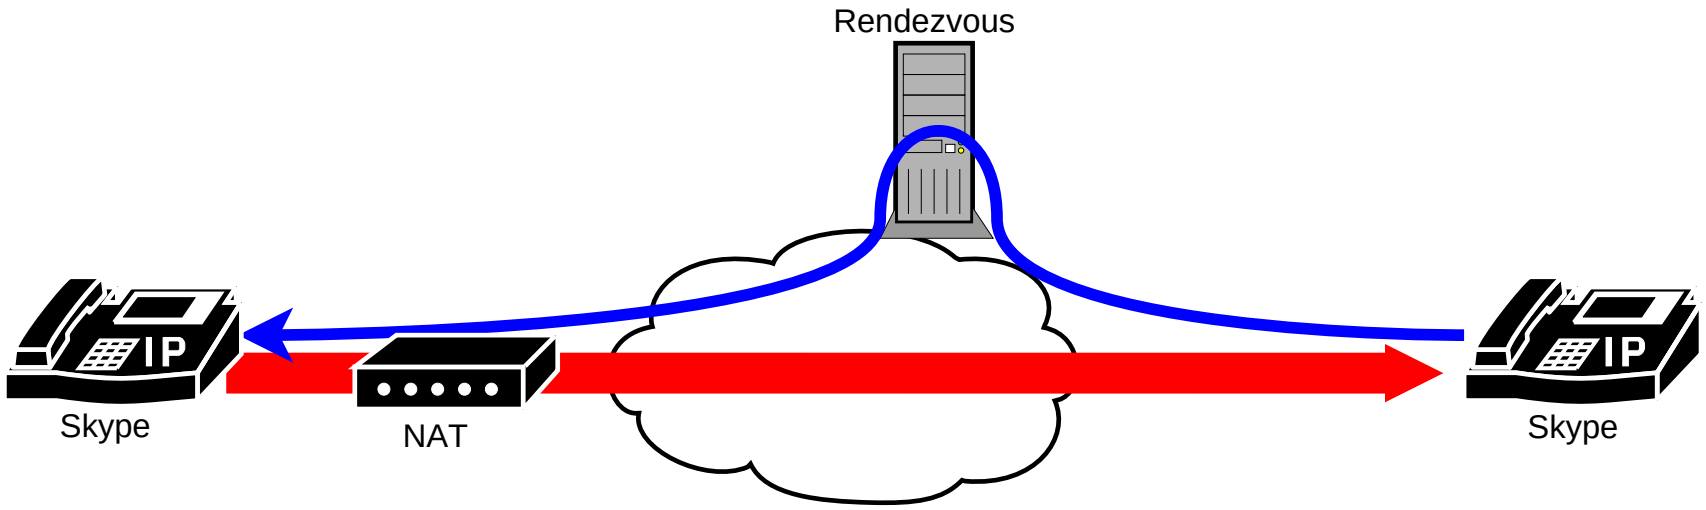

Level 1: Recipient NAT'ed

Solution: Use Rendezvous Service

Apps: Bittorrent, MSN, Yahoo, Skype, ...

NAT Traversal in Skype

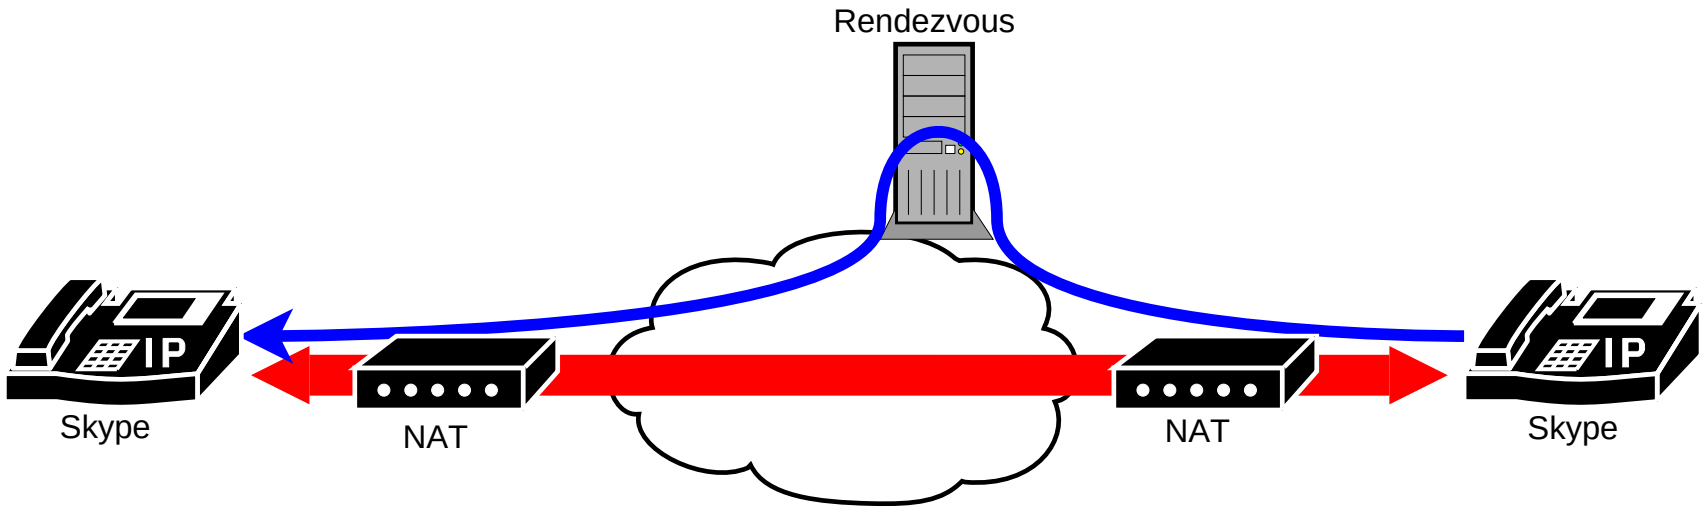

Level 2: Both NAT'ed (well-behaved NATs)

Solution: Use STUN (UDP) or STUNT (TCP)
Apps: MSN, Yahoo, Skype, ...

NAT Traversal in Skype

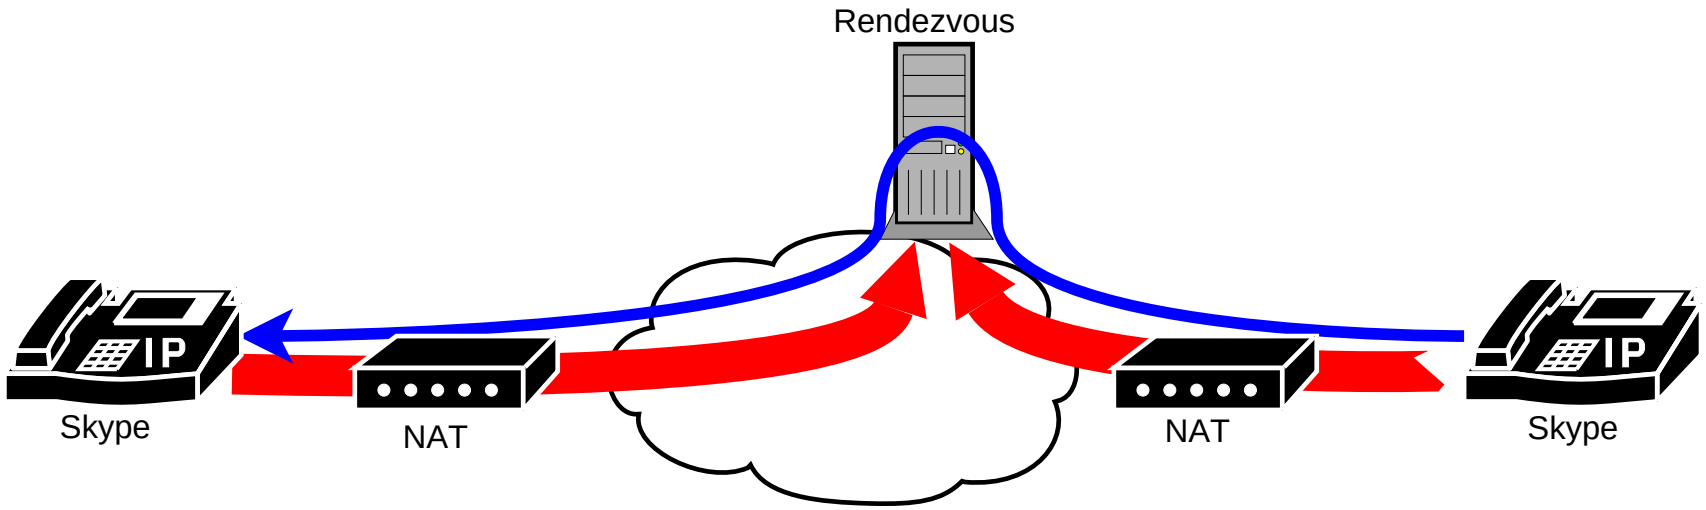

Level 3: Both NAT'ed (broken NATs)

Solution: Use TURN

Apps: MSN (limited), Yahoo (limited), Skype

NAT Traversal in Skype

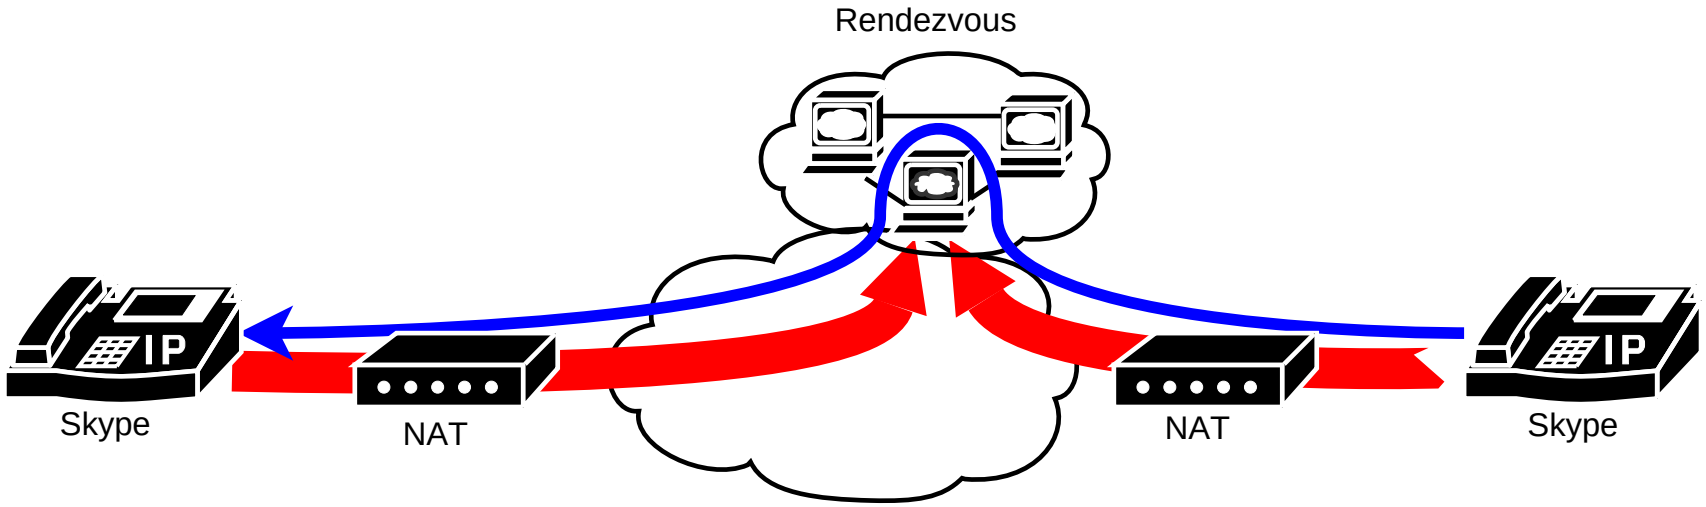

Level 3: Both NAT'ed (broken NATs)

Solution: Use TURN + P2P

Apps: Skype

Contributions

1. Properties of Skype's P2P Rendezvous-Relay Service
   - load (low), structure (yes), churn (low) ...
2. Skype user behavior
   - call-length (long), file-transfer (small), online activity (predictable) ...

Skype's P2P Network

- Inherent structure: [Baset et al. INFOCOM '06]
  - **Supernodes (SN):** no NAT, spare bandwidth
  - **Ordinary nodes (ON):** associate with one (or a few) supernodes. Don't contribute to overlay.

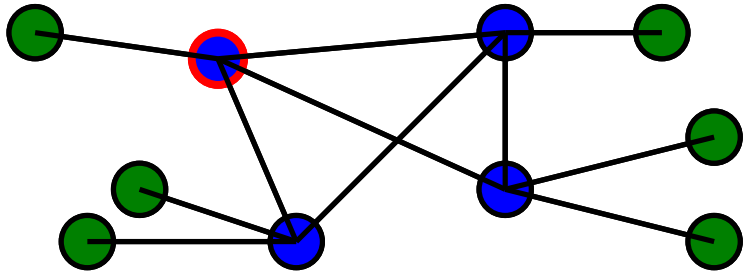

- Skype network: FastTrack (likely)
- Experiment: Capture traffic for a SN
  - at Cornell, 4.5 months, ∼13GB
  - Caveat: everything "encrypted", only one SN

Skype's P2P Network

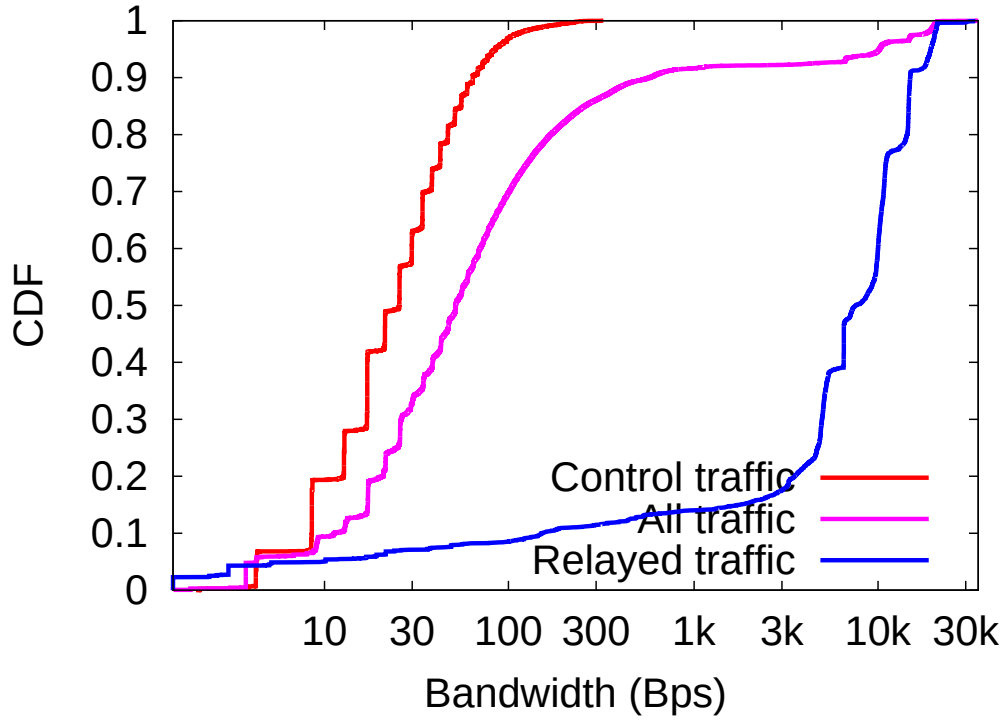

Supernode: Low Network Load

- Bandwidth: 51 Bps (median), 7.5 kBps (median relayed), 34 kBps (peak)
- Control traffic (∼75%), Relayed traffic (∼10%)

Skype's P2P Network

- Experiment: Estimate no. of active nodes
- Supernodes (SN):
  - Discover SN: crawl client SN cache
  - App-level ping: replay "hello"-packet
  - 30%–40% of SNs active, 250k found
  - Caveat: DHCP assigned SN address
- All Nodes (ON+SN):
  - Record number of active nodes
  - $\lesssim$4 million simultaneous users
  - Caveat: reported by proprietary client

Outrageous Opinion[1] #2

- Rough estimate: (just network, not CPU)
  - $\sim$1–2 GBps median relay-traffic
  - $\lesssim$20 well-provisioned boxes at PoPs
  - $\sim$ \$10 million per year
- Supernodes mostly at universities (EU, Asia, US)

Outrageous Opinion

Skype freeloads on university bandwidth

- Is there an ISP-friendly commercial model for P2P?

Skype's P2P Network

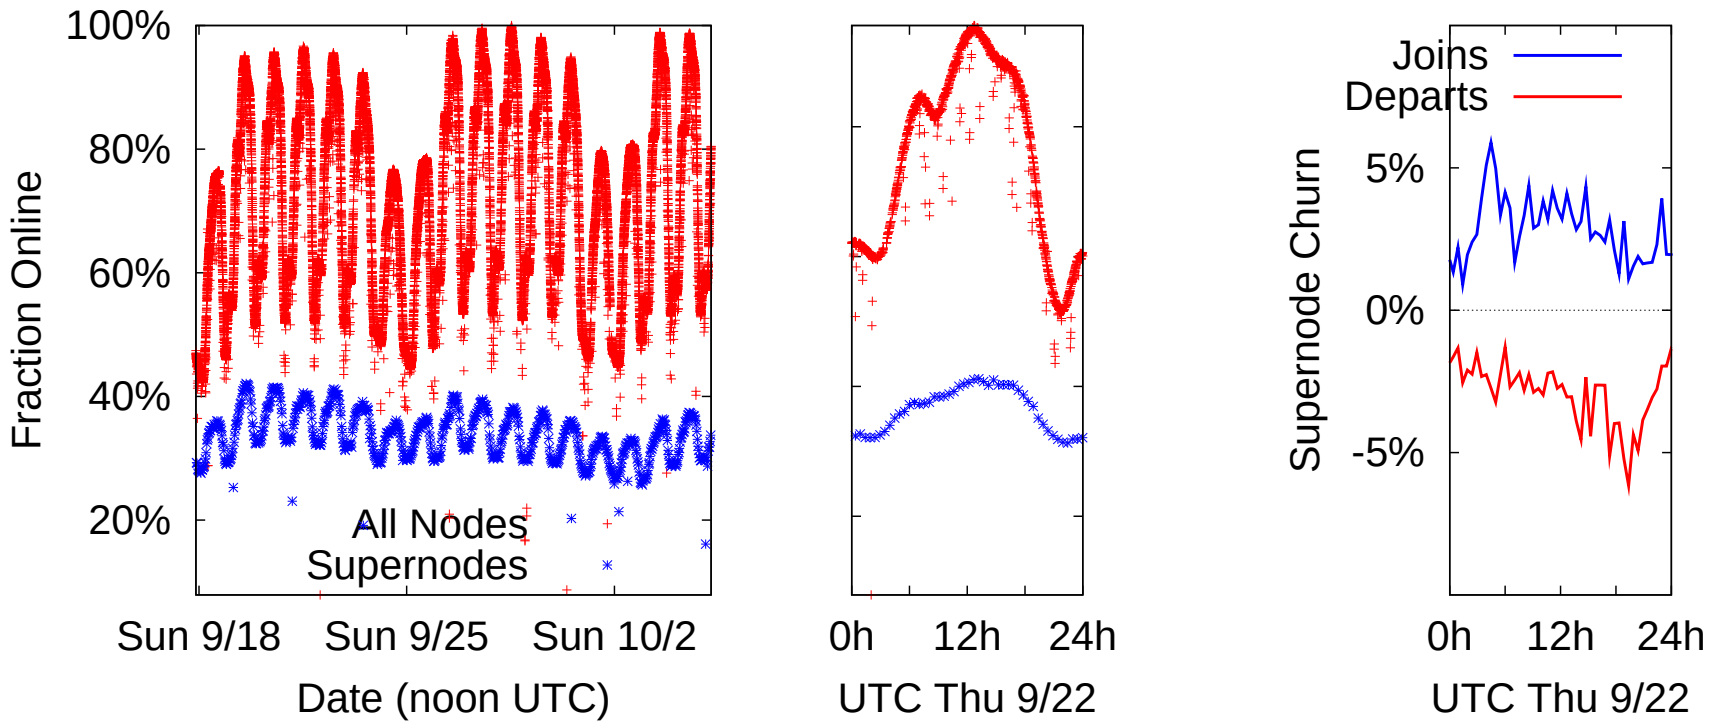

Supernodes: Diurnal, Work-week Patterns

- **Supernodes:** low churn, $\lesssim$5% turnover (over 30-min)

Outrageous Opinion[1] #3

- Handling Churn
  - a blood sport
  - Skype, perhaps intentionally, controls churn (based on NAT, bandwidth, maybe session history, ...)

Outrageous Opinion

Churn resistance is overrated. Control churn.

- Is there a price-performance tradeoff between controlling, and handling churn?

Skype Users

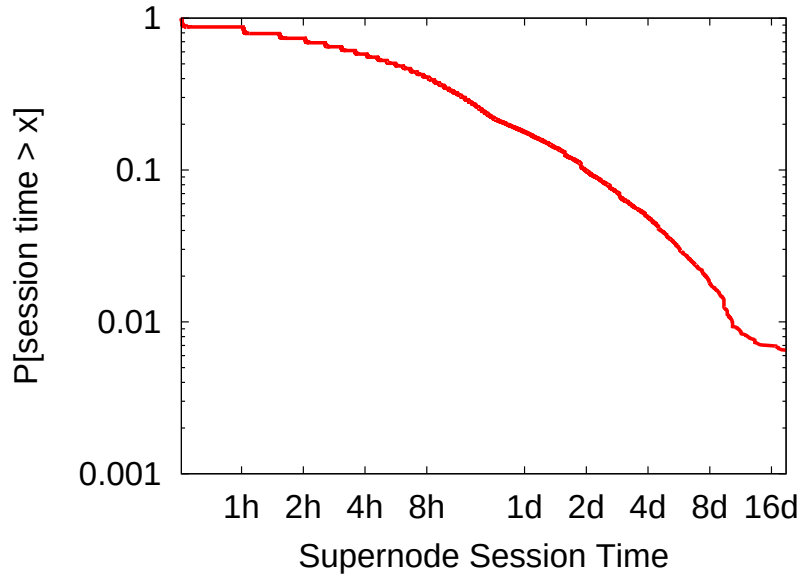

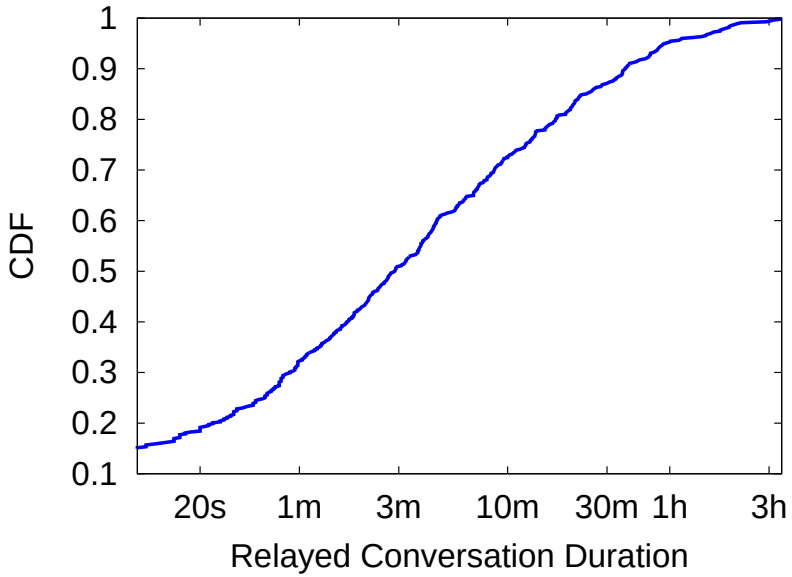

Skype users: differ from file-sharing, phone users

- **Session Length:** heavy-tailed, 5.5h median
- **File-Transfer:** 346kB median
- **Call-Length:** 12m 53s mean, 3h 26m max

Summary

1. Properties of Skype's P2P Rendezvous-Relay Service
   - Low Load: NAT traversal, many supernodes
   - Stable: exploits heterogeneity
   - ISP-unfriendly, taxes universities
2. Skype users
   - Different from file-sharing users (longer session lengths, smaller files transferred)
   - Different from phone users (longer calls)
3. Data available on request

Discussion

For a peer-to-peer application:

- Complexity and scalability of NAT traversal
- ISP-friendly commercial model
- Tradeoffs between controlling and handling churn

`www.cs.cornell.edu/~saikat/skype/`

Outrageous Opinions[2]

- NAT Traversal is essential for P2P
- Skype freeloads on university bandwidth
- Churn resistance is overrated. Control churn.

`www.cs.cornell.edu/~saikat/skype/`

[2] does not, necessarily, reflect the opinion of my co-authors, our employers, or anyone else.

Supernode Distribution

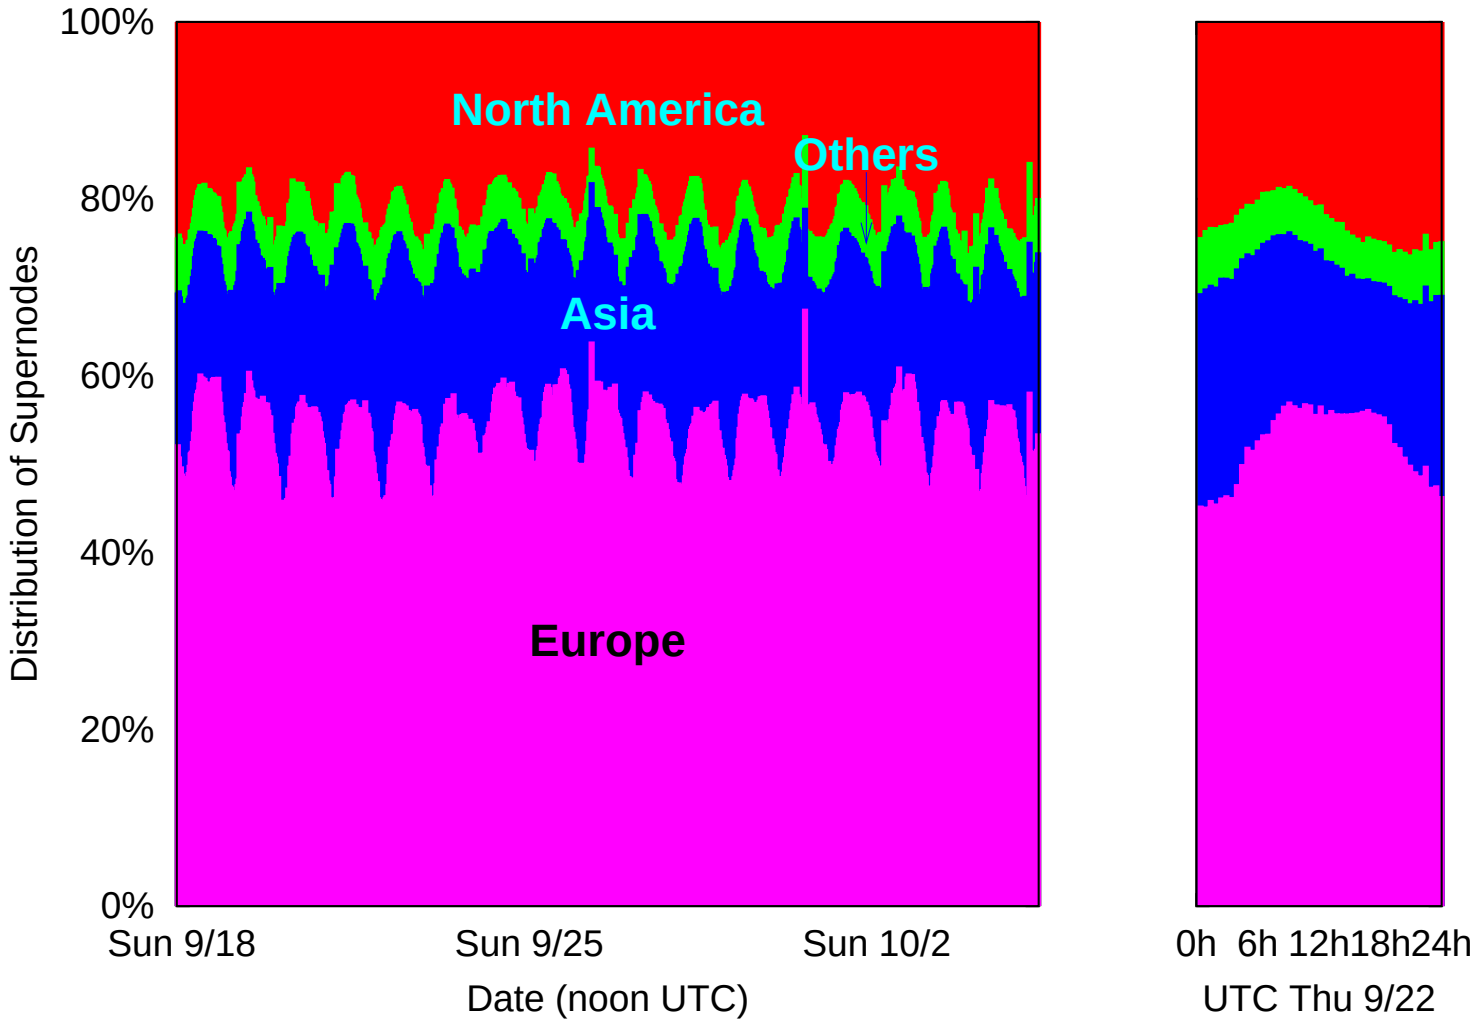

Filesize Distribution