Mobile IP (multimedia) without Mobile IP: SIP and RTP mobility

Henning Schulzrinne
Dept. of Computer Science
Columbia University
New York, New York
schulzrinne@cs.columbia.edu

Bell Laboratories

June 8, 1999

(Joint work Elin Wedlund)

Overview

- mobility: which layer?
- SIP overview
- SIP mobility
- RTP mobility
- Open issues

Mobility

- connected: move while connected
- plug-in: lap top
- synchronization: PalmPilot

Review: Mobile IP

- goal: make IP address changes invisible to upper layers
- “home address” of *mobile host* (MH)

home agent (HA): tunnels IP packets to current location;

corresponent host (CH): the host that talks to MH (mobile or fixed);

foreign agent (FA): router/MH that decapsulates packets

Mobile moves to new network

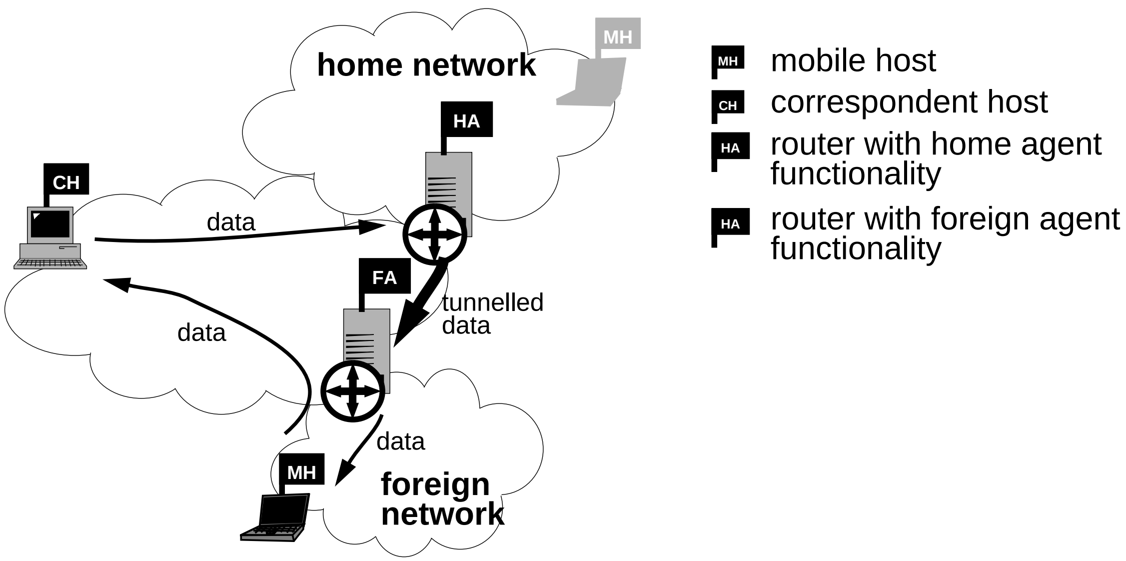

- extra HA–FA + encapsulation packet delay
- IP-in-IP encapsulation overhead of 20 bytes
- MH may send directly to CH, but may not get out

Mobile IP with route optimization

- CH tunnels directly to FA
- UDP-based *binding update* from HA
- old FA sends *binding warning* to HA ➠ HA sends *binding update*

Problems with route-optimized IP

- CH must change IP stack (tunnel!)
- CH must support binding update process
- only HA can send update ➠ delay
- MH has to rely on old FA to forward packets
- don’t send binding updates blindly

Why is there no mobile IP deployment?

- no incentive to install: few short-term visitors, few public access points beyond modems
- address-based packet filtering for return traffic
- operating system support lacking
- many synchronization-based ➠ partially disconnected operation

Application-Layer Mobility

Address assignment:

- DHCP
- IPv6 autoconfiguration (link-local addresses)

Web: "designed for mobility"?

- short transactions ➠ client can close
- restartable
- stateless at server (cookies)
- recoverable: "range" header

Application-Layer Mobility

Email:

- POP for email retrieval, local storage
- IMAP for server-storage
- web-based email access (`hotmail.com`)
- but: need "home" server for SMTP (spam prevention)

Application-layer mobility

News:

- download and read
- upload to local server

ftp:

- now: mostly large file transfers
- recent: restart capability ("smart download")

Application-layer mobility

file systems:

- NFS: UDP based
- even with TCP, stateless

telnet:

- type while walking/driving?
- decreasing importance

irc: needs continuous connectivity!

Mobile applications: radio and TV

- multicast: join at new location
- main problems: IGMP leave latency, IGMP traffic
- leave latency: depart without IGMP farewell
- smooth handoff if listener in next local network

What are truly mobile applications?

- clients only
- audio and video
  - 38% at home, 41% cars, 21% at work
  - three hours each weekday
- telephone: “number of mobile-phone subscriptions in Finland (2.9 mio.) has surpassed the number of traditional fixed-phone subscriptions.”

SIP Basics

Light-weight signaling: Session Initiation Protocol (SIP)

IETF MMUSIC working group (RFC 2543)

- light-weight generic signaling protocol
- part of IETF conference control architecture:
  - SAP for “Internet TV Guide” announcements
  - RTSP for media-on-demand
  - SDP for describing media
  - others: malloc, multicast, conference bus, …
- typical post-dial delay: 1.5 round-trip time (with UDP)

SIP Features, cont'd.

- 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 for H.323 experts

H.323	SIP + SDP
H.225.0 + RAS	SIP
H.245	SDP, SMIL, . . .
gatekeeper	proxy

SIP architecture: peer-to-peer

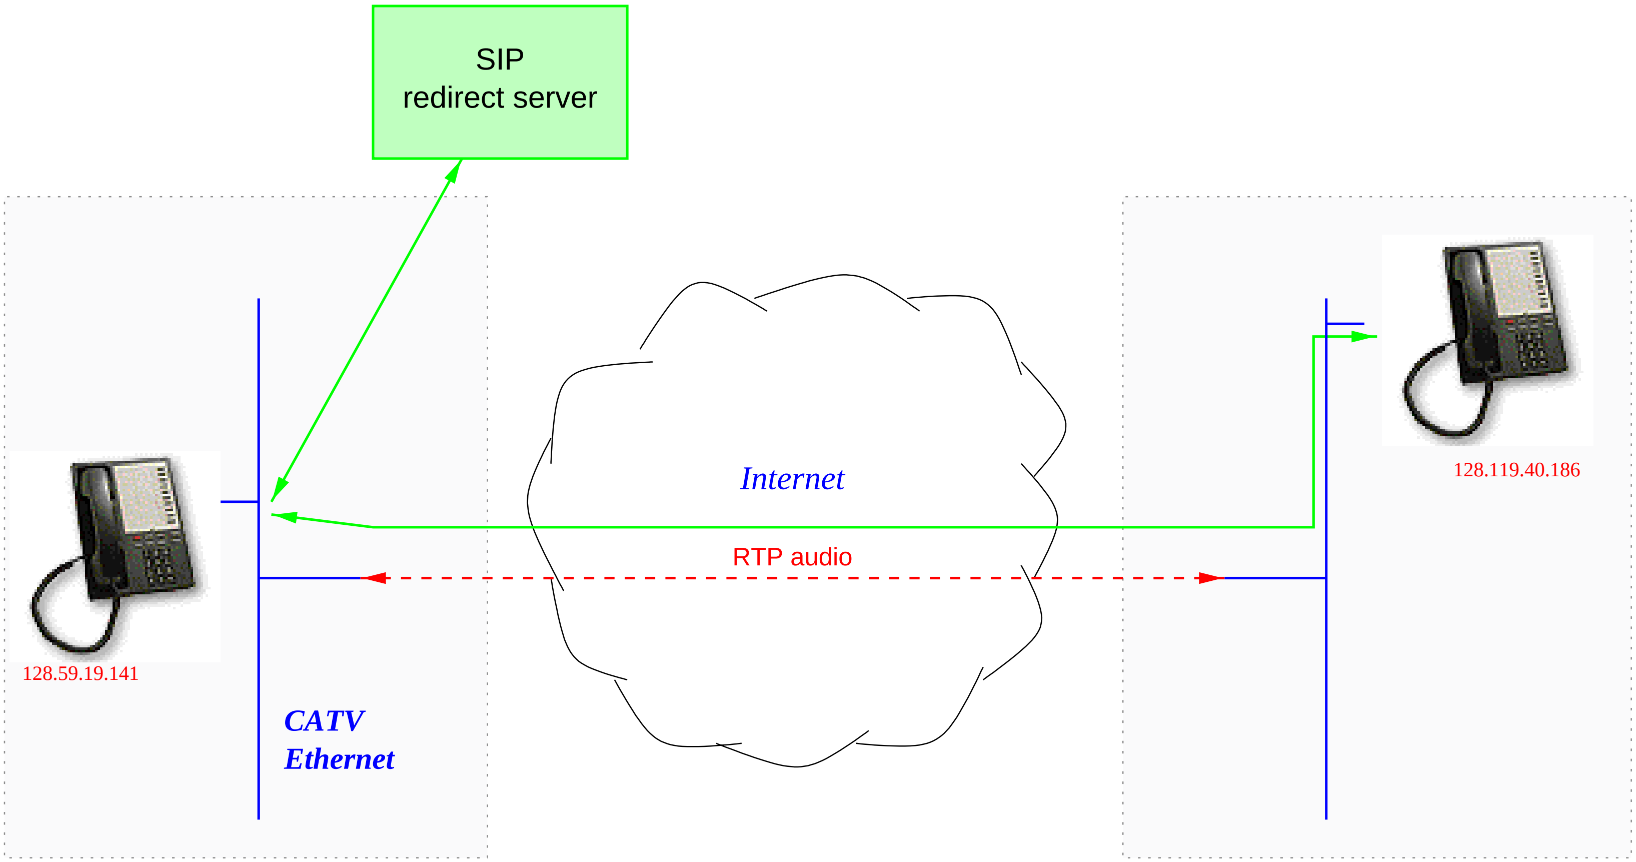

SIP architecture: carrier

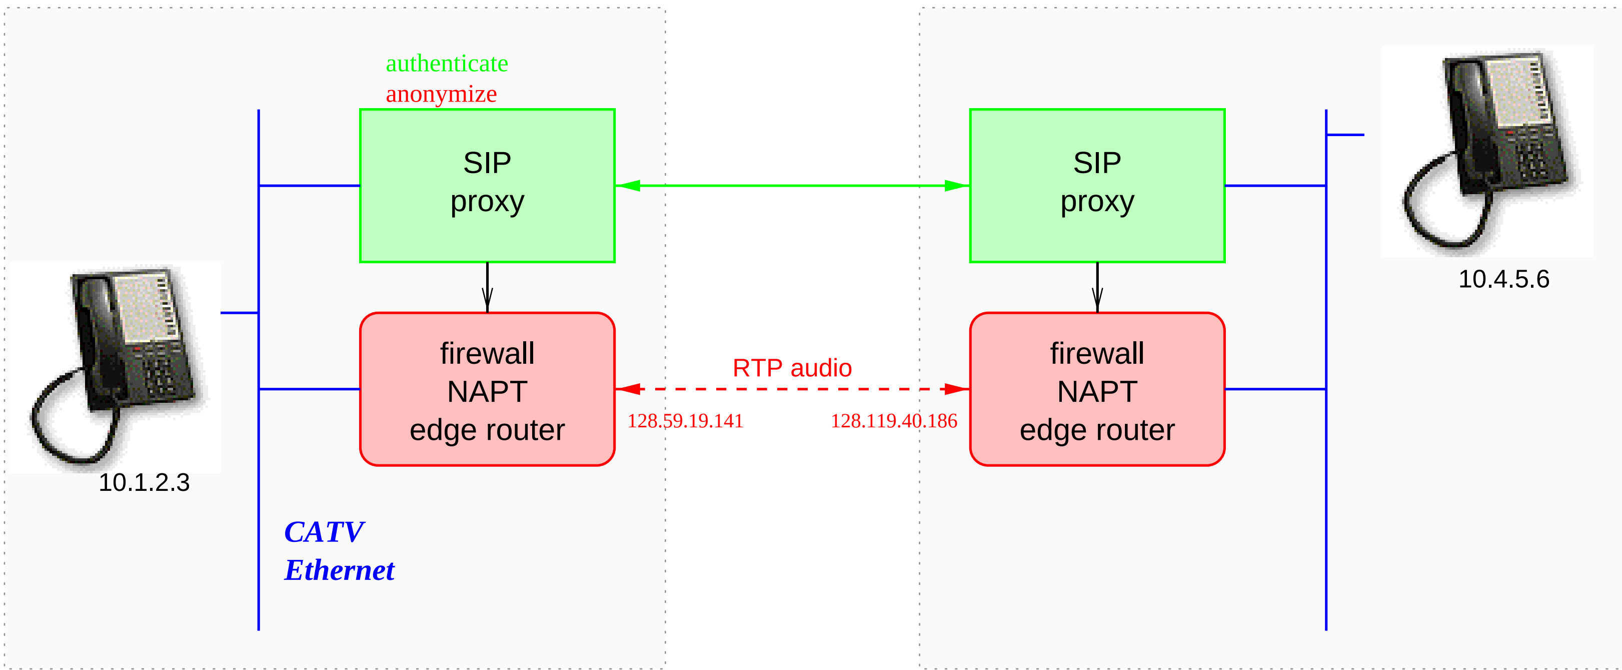

SIP addresses food chain

yellow pages		"president of the United States"
		↓ *WWW search engines*
common names		"Bill Clinton, Whitehouse"
		↓ *directory services*
host-independent		president@whitehouse.gov
	↙ *SIP*	↘ *SIP*
host-specific	sip:bubba@oval.eop.gov	sip:+1-202-456-1111@net2ph.com
	↓ *DNS*	
IP address	198.137.241.30	

SIP: basic operation

1. use directory service (e.g., LDAP) to map name to *user@domain*
2. locate SIP servers using DNS SRV, CNAME
3. called server may map name to *user@host*
4. callee accepts, rejects, forward ($\rightarrow$ new address)
5. if new address, go to step 2
6. if accept, caller confirms
7. …conversation …
8. caller or callee sends BYE

SIP–DNS interaction

extended email-like domain resolution ➠ try until success:

1. try SRV DNS record for "sip.udp" and "sip.tcp" in domain, with priority and weights for randomized load balancing
2. DNS CNAME or A record
3. may try SMTP EXPN command to get new address; goto (1)
4. if all else fails, send SIP request via MIME

SIP operation in proxy mode

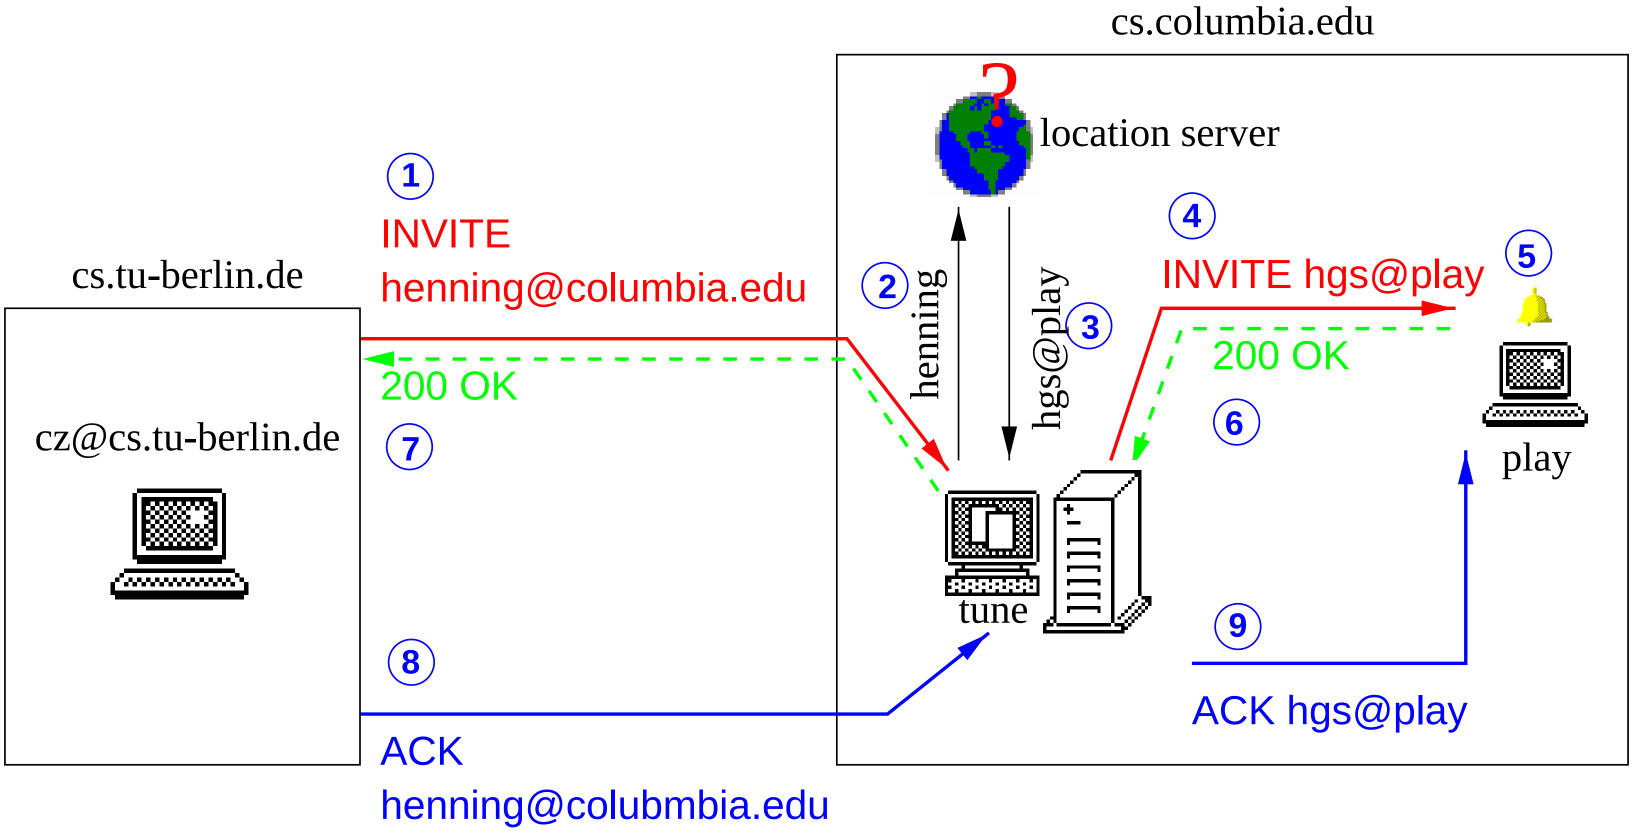

SIP operation in redirect mode

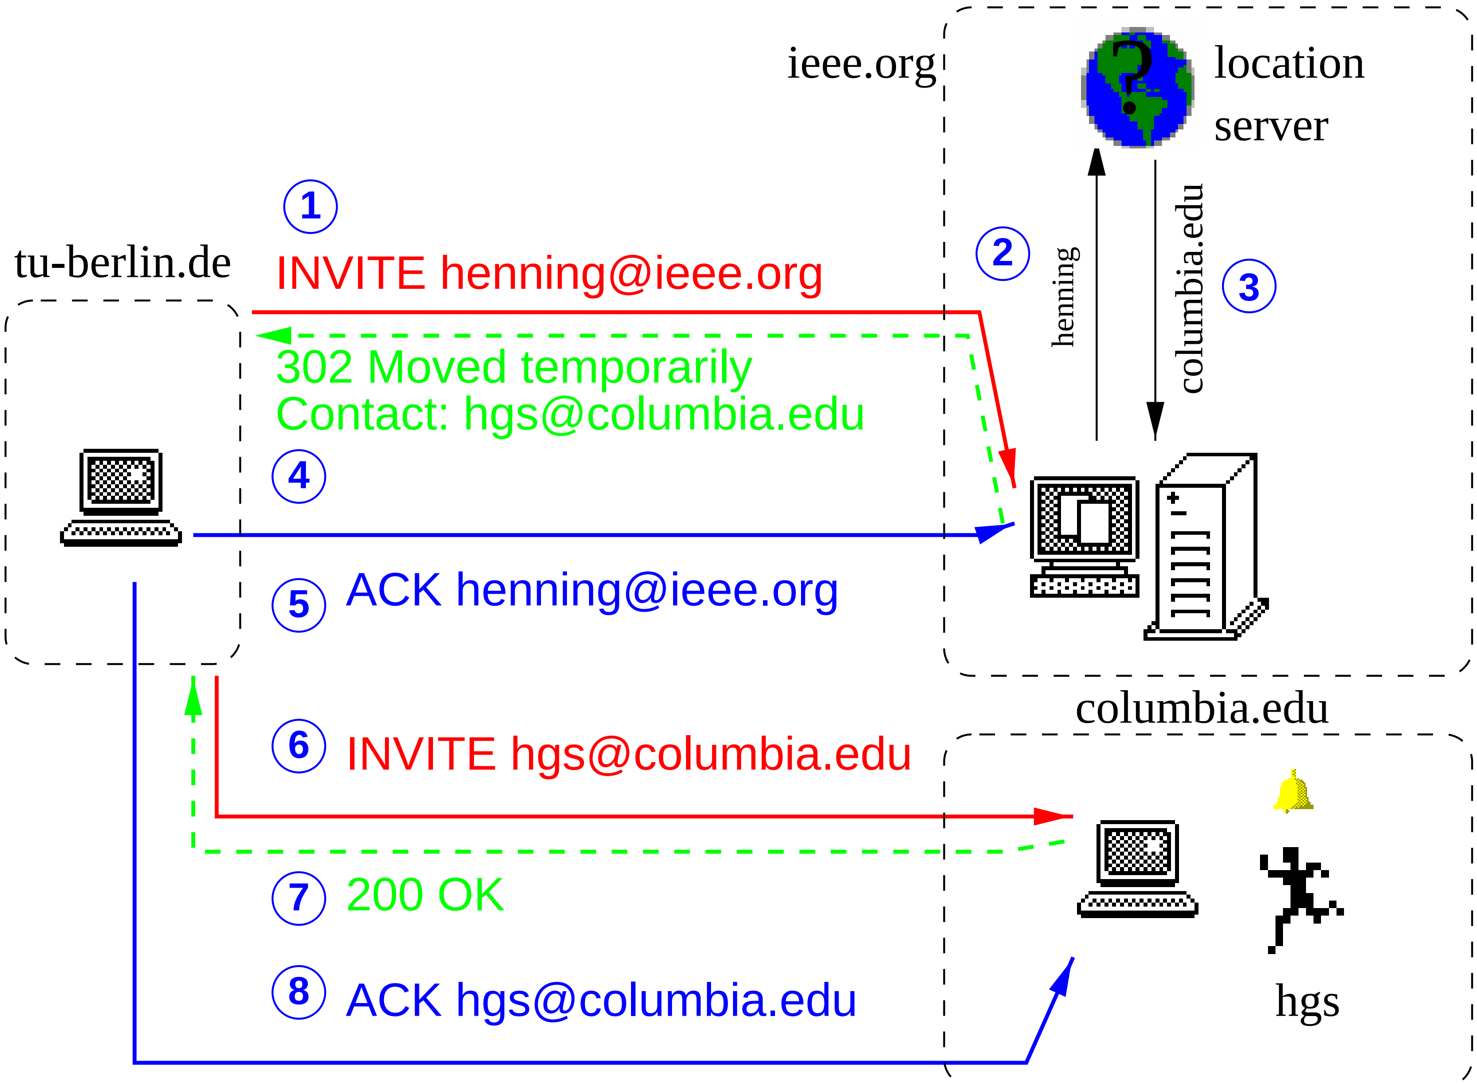

SIP operation in redirect mode

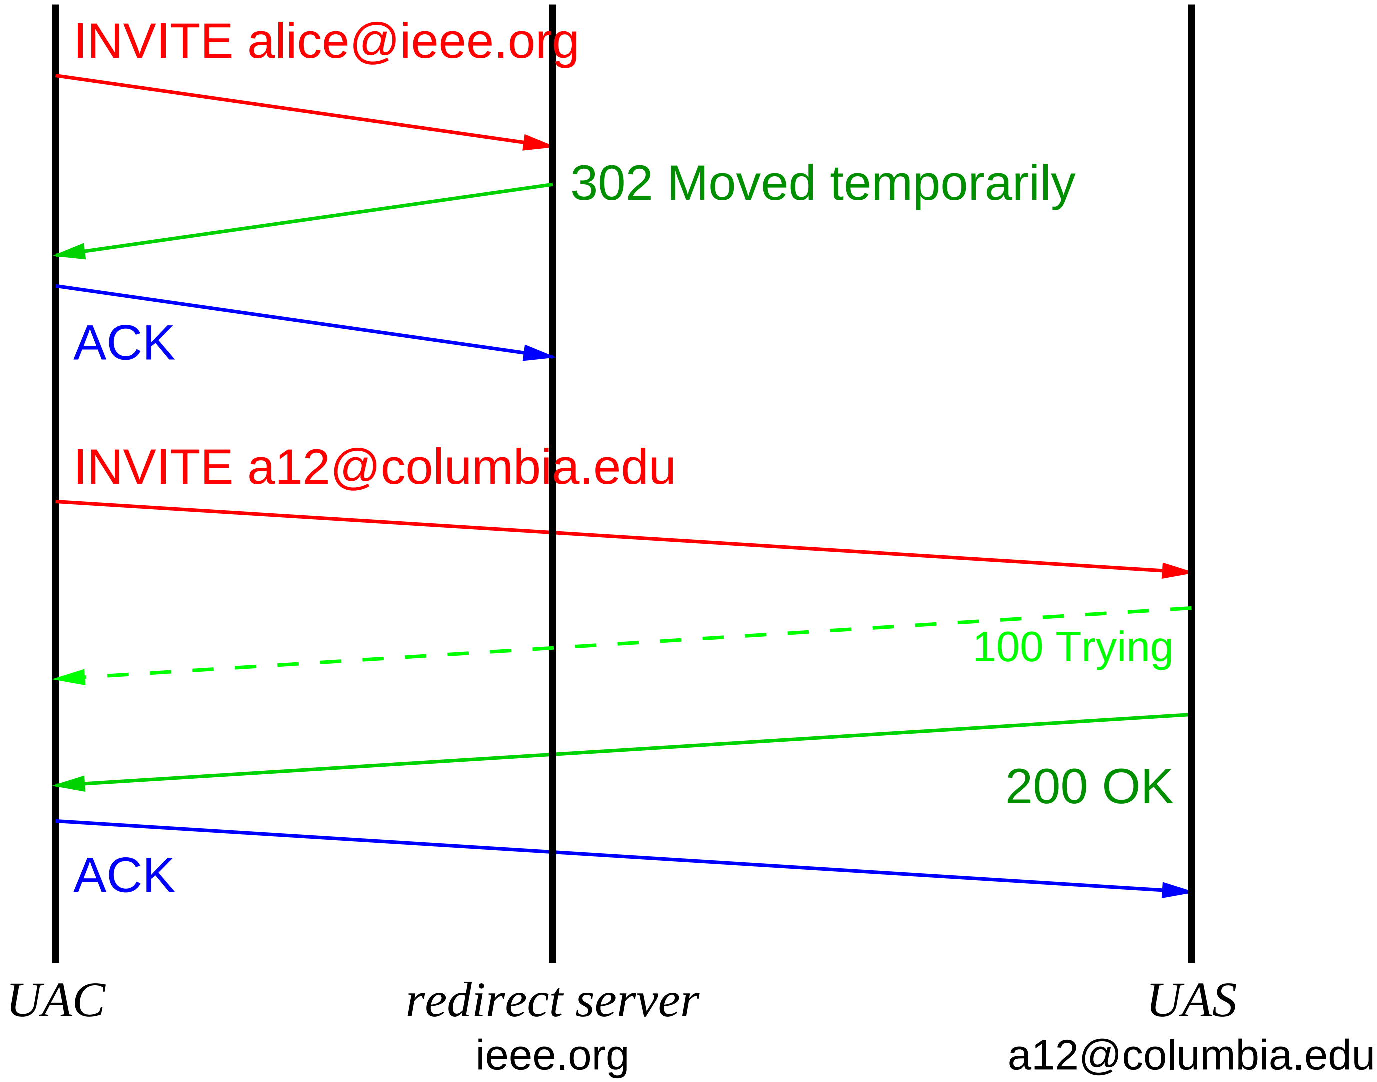

Interaction with resource reservation

avoid “fast busy” after ringing ➠ interleave

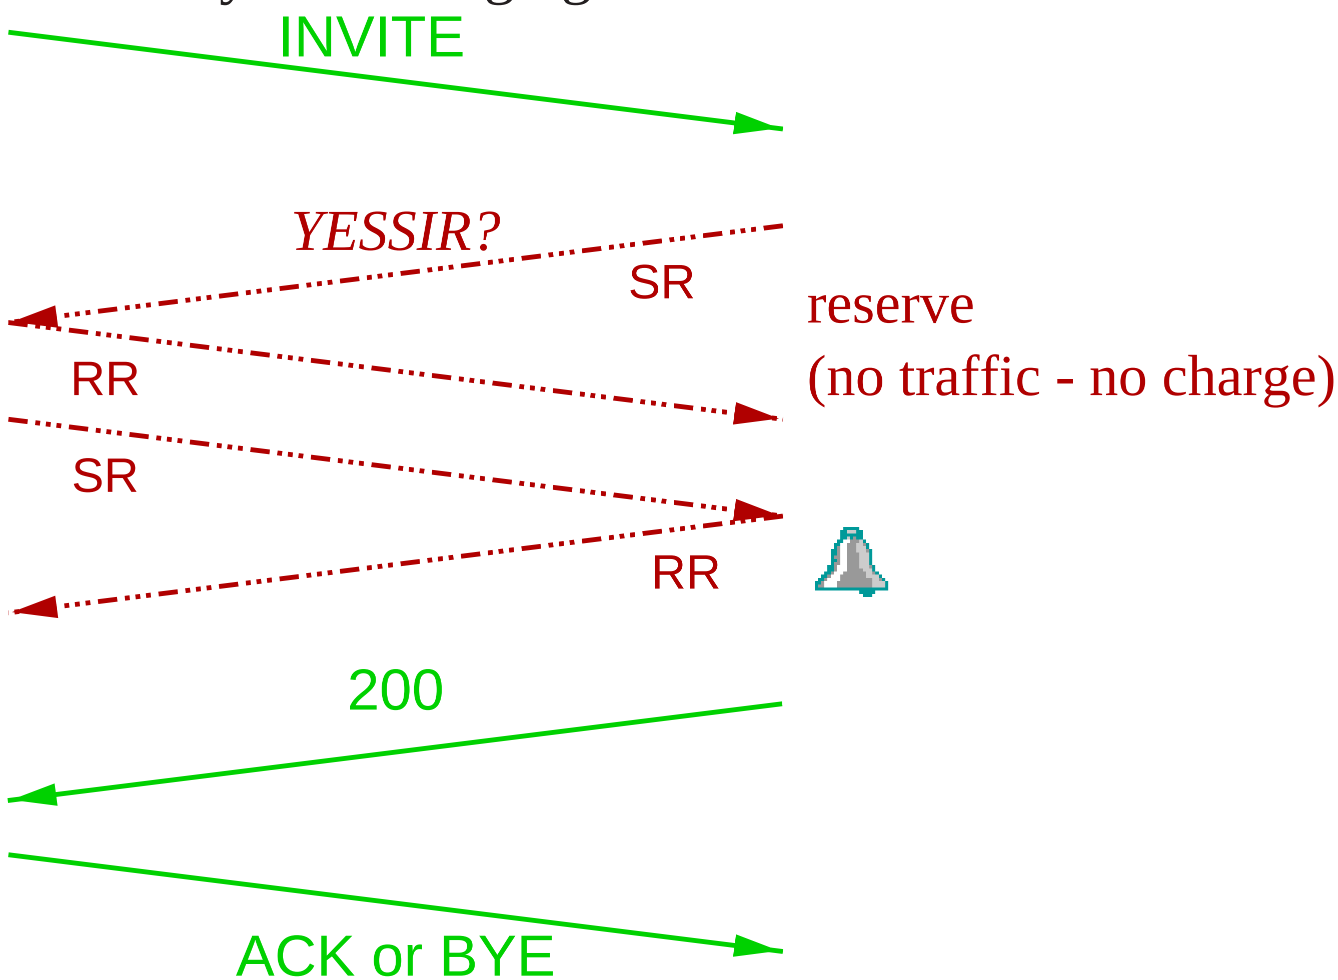

Invitation modes

invitation	conference	
	unicast	multicast
unicast	telephony	MBone session
multicast	reach first	dept. conference

➠ SIP for all modes, SAP also for multicast/multicast

SIP servers and clients

UAC: user-agent client (caller application)

UAS: user-agent server ➠ accept, redirect, refuse call

redirect server: redirect requests

proxy server: server + client

registrar: track user locations

often combine registrar + proxy or redirect server

Proxy and redirect servers

proxy: may *fork* requests ➠ parallel or sequential search

- *near-end proxy:* outgoing calls ➠ address lookup, policy, firewalls
- *far-end proxy:* closer to callee ➠ callee firewall, call path hiding

redirect server: lower state overhead, more messages

SIP requests and responses

- HTTP look-alike
- provisional and final responses:
  - 1xx = searching, ringing, queueing, . . .
  - 2xx = success
  - 3xx = forwarding
  - 4xx = client mistakes
  - 5xx = server failures
  - 6xx = busy, refuse, not available anywhere

SIP protocol request

```
INVITE sip:schulzrinne@cs.columbia.edu SIP/2.0
From: Christian Zahl <sip:cz@cs.tu-berlin.de>
To: Henning Schulzrinne <sip:schulzrinne@cs.columbia.edu>
Via: SIP/2.0/UDP 131.215.131.131, SIP/2.0 foo.com
Call-ID: 3678134014@cloud9.cs.tu-berlin.de
Content-Type: application/sdp
Content-Length: 187
CSeq: 8348 INVITE
Subject: New error codes
```

session description

SIP URLs

sip:[*user:pw@*]*host*:[*port*]
;transport=UDP;maddr=224.2.0.1

- used in Request-URI, Contact headers (redirect, registration), web pages
- transport and maddr specify transport
- can specify methods, header and body in web pages, email
- example: `sip:a.g.bell@belltel.com`

SIP/RTP mobility overview

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

Review: DHCP

1. client broadcasts `DHCPDISCOVER`
2. server offers address: `DHCPOFFER`
3. client broadcasts `DHCPREQUEST`
4. server acks via `DHCPACK`

SIP 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)

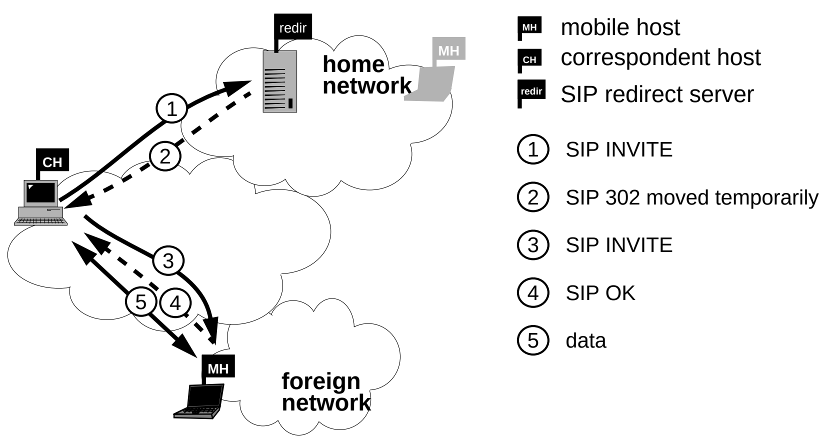

SIP mobility: mid-call

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

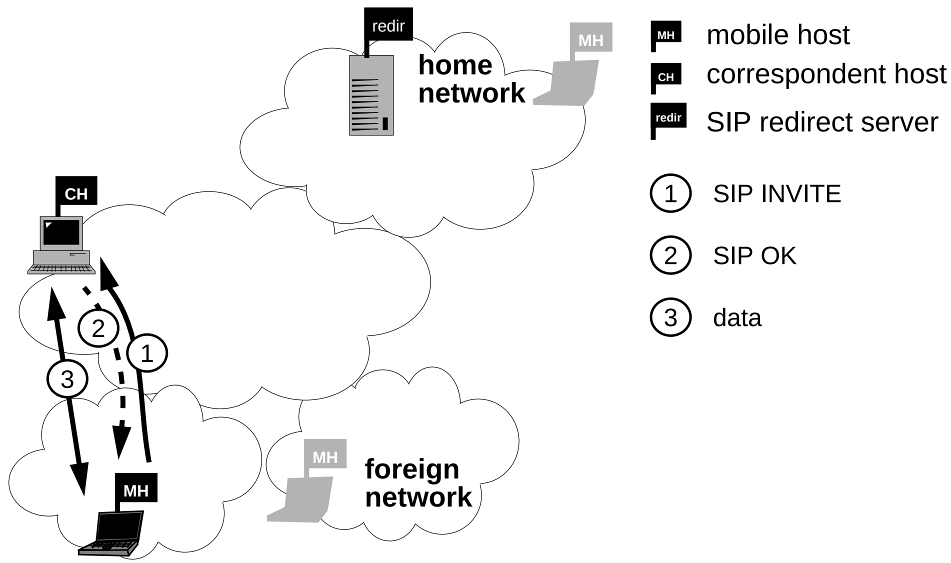

SIP mobility: simultaneous moves

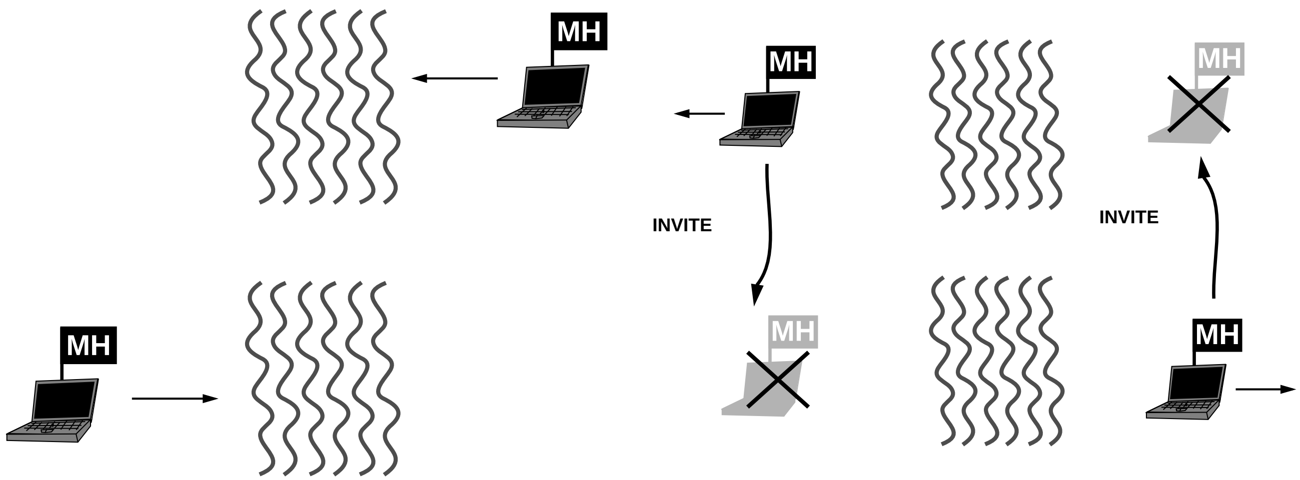

➠ need addition to protocol behavior: try well-known address

SIP mobility: security

- telephone-like: random Call-ID, sequential CSeq
- password: HTTP “basic” authentication ➠ plain-text password no better than Call-ID
- challenge-response: HTTP “digest” authentication ➠ possibly more messages for updates
- true crypto: PGP, S/MIME
- possibilities: shared secret, ipsec, Lamport’s hash (OTP)

SIP mobility: multi-stage registration

Send this to local registrar, `ca.move.com`

```
REGISTER sip:ny.move.com SIP/2.0
Contact: sip:me@ca.move.com
  ;expires=3600
```

- gets proxied like INVITE to `nyc.move.com` ➠ hierarchical LR
- only updates home registrar when MH leaves Calif.
- first INVITE goes to `ca.move.com`

Packet loss recovery

- CH's packets sent to old MH address lost
- MH indicates seq. no. of last packet received in SDP
- CH retransmits missing packets
- artificial silence period?
- proxy recovery?

RTP mobility support

- RTP uses 32-bit random *synchronization source* (SSRC)
- stays constant even if IP address changes
- RTCP SDES CNAME binds to global identity
- very fast, binding update retransmitted every 50 ms…
- loop/collision detection: if two sources with same SSRC, keep one
- need to modify RTP discard behavior: keep *new* IP address

RTP mobility

- CH accepts RTP packets from new IP address
- “mobile RTP” CH:
  - send own RTP packets (also) to the new address
  - allows smooth hand-offs
  - terminate double sending when receiving SIP update

Comparison SIP vs. IP mobility

	IP	SIP	RTP
constant	IP address	SIP url `user@home.com`	SSRC/SDES
update	binding upd.	re-INVITE	audio packet
security	via HA	call-ID, crypto	crypto
impl.	OS	app.	app.
reg.	one (*)	proxy	none

Open Issues/Problems

- latency: DHCP uses ARP for confirmation ➠ 150 ms latency
- soft handover ➠ multiple IP addresses?
- multi-stage location register
- interaction with RTP loop detection
- division of labor: do we need link-layer mobility?

Conclusion

- mobile data mostly done
- mobile IP: infrastructure changes, overhead
- Internet telephone mobility ➠ SIP
- mobile media-on-demand terminals: RTP or RTSP REDIRECT

More information

Internet and telecom statistics:
http://www.cs.columbia.edu/˜hgs/internet

Papers: http://www.cs.columbia.edu/˜hgs/research/irt

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

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

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