Internet Telephony: Status and Directions

Henning Schulzrinne
Columbia University, New York
schulzrinne@cs.columbia.edu

Overview

- new Internet services: “telephone”, “radio”, “television”
- why Internet telephony?
- why not already?
- Internet telephony modalities
- components needed:
  - data transport
  - resource reservation
  - signaling
  - service location

New Internet services

- tougher: replacing dedicated electronic media
- typewriter model of development
- yet another convergence?

The phone works — why bother with VoIP?

user perspective

- variable compression: tin can to broadcast quality
- security through encryption
- caller, talker identification
- better user interface
- internat. calls: TAT transatlantic cable = $0.03/hr
- no local access fees (3.4c)
- easy: video, whiteboard, . . .

carrier perspective

- silence suppression ⇛ traffic ↓
- shared facilities ⇛ management, redundancy
- advanced services (simpler than AIN and CTI)
- operational advantages
- cheaper switching
- fax as data

The new phone companies

- separation bit carriage $\leftrightarrow$ services
- anybody with Internet connection can provide services (ACD, 800, 900, directory, . . .)
- distinction “in” vs. “out” of network not useful
- incremental start-up investment not large
- new players:
  - cable companies ➠ no new infrastructure, but mostly one-way
  - electric utilities ➠ need line management anyway
  - Qwest, IXC (resell to ISPs), . . .

Internet telephony as PBX replacement

global Internet not quite ready ➠ try as PBX

- have mission-critical LAN, PCs anyway
- usually ample (if switched) bandwidth, low latency
- packet switching is cheaper
- network PCs $\overset{\$}{=}$ ISDN phones
- no need for billing

Internet telephony services

- voice mail $\longrightarrow$ email
- calendar integration
- user-programmable call processing logic
- call first available sales person (ACD)
- call whole department
- web IVR
- return web page with favorite “on hold” music

Internet telephony services

- camp-on without holding a line
- short message service (“instant messaging”)
- schedule call into the future
- call with expiration date
- add/remove parties to/from call ➠ mesh
- “buddy lists”

Switching costs

switching method	ports	Gb/s	cents/kb/s	$/interface
10BaseT hub	16	0.16	0.1	9.4
100BaseT hub	16	1.6	0.05	46
10BaseT switch	24	0.24	1.2	121
100BaseTX switch	8	0.80	0.15	156
router		2.1	16.0	
local ATM switch	16	2.48	1.0	1581
PBX	256	0.02	218.	140
5ESS local (no AIN)	5,000	0.32	469.	300
5ESS local (AIN)	20,000	1.28	273.	175
4ESS toll	100,000	6.40	7.8	

Telephone costs

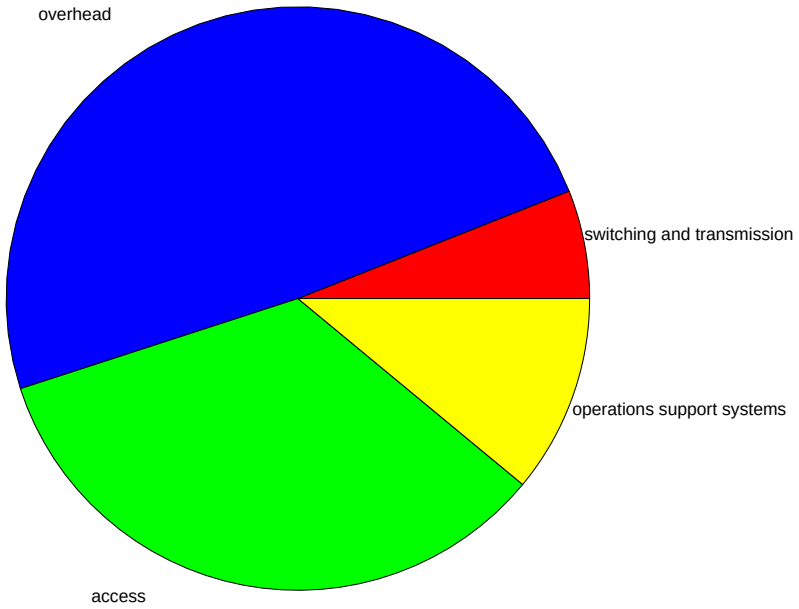

infrastructure	10-23%
switching and transmission	6%
overhead	49%
access	34%
operations support systems	11%

Transport costs

network	$/min	$/MB
wholesale telephone	0.01–0.02	
U.S. domestic interstate consumer rates	0.05–0.15	
U.S. domestic intrastate consumer rates	0.05–0.25	
modem		0.25 – 0.50
private line		0.50 – 1.00
frame relay		0.30
MCI frame SVC		0.05
Internet		0.04 – 0.15
Internet modem		0.33
Internet backbone		0.01

1 minute voice = 480 kB with silence suppression, 1 MB without

Phone usage

“Free” phone calls does not mean unbounded increase:

year	lines (millions)	local calls min/day/line	local calls min/day/person
1980	102.2	39	17.5
1988	127.1	39	20.2
1996	166.3	40	25.1

Why aren't we using it now?

Internet capacity $\ll$ phone traffic:

world phone traffic	600	Gb/s	U.S. total	368	Gb/s
international traffic	13	Gb/s	U.S. interstate	55	Gb/s
			AT&T long distance	61	Gb/s
MCI Internet	1.8	Gb/s			
public Internet (late 1997)	75	Gb/s			

- unpredictable sound quality, reliability
- doesn't work well for dial-up users
- no cheap Internet devices
- 640 M phone lines, 122 M in U.S. ➠ gateways
- no billing infrastructure

Projections

- MCI: “80% data, 20% voice”
- “AT&T could lose $350 million in international calls by 2001”
- “By 2002, the Internet could account for 11% of U.S. and international long-distance voice traffic”
- “Up to 10% of the world’s fax market, which generates $45 billion in telecom revenue a year, will move to Internet in 2 or 3 years”
- May 1999: BT builds IP phone network in Spain
- but: cable modems only 250,000 to 275,000 users in US, 10% of Internet users by 2000

Data vs. Voice Traffic

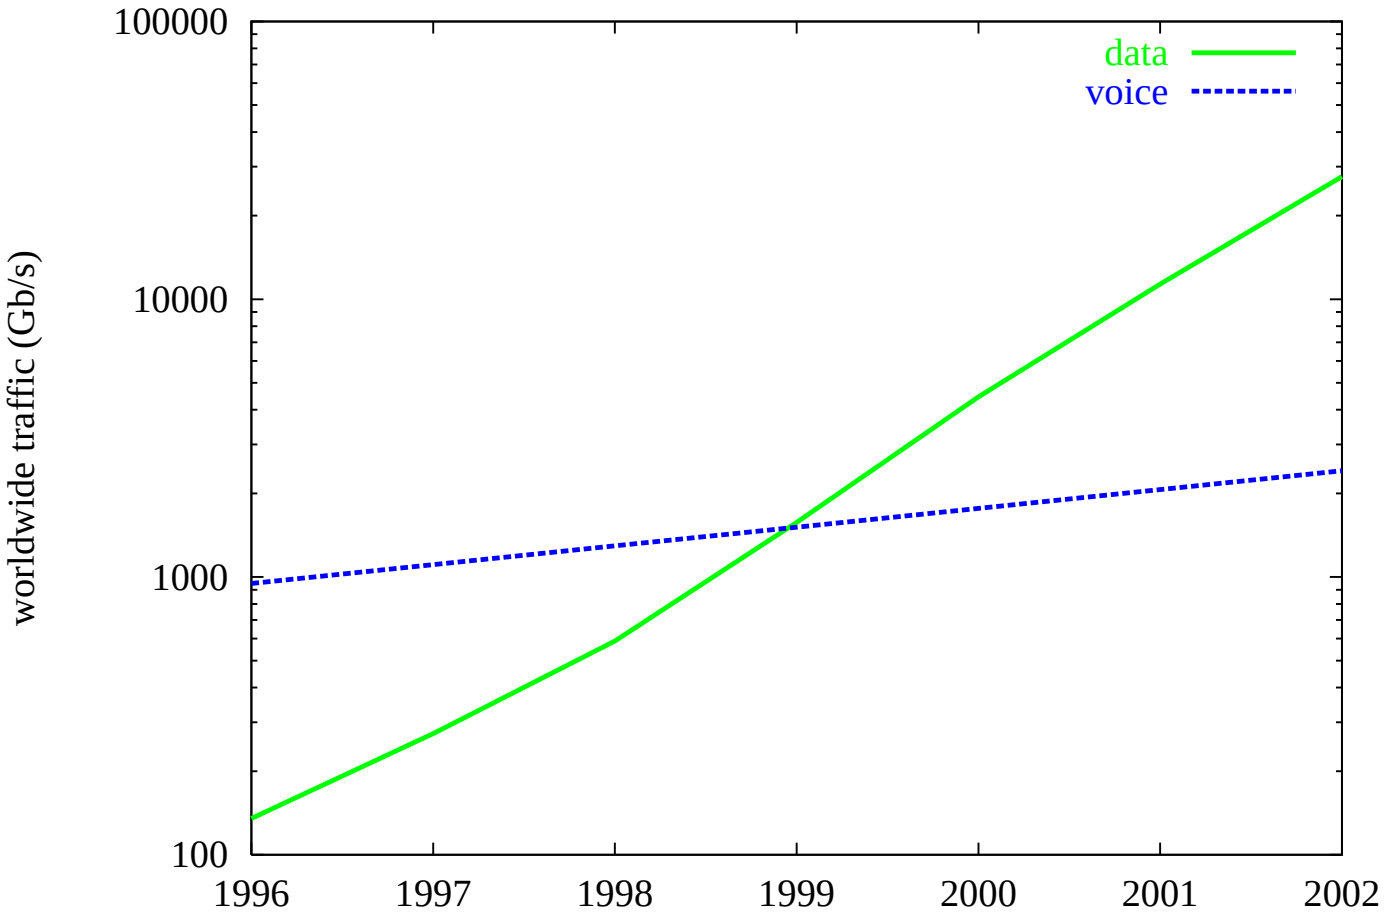

Components for Internet Multimedia

Internet multimedia protocol stack

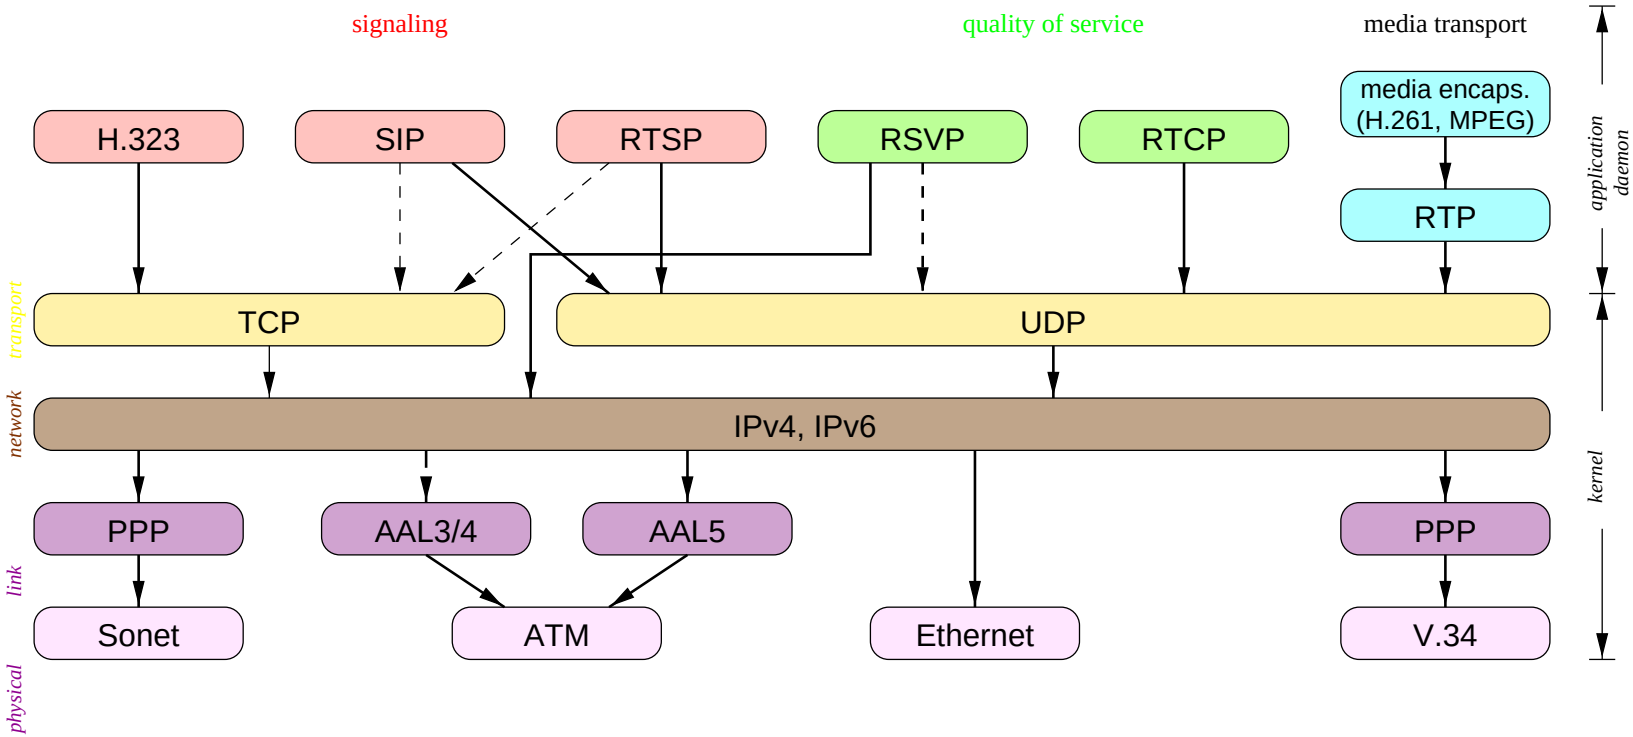

Components for Internet Multimedia

multicast: routing, address allocation

data transport: RTP

resource reservation: RSVP, YESSIR, diff-serv

“TV” – announcing multicast sessions: SAP

“phone” – session setup for conferences/telephony: SIP

“VCR” – control of streaming media: RTSP

local applications: conference bus

policy issues: billing, firewall access, clearing houses

Applications for Multicast

- audio-video distribution (1-to-many) and symmetric (all-to-all)
- distributed simulation (war gaming, multi-player Doom, . . .)
- resource discovery (where's the next time server?)
- file distribution (stock market quotes, new software, . . .)
- network news (Usenet)

Host group model

Deering, 1991:

- senders need not be members;
- groups may have any number of members;
- there are no topological restrictions on group membership;
- membership is dynamic and autonomous;
- host groups may be transient or permanent.

IP Multicast: Problems

- multicast routing ➠ state $\propto S, G$
- proposals:
  - DVMRP, PIM-DM for dense groups
  - PIM-SM or CBT for sparse groups (“core”)
- overlay networks (Mbone) hard to maintain
- billing and charging (satellite TV problem)
- multimedia applications mostly on-demand

Multicast address allocation

- about 268 mio. “class D” addresses
- can’t have FCC assign channels
- hierarchical borrowing, using DHCP locally
- IETF malloc WG

Data transport – RTP

Real-Time Transport Protocol (RTP) = data + control

data: timing, loss detection, content labeling, talkspurts, encryption

control: (RTCP) ➠ periodic with $T \sim$ population

- QOS feedback
- membership estimation
- loop detection

RTP functions

- segmentation/reassembly done by UDP (or similar)
- resequencing (if needed)
- loss detection for quality estimation, recovery
- intra-media synchronization: remove delay jitter through playout buffer
- intra-media synchronization: drifting sampling clocks
- inter-media synchronization (lip sync between audio and video)
- quality-of-service feedback and rate adaptation
- source identification

Resource Reservation

- *can't* compensate for lack of bandwidth or reliability
- *can* provide incumbency protection
- receiver makes requests ➠ RSVP
- sender makes requests ➠ YESSIR
- issues: scaling (state), security, complexity

Internet telephony modes

- tail-end hop off ➠ callee has phone
- front-end hop on ➠ caller uses phone
- Internet in the middle: per-call, multiplexed

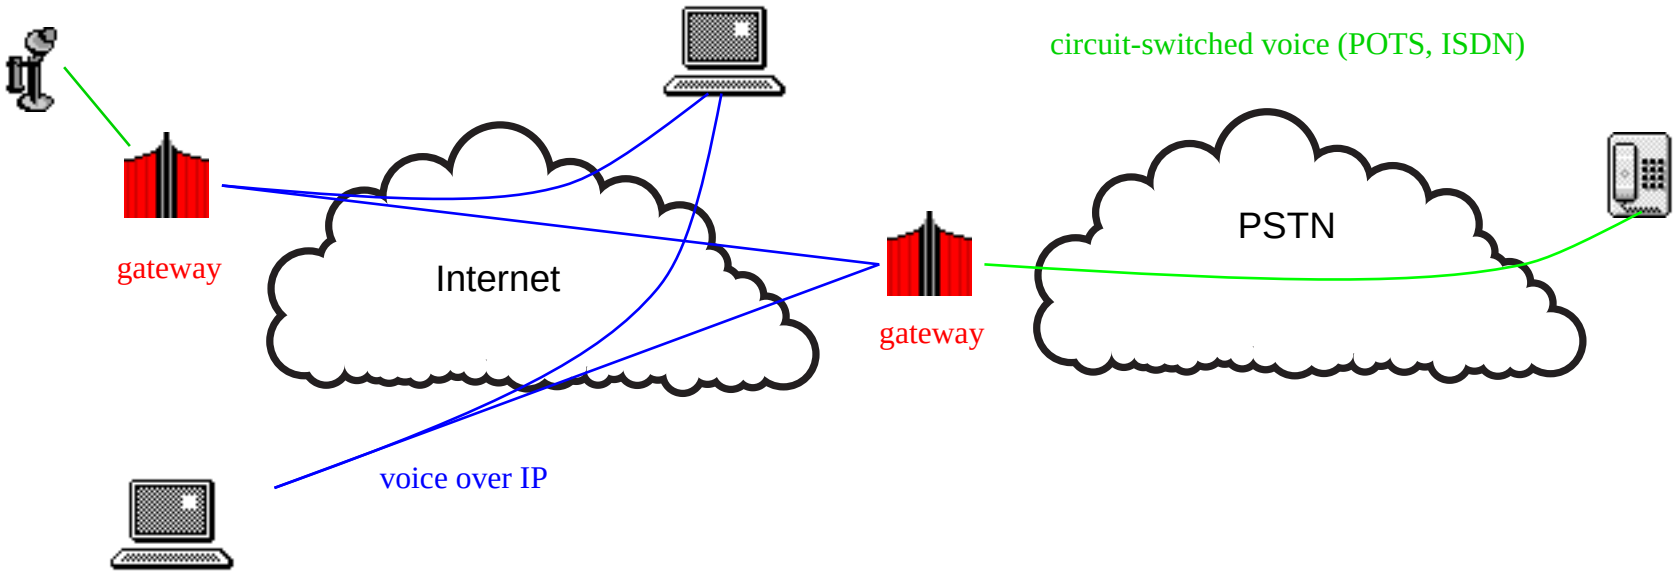

Internet “signaling”

all non-data (“out-of-band”) functions:

routing: unicast; DVMRP, PIM, CBT for multicast $\surd$

quality of service: RSVP, RTCP, diff-serv $\surd$

user Contact: map name to location (IP address)

call set-up/teardown: SIP, H.323

policy, billing: “vertical” protocols

Architecture

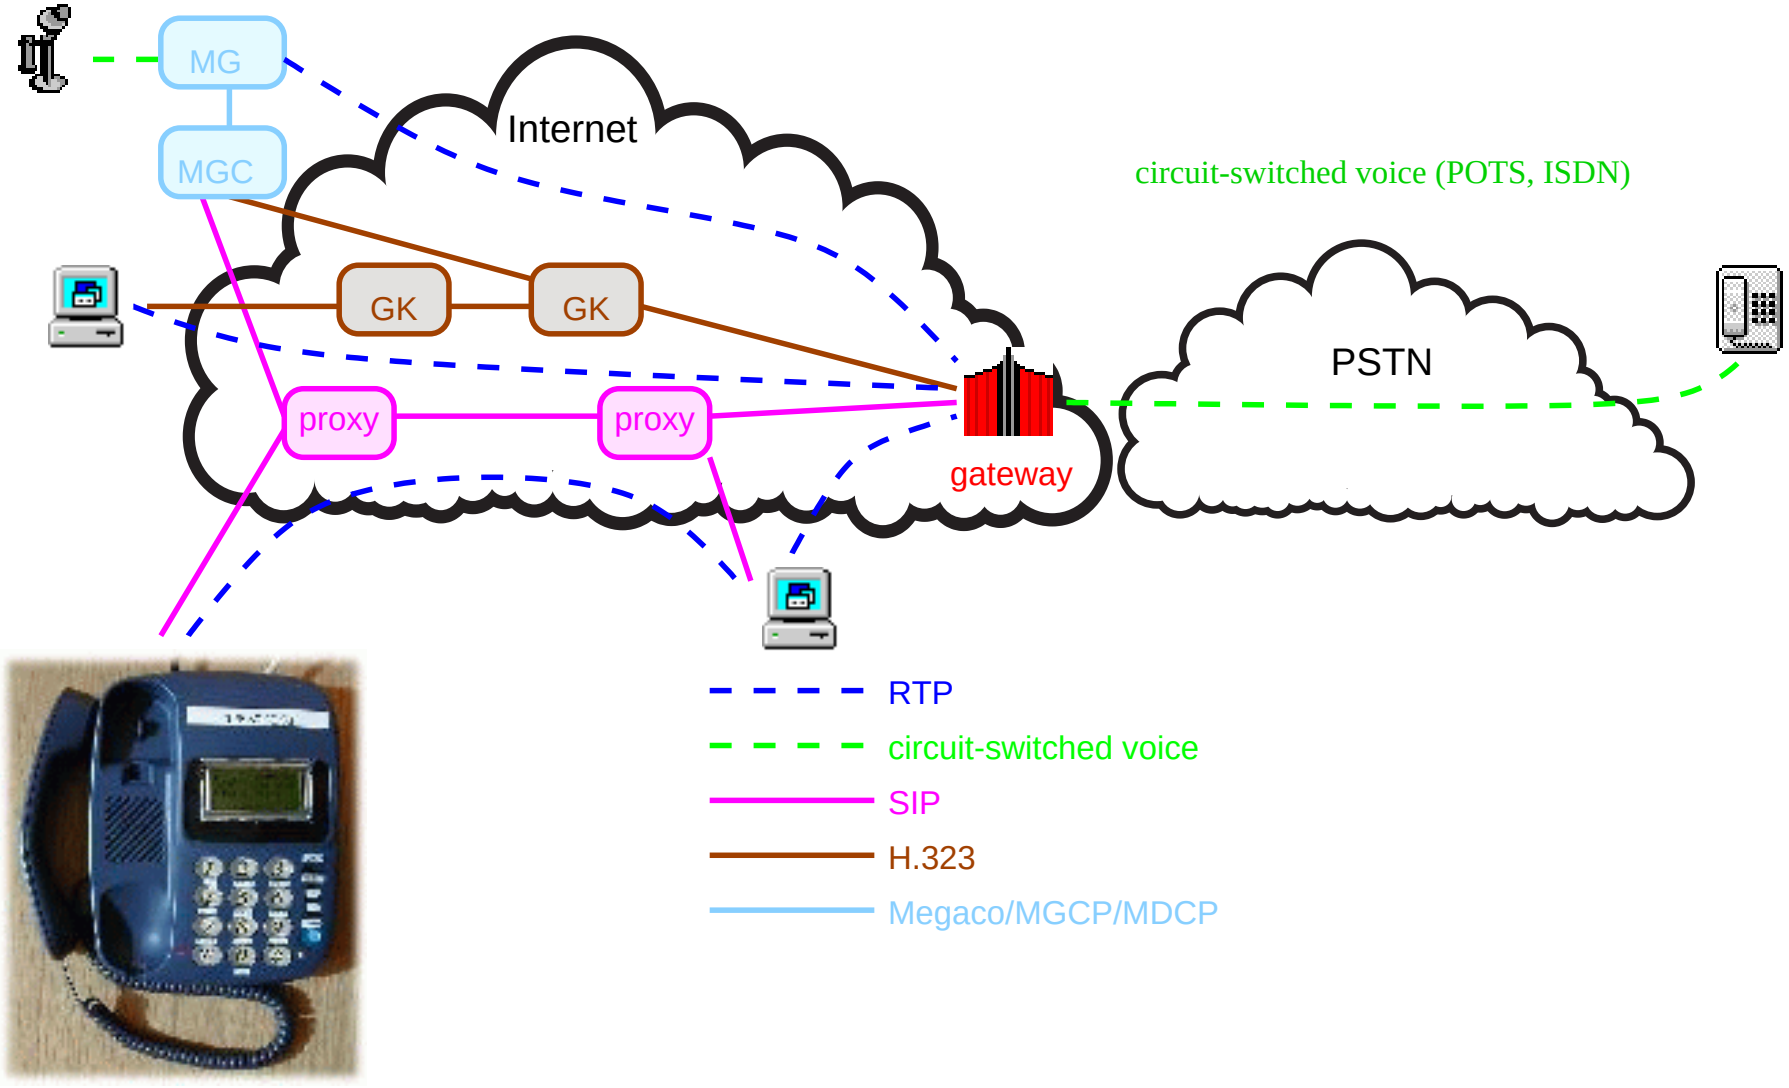

Differences: Internet Telephony ↔ POTS

- separate control, transport (UDP) ➠ no triangle routing
- separate connectivity from resource availability
- separate services from bit transport
- datagram service ➠ less bootstrapping
- in-band signaling ➠ higher speed
- features "network" → end system: distinctive ringing, caller id, speed dialing, number translation, . . . ➠ scaling
- features: intra-PBX = inter-LATA and general
- protocols: user-network = network-network signaling

Internet Telephony

- multimedia basically free (unlike ISDN)
- minimal extensions: signaling, not “stove pipe”
- leverage existing work: email, HTTP security, URIs, HTML, cgi, . . .

Light-weight signaling: Session Initiation Protocol (SIP)

IETF MMUSIC working group (RFC 2543)

- light-weight generic signaling protocol
- typical post-dial delay: 1.5 round-trip time (with UDP)
- network-protocol independent: UDP or TCP (or AAL5 or X.25)

SIP functionality

- call user
- re-negotiate call parameters
- manual and automatic forwarding
- call center: reach first (load distribution) or reach all (department conference)
- *personal mobility* (complements data link/IP mobility) ➠ change of terminal (PC, digital cordless, palmtop), location
- terminate and transfer calls

Service creation: Call Processing Language

- incoming and outgoing
- “if somebody is trying to call for the 3rd time, allow mobile”
- “try office and lab in parallel, if that fails, try home”
- “allow call to mobile if I’ve talked to person before”
- users and administrators
- not quite like cgi: multiple responses? timers?
- Tcl, Java?

Real-Time Streaming Protocol (RTSP)

remote-control streaming media ➠

- “rough” synchronization (fine-grained ➠ RTP sender reports)
- virtual presentations = synchronized playback from several servers ➠ command timing
- load balancing using redirection at connect, during stream
- supports any session description
- device control ➠ camera pan, zoom, tilt
- caching: similar to HTTP, except “cut-through”

Open Operational Issues

- billing
- finding the nearest gateway to the Internet (➠ GLP)
- mapping E.164 (phone) numbers to IP addresses
- controlling phones through the Internet (PINT)
- 911 services
- CALEA
- anonymity and certified identity

Billing

- simplification: email/web delivery, credit card payment
- what to bill for?

  transport services: volume, time, reserved resources; "free upgrades"

  signaling services: filtering, forwarding, scripting, mobility, . . .

  storage services: voice mail

  gateway services: PSTN gateways

Emergency (911) services

- U.S.: dial “911” anywhere $\longrightarrow$ nearest Public Safety Answering Points
- look up street address from telephone company database
- but. . .
  - IP address dynamically assigned
  - may not be correlated to geography
  - dial-in from hotel, remote sites?
  - prevent services: hanging up, transfer, hold, . . .

Emergency services

- advantages:
  - multimedia (video, medical data, . . .)
  - medical database access, with authentication token
  - remote activation of medical devices
- solutions:
  - enclose (signed) location information with call
  - IP address → provider → lookup (RADIUS) ➠ needs authenticated protocol
  - GPS

Conclusion

- transition of separate circuit-switched ➠ IP-based applications
- packets from the inside out or the outside in?
- IP over ATM, Sonet, WDM?
- IPv6 or NATs?
- “the end of distance” or tiered IP service?