Ubiquitous Streaming Media and Telephony

Henning Schulzrinne
Internet Real-Time Lab (IRT)
Dept. of Computer Science
Columbia University
New York, New York
schulzrinne@cs.columbia.edu

Bell Labs, Murray Hill)

July 14, 2000

Joint work with Jonathan Lennox, Dave Marples (Telcordia), Stan Moyer (Telcordia), Maria Papadopouli, Jonathan Rosenberg, and Elin Wedlund

Overview

- Internet telephony vs. streaming media
- Internet telephony signaling
- signaling and RPC
- signaling and presence
- networked appliances
- mobility – more than just wireless terminals
- alternate wireless architectures

Internet Telephony vs. Streaming Media

	streaming	VoIP
transport	RTP	RTP
delay	< 2 s	< 0.2 s
setup	RTSP	SIP
description	SDP	SDP
end system	S$\leftrightarrow$C	C$\leftrightarrow$C

intersections: voicemail, IVR, MP3 telephone, …

SIP: Session Initiation Protocol

IETF-standardized *peer-to-peer* signaling protocol (RFC 2543):

- locate user given email-style address
- set up session
- (re)-negotiate session parameters
- manual and automatic forwarding (“name/number mapping”)
- *personal mobility* ➠ different terminal, same identifier
- “forking” of calls: one call, multiple destinations
- terminate and transfer calls

SIP features

- provides call control (hold, forward, transfer, media changes, . . .)
- leverages web infrastructure: security, “cgi-bin”, electronic payments, PICS, cookies, . . .
- web-oriented: return HTML pages (“web IVR”)
- network-protocol independent: UDP or TCP (or AAL5 or X.25)
- easily extends to presence information (“buddy lists”) and event notification

SIP architecture: inbound and outbound proxy

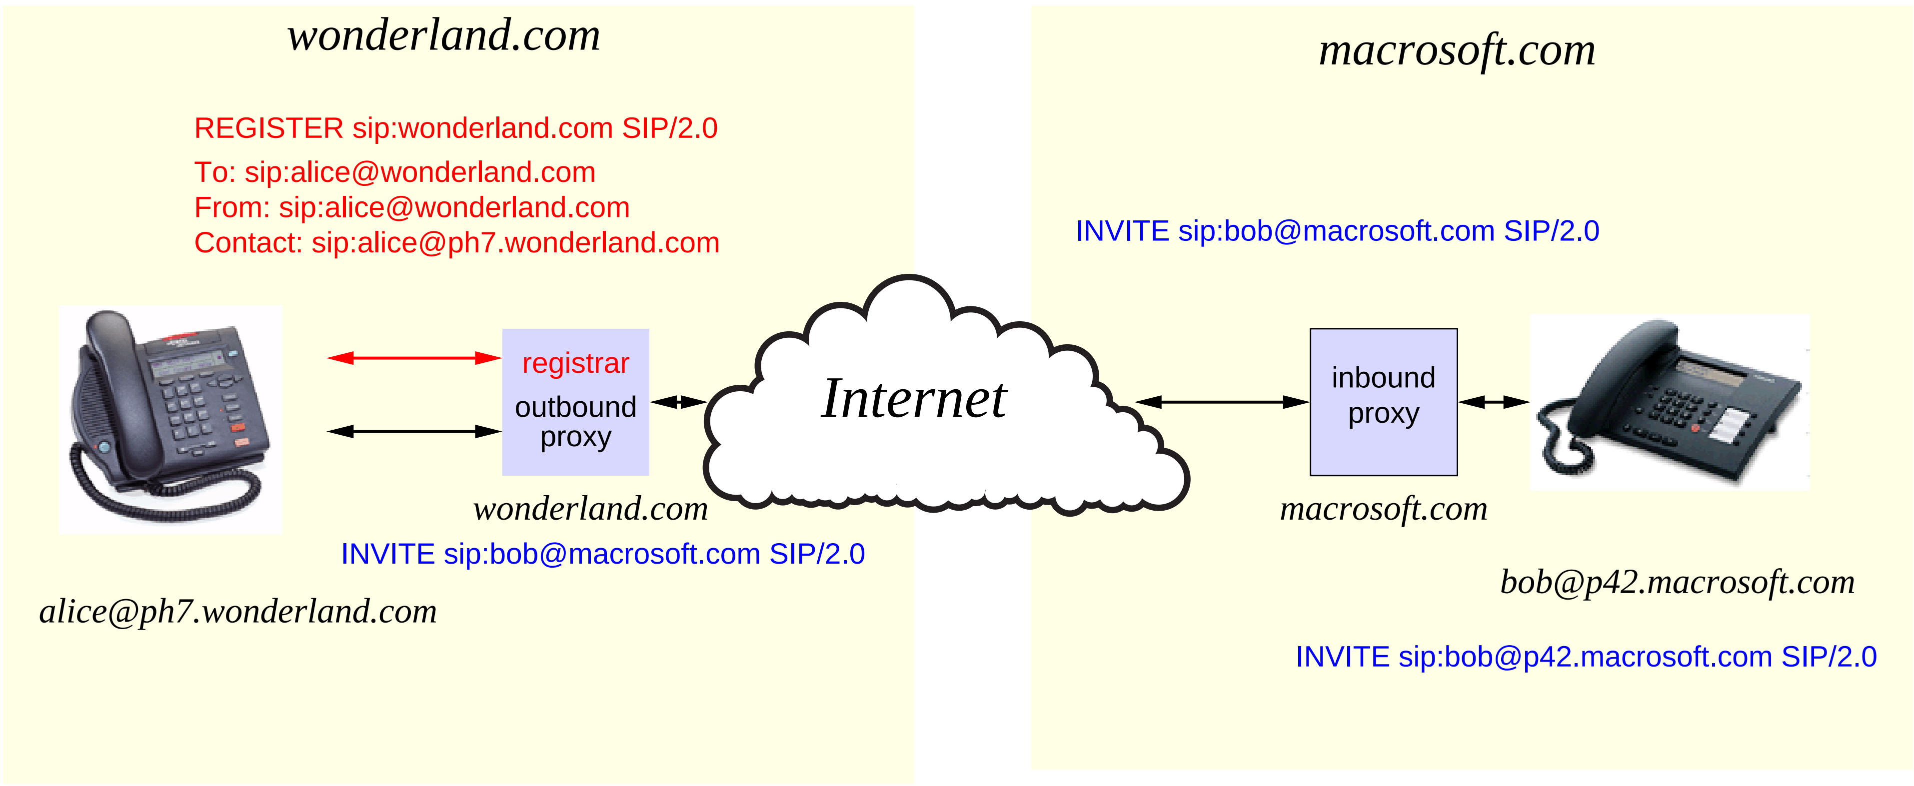

SIP – more detail

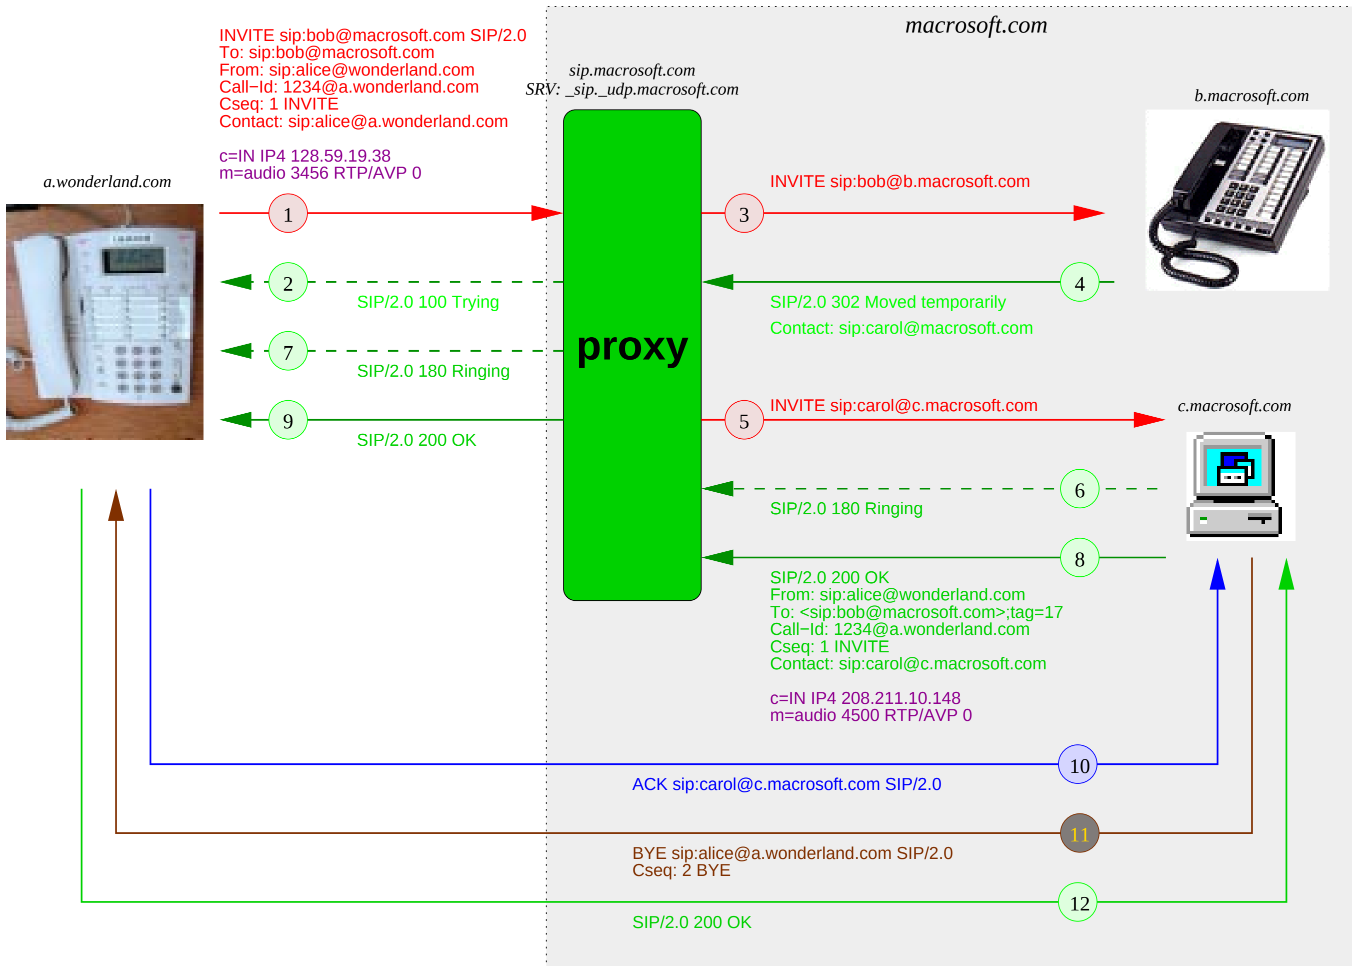

What is SIP good at?

- session setup = “out of band”
- resource location via location-independent identifier (“user@domain”, tel)
- particularly if location varies rapidly or filtering is needed (i.e., is inappropriate for DNS and LDAP)
- real-time: faster than email
- reach multiple end point simultaneously or in sequence = *forking*
- possibly hide end-point location
- delayed final answer (“ringing”) $\longleftrightarrow$ RTSP

What is SIP not meant for?

- bulk transport: media streams, files, pictures, ...
- asynchronous messaging ("email")
- resource reservation
- high-efficiency general-purpose RPC

SIP and Corba

	SIP	Corba
data	optional fields	versioning hard
	two-level hierarchy	general, C-like
hiding	dynamic	directory-based
multiple	forking proxy	no
transport	UDP, TCP, . . .	TCP
strength	inter-domain	inter-domain
generality	session set-up	RPC, events, . . .

SIP servers can benefit from Corba *locally* for user location and service creation

SIP and XML

- XML will play increasing role in SIP-enabled systems:
  - call processing language (CPL)
  - presence information for SIP as presence protocol
  - device configuration, buddy lists
  - possibly, future version of Session Description Protocol (SDP)
  - back-end for proxy services (e.g., Parlay over SOAP)
- but not appropriate everywhere:
  - can be verbose
  - hard to parse without generic (bulky) parser

Signaling and presence

- signaling: probe for presence
- presence: indicate presence
- presence $\subset$ event notification
- events: basement flooded

SIP for presence

- SUBSCRIBE to events, NOTIFY subscribers
- presence user agents (PUAs) provide information about events, e.g., through SIP REGISTER
- need to deal with transient user agents $\longrightarrow$ presence agent
- same name-based ("application layer") routing
- instant messaging: MESSAGE — allows mobility, independent of connection

SIP network appliances – the SIP toaster

- many in-house control and location technologies: X10, Jini, VESA home networks, . . .
- use MESSAGE + XML device control description for
- may use SLP locally: `d=lamp,r=kitchen`
- devices can move across networks, SIP address stays the same
- new address book entry: `washer@hgs.home.net`
- can subscribe to events: “laundry done”

Columbia e*phone

- DSP for voice coding *and* signaling ➠ limited memory (e*phone: 512 kB SRAM)
- only need minimal IP stack (IP/UDP/RTP, DHCP, SIP, tftp, DNS), not TCP
- also, MP3 radio
- sensor interfaces to the world: chair, IR, temperature, . . .

Columbia e*phone

Mobility in an IP environment

Terminal mobility: terminal moves between subnets

Personal mobility: different terminals, same address

Service mobility: keep same services while mobile

Terminal mobility

- domain of IEEE 802.11, 3GPP, mobile IP, ...
- main problems:
  - handover performance
  - handover failure due to lack of resources in new network
  - authentication of redirection

Personal mobility

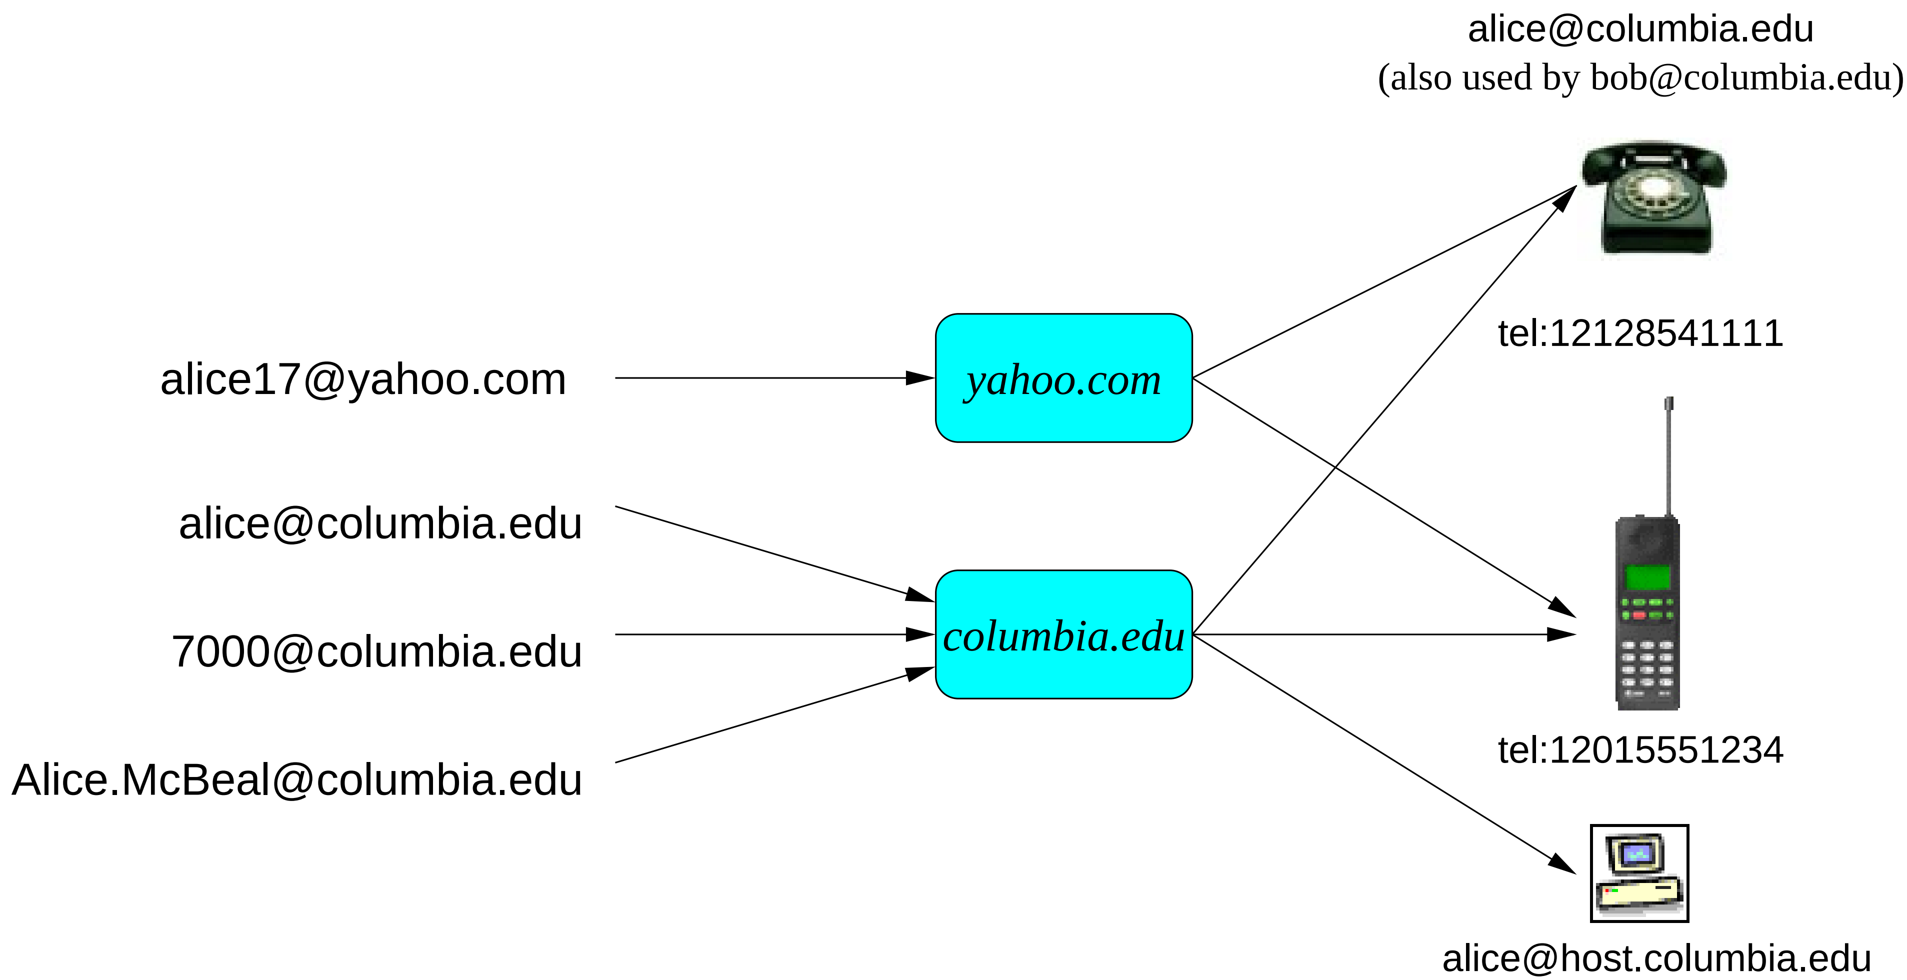

Personal mobility

- switch between PDA, cell phone, PC, Ethernet phone, Internet appliance, . . .
- several "generic" addresses, one person/function, many terminals
- e.g., `tel:2129397042`, `hgs@cs.columbia.edu`, `schulzrinne@yahoo.com` or `support@acme.com`
- SIP is designed for that – proxying and redirection does translation
- but: need mapping mechanisms to recognize registrations as belonging to the same person
- some possible solutions:
  - dip into LDAP personnel database or `/etc/passwd` to match phone number and variations of name (*J.Doe*, *John.Doe*, *Doe*)
  - need dialing plan to recognize `7042@cs.columbia.edu` and `tel:2129397042` as same

Service mobility

Examples:

- speed dial & address book
- media preferences
- special feature buttons (voice mail, do-not-disturb)
- incoming call handling instructions
- buddy lists

$\longrightarrow$ independent of terminal (including pay phone!), across providers

Service mobility

- REGISTER can retrieve configuration information (e.g., speed dial settings, distinctive ringing or voice mail settings)
- but needs to be device-independent
- most such services (e.g., voicemail forwarding, call filtering) should remain on server(s)

Separate issue: how does the payphone (or colleague's phone) recognize you?

- PDA (IR)
- i-button
- fingerprint
- speech recognition, . . .

One device, but changing set of owners!

Service mobility – call handling

- need uniform basic service description model $\longrightarrow$ Call Processing Language (CPL)
- CPL = XML-based flow graph for inbound & outbound calls
- CPL for local call handling
- update CPL from terminal: add telemarketer to block list
- harder: synchronize CPL changes across multiple providers
- one possibility: REGISTER updates information, but device needs to know that it has multiple identities
- merging of call logs

Terminal mobility – details

- move to new network ➠ IP address changes (DHCP)
- mobile IP hides address changes
- but: little deployment
- encapsulation overhead
- dog-legged routing
- may not work with IP address filtering

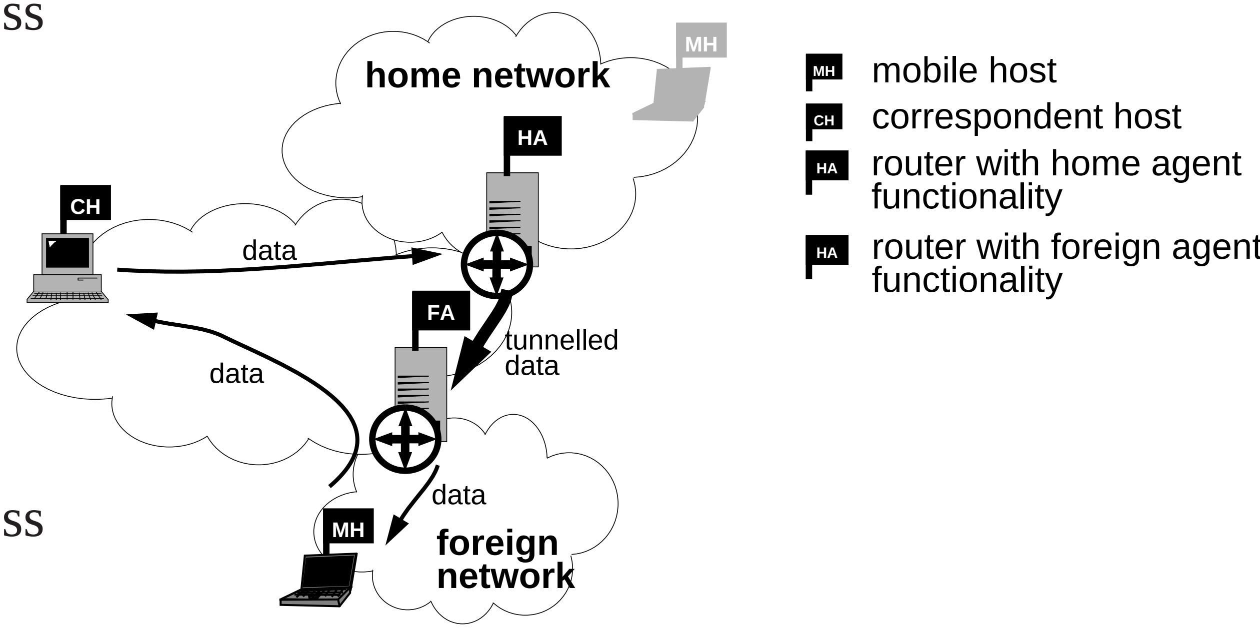

SIP terminal mobility overview

- pre-call mobility ➠ SIP proxy, redirect
- mid-call mobility ➠ SIP re-INVITE, RTP
- recovery from disconnection

SIP terminal mobility: pre-call

- MH acquires IP address via DHCP
- optional: MH finds SIP server via multicast REGISTER
- MH updates home SIP server
- optimization: hierarchical LR (later)

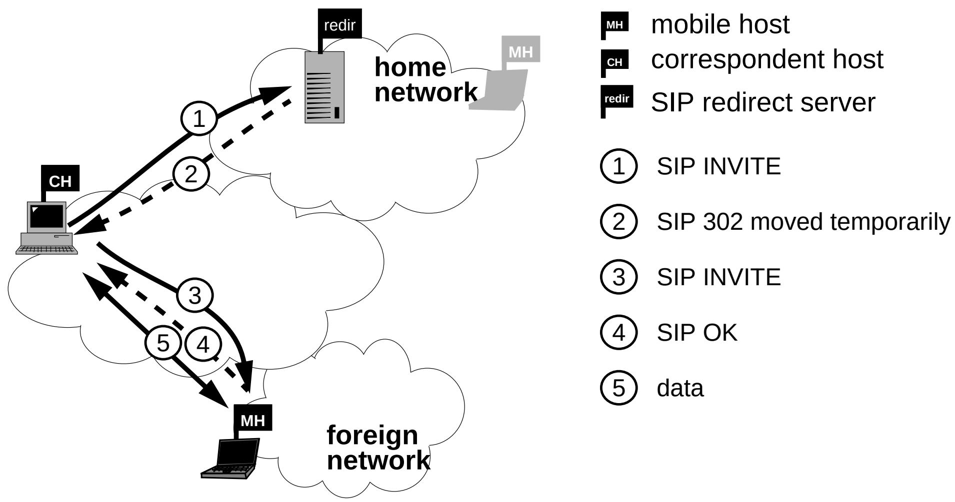

SIP terminal mobility: mid-call

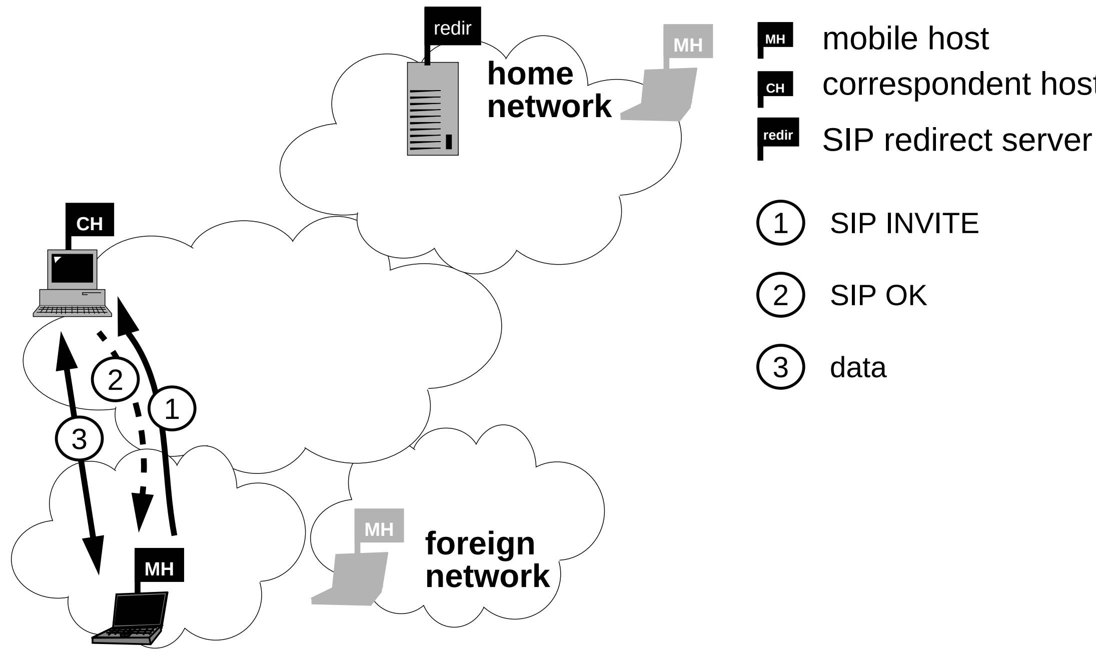

- MH→CH: new INVITE, with Contact and updated SDP
- re-registers with home registrar

SIP terminal mobility: multi-stage registration

Don’t want to bother home registrar with each move

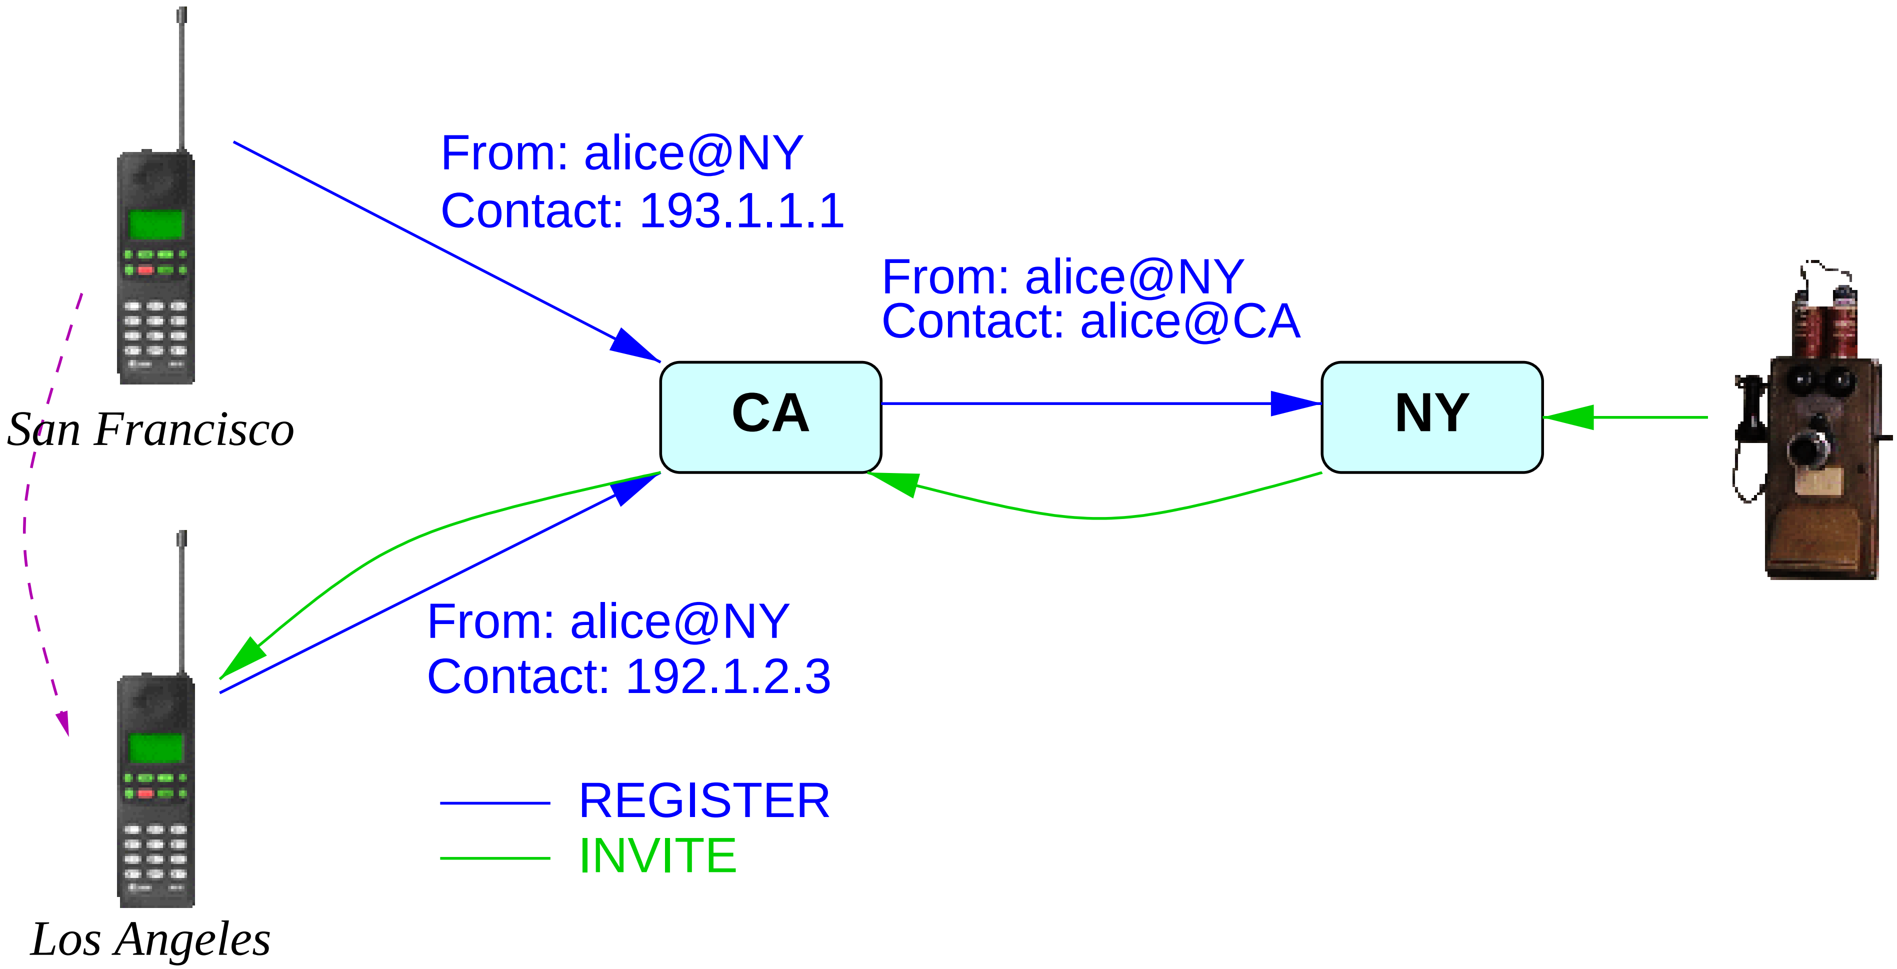

SIP and mobility: issues

- doesn't work for TCP applications – solutions:
  - punt: "don't walk while telnet'ing"
  - application-layer awareness: restart web, email, ftp transfer – need for deep fade anyway…
  - NAT-style boxes controlled by SIP (see Telcordia ITSUMO project)
- but: works nicely for "vertical handoff" between different technologies - e.g., transfer call from mobile handset to office videophone when arriving at work

Alternative wireless architectures

- 3G spectrum is expensive – UK: \$300 for every adult for single company (out of 5), \$1,200 assuming 25% sign up
- can build license-free infrastructure for that . . .
- idea: embed IEEE 802.11 11 Mb/s transmitter in every DSL or cable modem
- range: 80’ (indoor) to 500’ (open)
- complications:
  - need dual-mode for highway use $\longrightarrow$ vertical hand-off
  - each house may have a different network provider
  - fast IP-level hand-off, can’t rely on micro-mobility

Partial connectivity: *7 Degrees of Separation*

- BlueTooth and 802.11 also usable for peer-to-peer communication
- architecture: occasional data “filling stations”
- users exchange data – newspaper articles, MP3 files, . . . – when within wireless range
- each mobile is a small web server and search engine
- match according to topic (“Yankees”) or URL (“www.nytimes.com”)
- Manhattan: 4,434 people/km^2 $\longrightarrow$ 25 people in square km = 0.6% interested
- about 8% of the NYC reads the *NY Times*
- assumption: start with single *data holder*

Average full propagation delay for random way model

walk in straight line, then stop and change directions

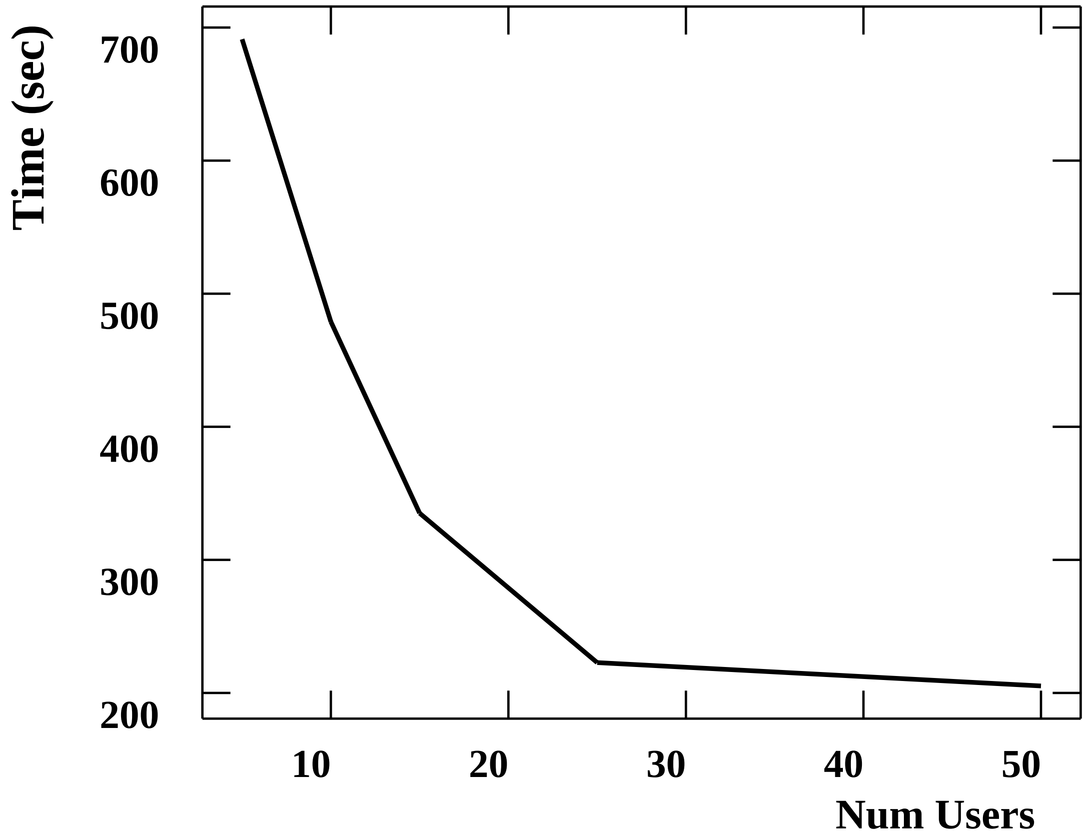

querying interval: 60 s

Average delay for random way model

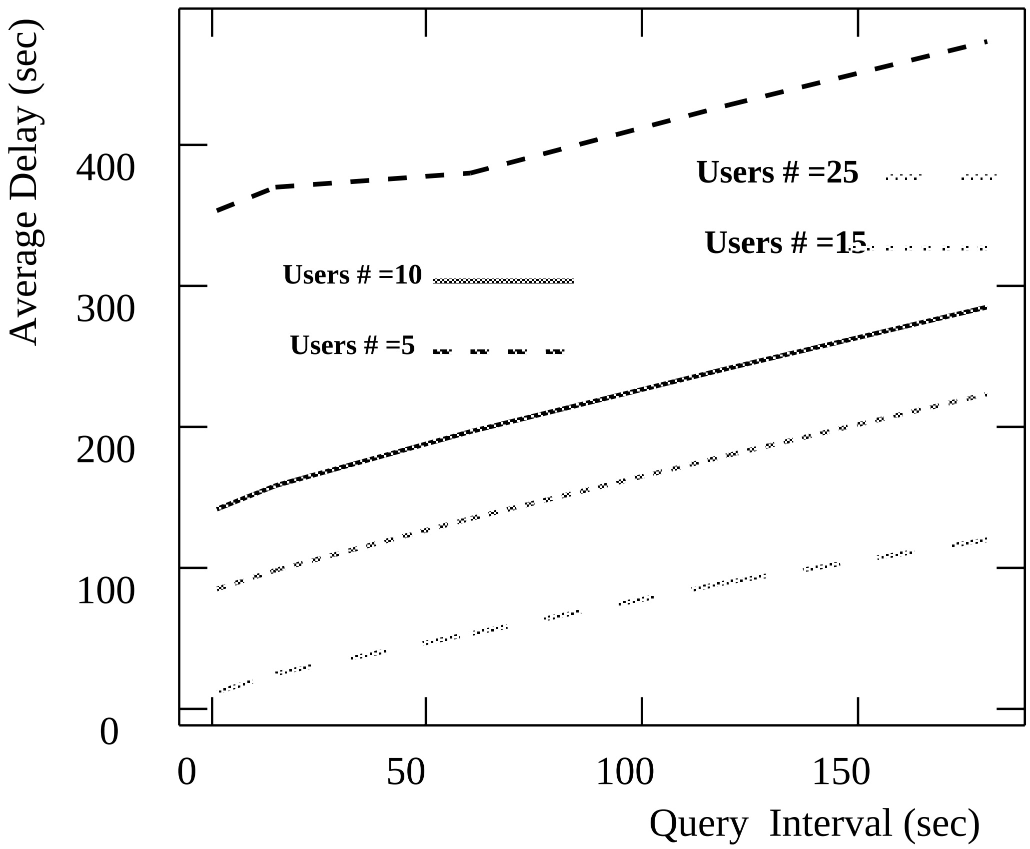

Subway model (IRT)

- Poisson process with a $1/\lambda$ of $60 \ldots 180\,\mathrm{s}$
- train with six cars arrives every 5 minutes
- stops for $45\,\mathrm{s}$
- ten stops
- time between stations uniformaly between $168 \ldots 210\,\mathrm{s}$
- ride between 2 and 6 stations
- data exchange on subway platform and in train car

Percentage of data holders

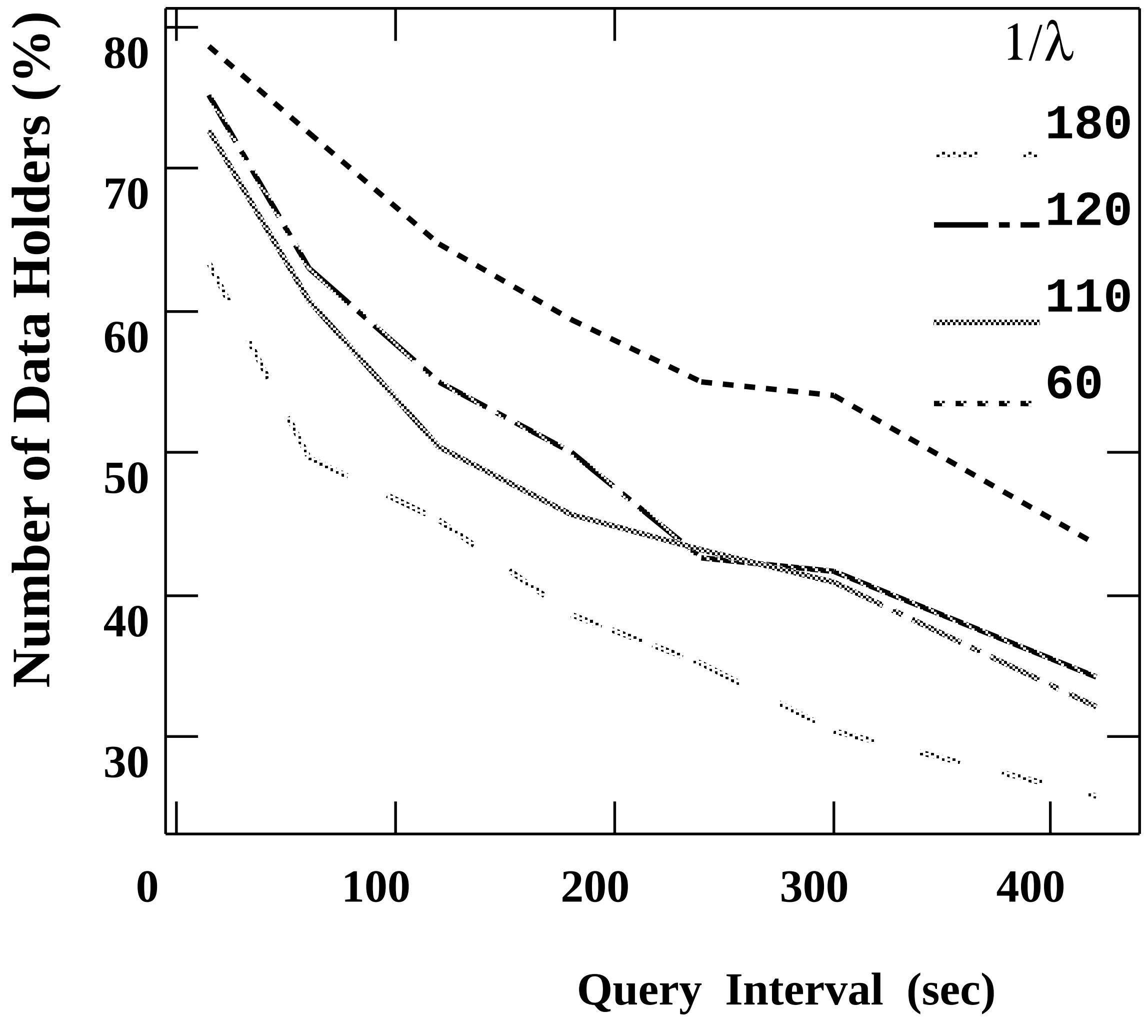

Delay for subway model

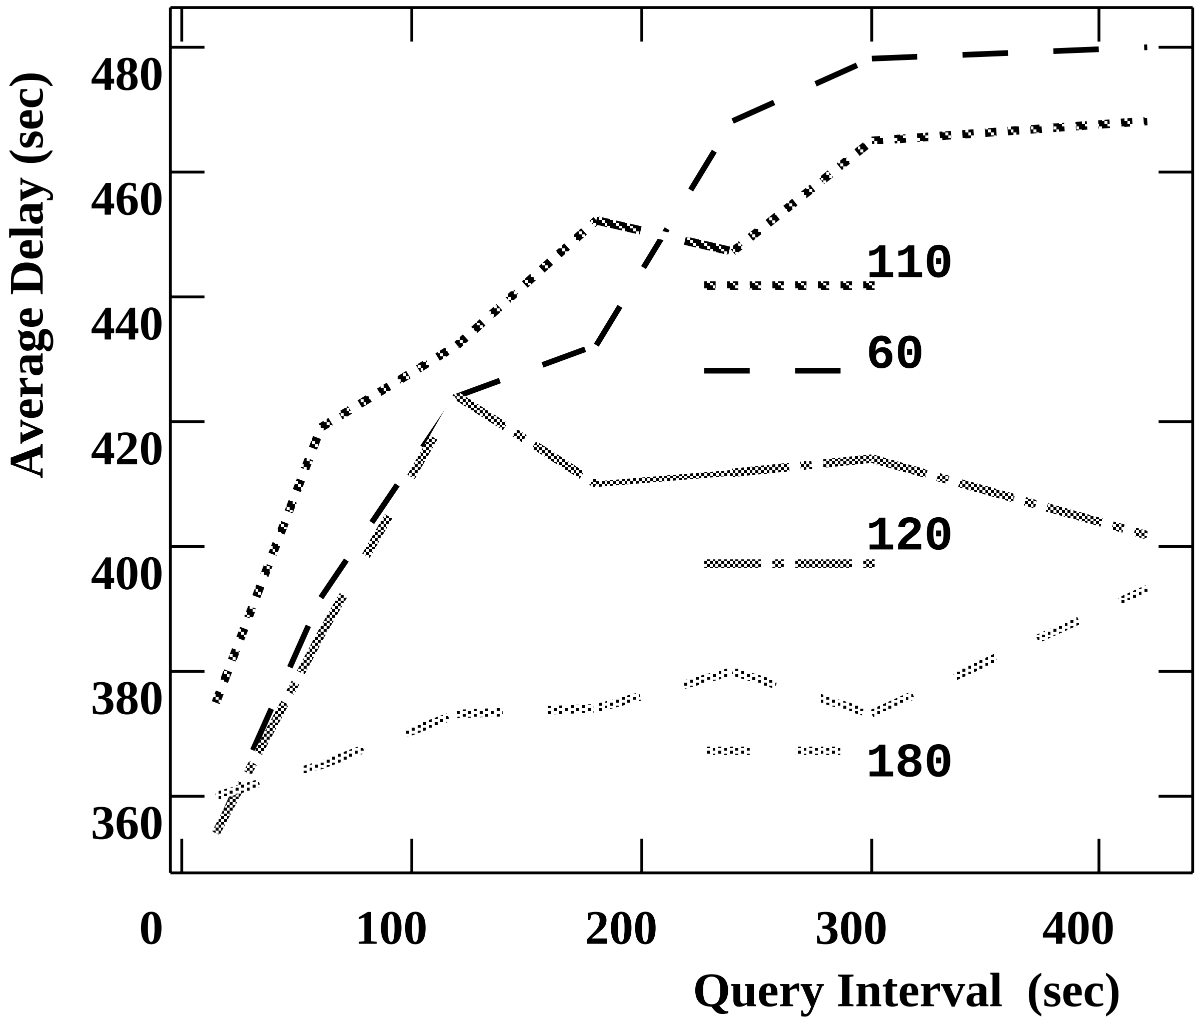

Epidemic model

- $t = 0$: one data holder ("infected"), $N - 1$ queriers ("susceptibles")
- in interval h, infected will transmit with $h\alpha + o(h)$
- pure birth process with $\lambda_k = (N - k)N\alpha$ for k data holders
- $E[T] = \frac{1}{\alpha} \sum_{i=1}^{N-1} \frac{1}{i(N-1)}$

Epidemic model

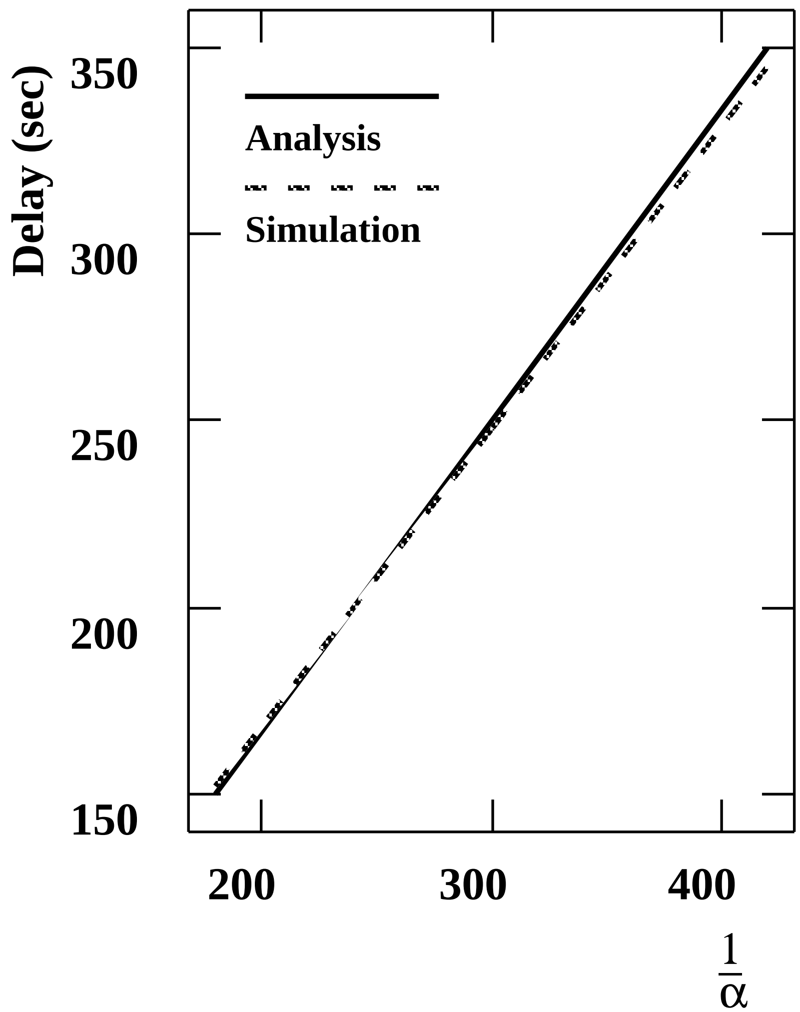

Conclusion

- streaming and Internet telephony should share end systems
- Internet telephony is an Internet application, not telephony on packets
- signaling and presence are duals of each other
- mobility is more than just wireless handsets
- terminal, personal and service mobility
- SIP enables all three, but likely to be hybrid solutions
- high-speed wireless more likely in islands (e.g., Wayport)

For more information. . .

SIP: http://www.cs.columbia.edu/sip

RTP: http://www.cs.columbia.edu/˜hgs/rtp

RTP: http://www.cs.columbia.edu/˜hgs/rtsp

Papers: http://www.cs.columbia.edu/IRT