Industrial Strength and Mobile Internet Telephony

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

Siemens Visit (Munich)

October 27, 2000

Joint work with Jonathan Rosenberg, SIP IM/presence group, Telcordia, Columbia IRT research group

Overview

- industrial-strength VoIP and presence services
  - scaling
  - redundancy and fault tolerance
  - network management
  - administration
  - integration
- using SIP for supporting facets of mobility

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

- user agent = UAC + UAS
- often combine registrar + (proxy or redirect server)

Design Goals

- 5-nines reliability
- scalability to major domains like aol.com, siemens.com or t-online.de
- commodity unreliable hardware (PCs)
- commodity software for databases and directories
- avoid clustering software

Scaling

- SIP signaling primarily handled by SIP proxies, with associated registrars and location servers
- critical – common infrastructure for IM/presence, VoIP, conferences, mobile networks, . . .
- SIP proxies do not switch voice, but
  - route calls – mobility
  - implement policies
  - programmable logic
- far higher variability than classical switches: execute subscriber-defined code during call signaling:
  - sip-cgi scripts (similar to web cgi-bin scripts)
  - CPL scripts – XML-based call logic

Scaling

- call routing: no “area codes” ➠ email-style addresses, with all att.com through single (logical) proxy
- but: easier to scale due to higher signaling bandwidth
- transmission delay: 288 μs/message for 10 Mb/s Ethernet (typical: 360 bytes)

Scaling or How Many Calls can a SIP Switch Switch?

Some metrics:

- BHCA – 750,000 to 2.5 million busy hour call attempts for large class-5 switches = 3.6 ms/call
- AT&T: 280 million calls a day = 0.3 ms/call
- Yahoo: 780 million page views/day
- AOL: 110 million emails/day
- AOL: 500 million IM/day
- web server: about 1,500 to 3,000 static requests/second

Signaling Load Components

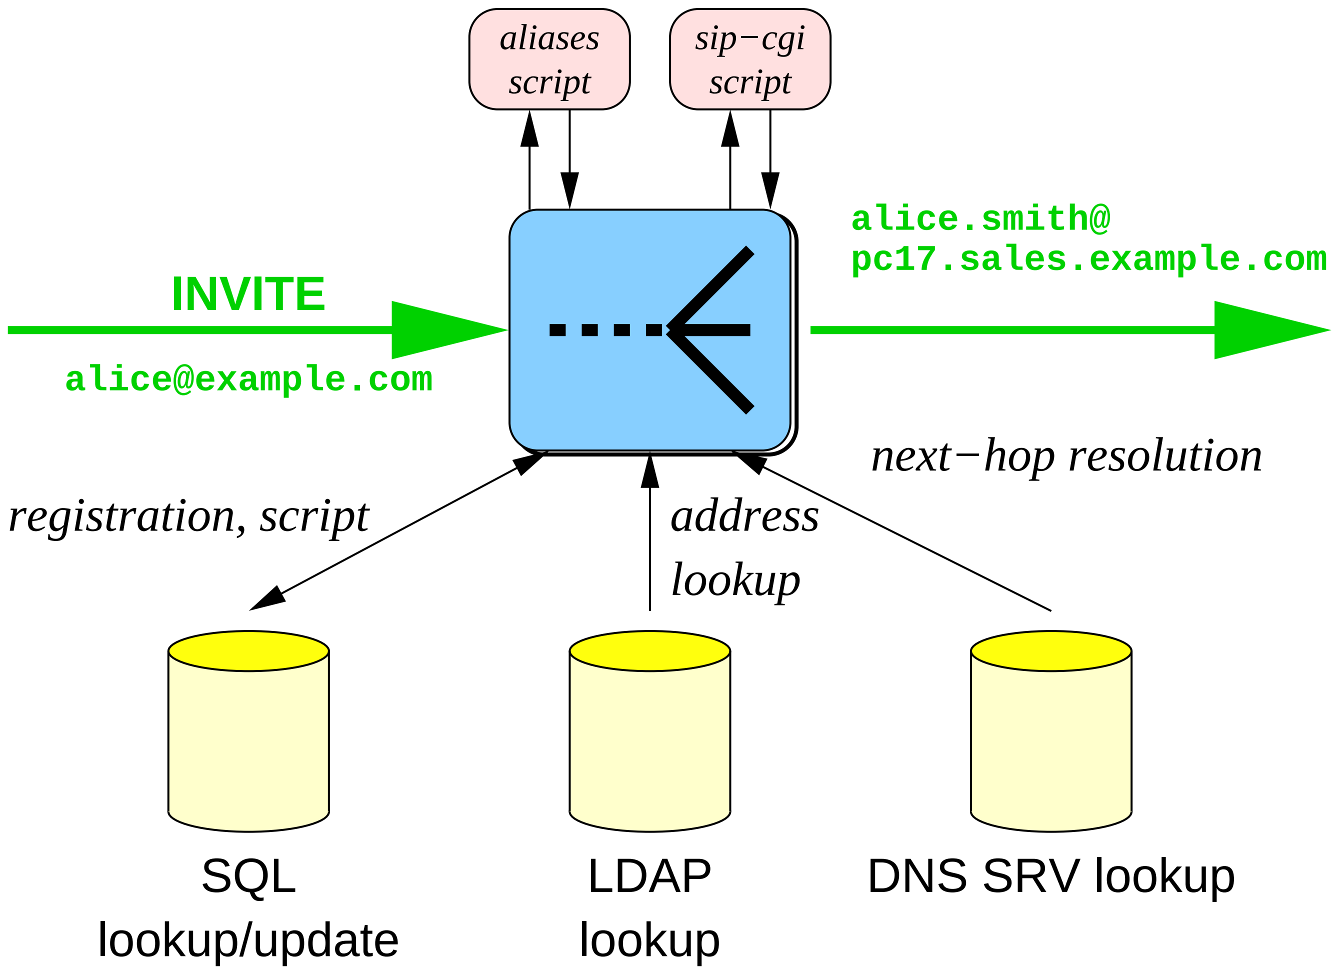

Typical Signaling Processing Steps

1. parse incoming SIP request
2. possibly invoke a generic administrative script
3. map aliases (e.g., `peter.ford` ⟶ `pf`) in local database to canonical identifier
4. check registration in LDAP or via SQL query
5. invoke per-user cgi script
6. translate host name
7. forward request, response
8. log request

SIP Scaling Differs From Other Internet Protocols

- not CPU-bound ➠ delay $\neq$ 1/throughput
- low byte volume ➠ easy to physically distribute for redundancy and load distribution
- servers can easily be shared among domains

Signaling Load Distribution

ease depends on service model: SIP proxy, redirect, registrar

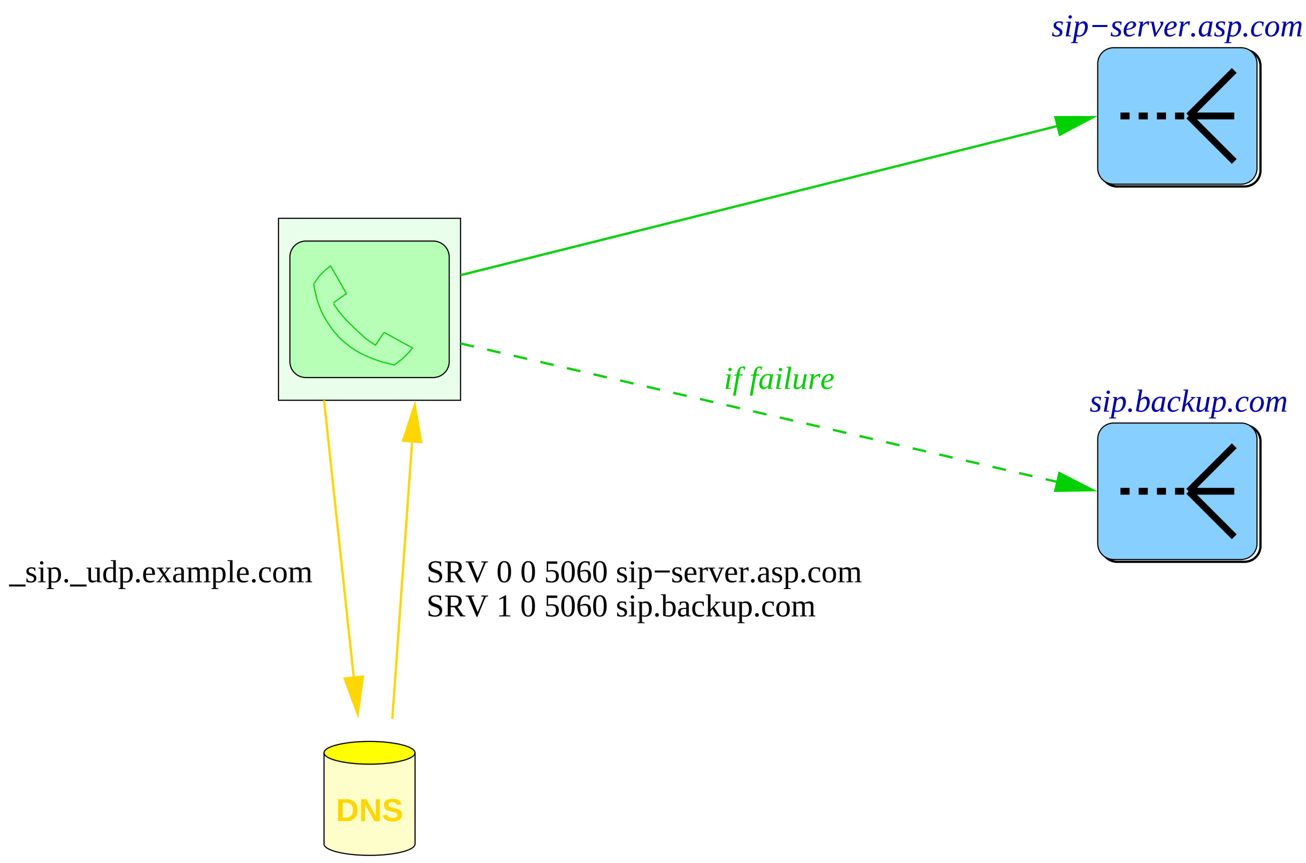

DNS SRV Records

- DNS SRV records: priority and weight

```
_sip._tcp        SRV 0 0 5060 sip-server.cs.columbia.edu.
                 SRV 1 0 5060 backup.ip-provider.net.
_sip._udp        SRV 0 0 5060 sip-server.cs.columbia.edu.
                 SRV 1 0 5060 backup.ip-provider.net.
```

- clients try hosts in order of priority, then balance requests randomly scaled according to weight

Signaling Load Distribution

- does *not* take current load into account
- hot spots?
- SIP allows per-transaction routing of requests, with Route header for routing subsequent transactions
- Route can be either specific domain or IP address OR SRV
- proposal to allow Route also for first request
- if call state, more difficult to fail-over mid-call ➠ need back-end state synchronization

Other Load Components

Full characterization requires dimensioning other servers:

- SQL or in-memory databases for authentication and registration
  - storage requirement depends on Contact length
  - from $\approx$ 50 to 1,000s bytes/client
- LDAP servers – about 180 searches/second?
- media servers for voicemail and IVR
- conferencing servers – primarily media/computation-limited

With roughly hourly SIP registration updates, writes can dominate – campus with 20,000 devices ➠ 5.5 updates/second

Fault Tolerance

- failure of proxies does not affect (most) existing calls
- possible exceptions: firewall proxies
- mid-call requests via Route can use different server, if DNS SRV used as address
- registration information:
  - is refreshed roughly hourly
  - multicast
  - forking registrations
  - our SLP synchronization work?
  - recovery after reboot ➠ persistent memory
- PSTN gateway location ➠ TRIP

Administration

- phone administration across platforms
- local user registration
  - anybody can register
  - web page
  - inherit from other database (AAA, RADIUS, LDAP, /etc/passwd, ...)

Administering Authentication

- PGP or S/MIME certified by third party
- carrier-based authentication, signed by proxy ➠ “DT certifies that this customer is called Lieschen Müller” or “this caller is calling from the premises of Visa”
- per-callee user name(s) and passwords: “friends/secret”
- per-domain identities with global identifiers

Example: Columbia Internet Extensible Multimedia Architecture

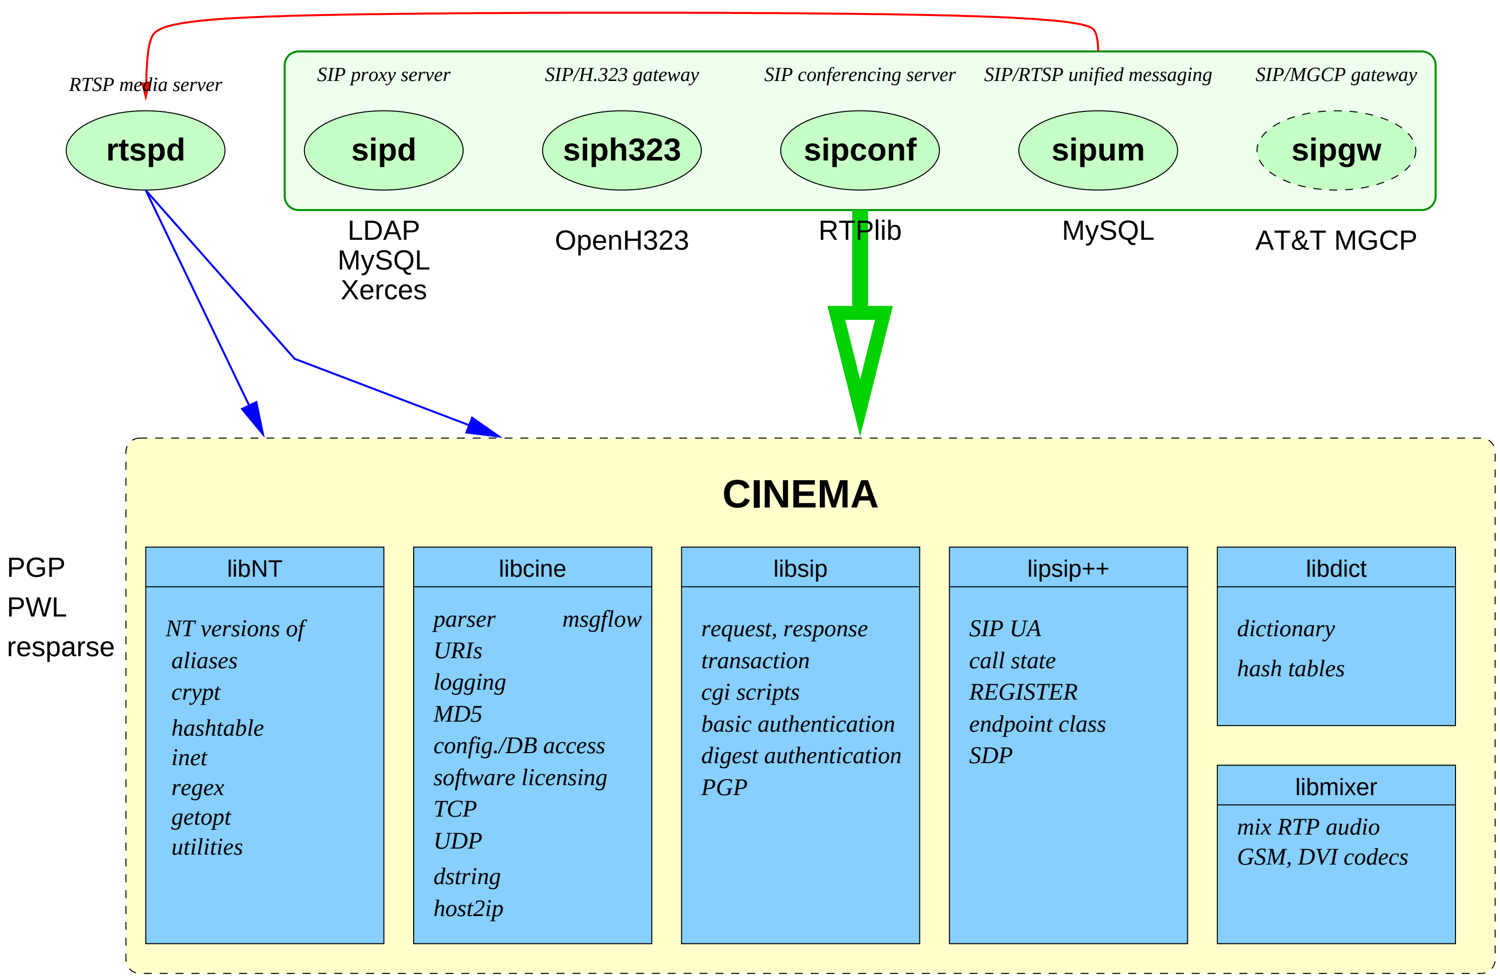

Single Sign-On

Uses per-domain identities

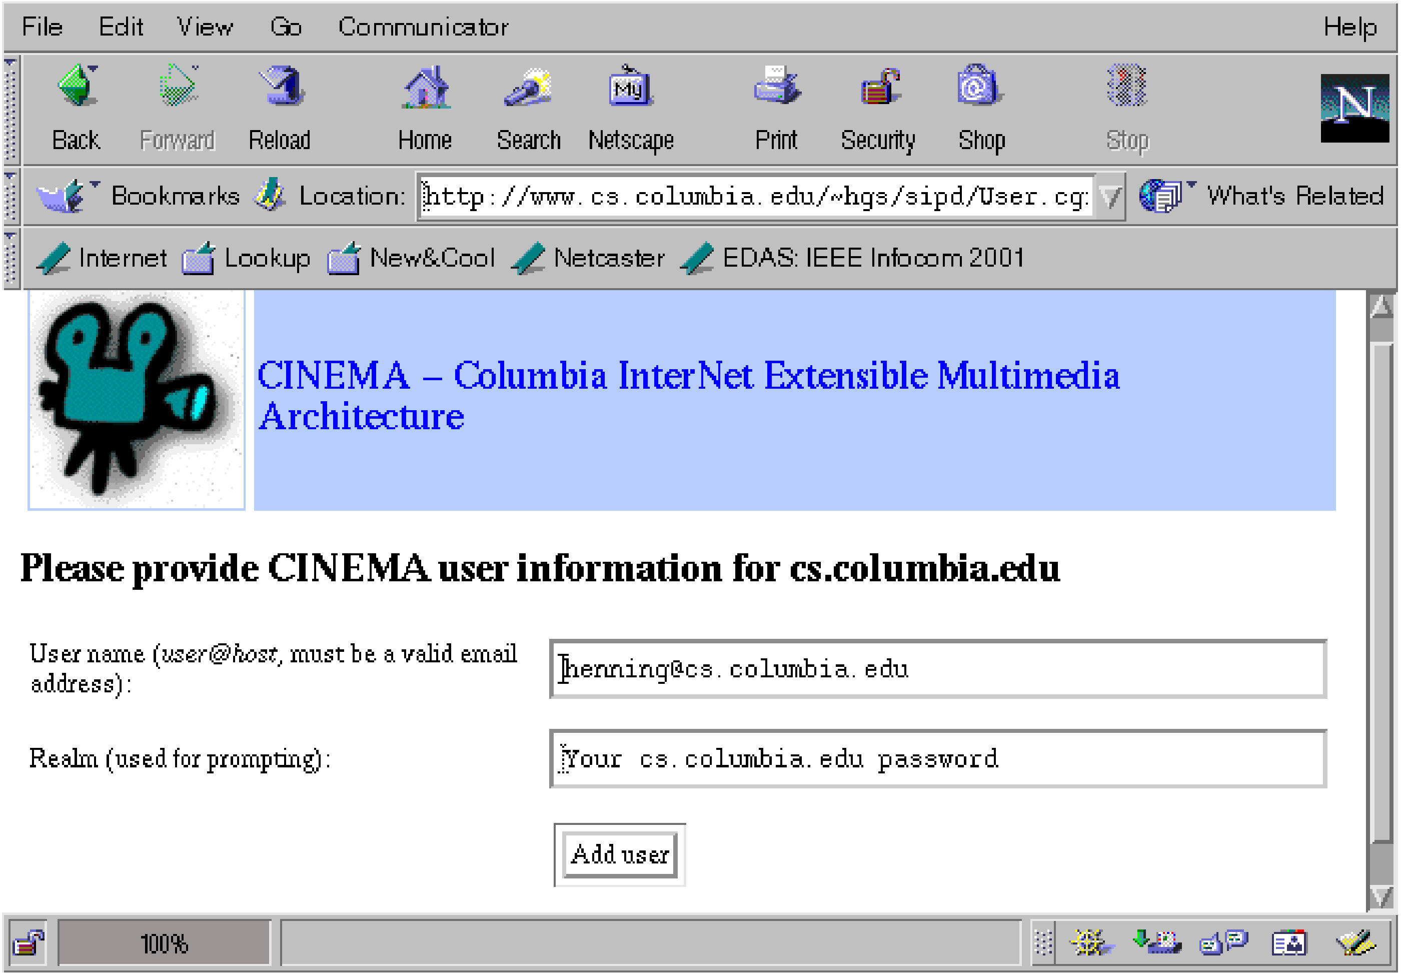

CINEMA Registration

Email send to `henning@cs.columbia.edu`:

```
Subject: Your CINEMA registration
Date:    Tue, 24 Oct 2000 21:48:09 -0400 (EDT)
From:    <CGI.script.-.do.not.reply@cs.columbia.edu>
To:      henning@cs.columbia.edu

Your new CINEMA password for cs.columbia.edu is
"deduct.transversal.desert".
The realm is "Password for cs.columbia.edu".
```

User Administration

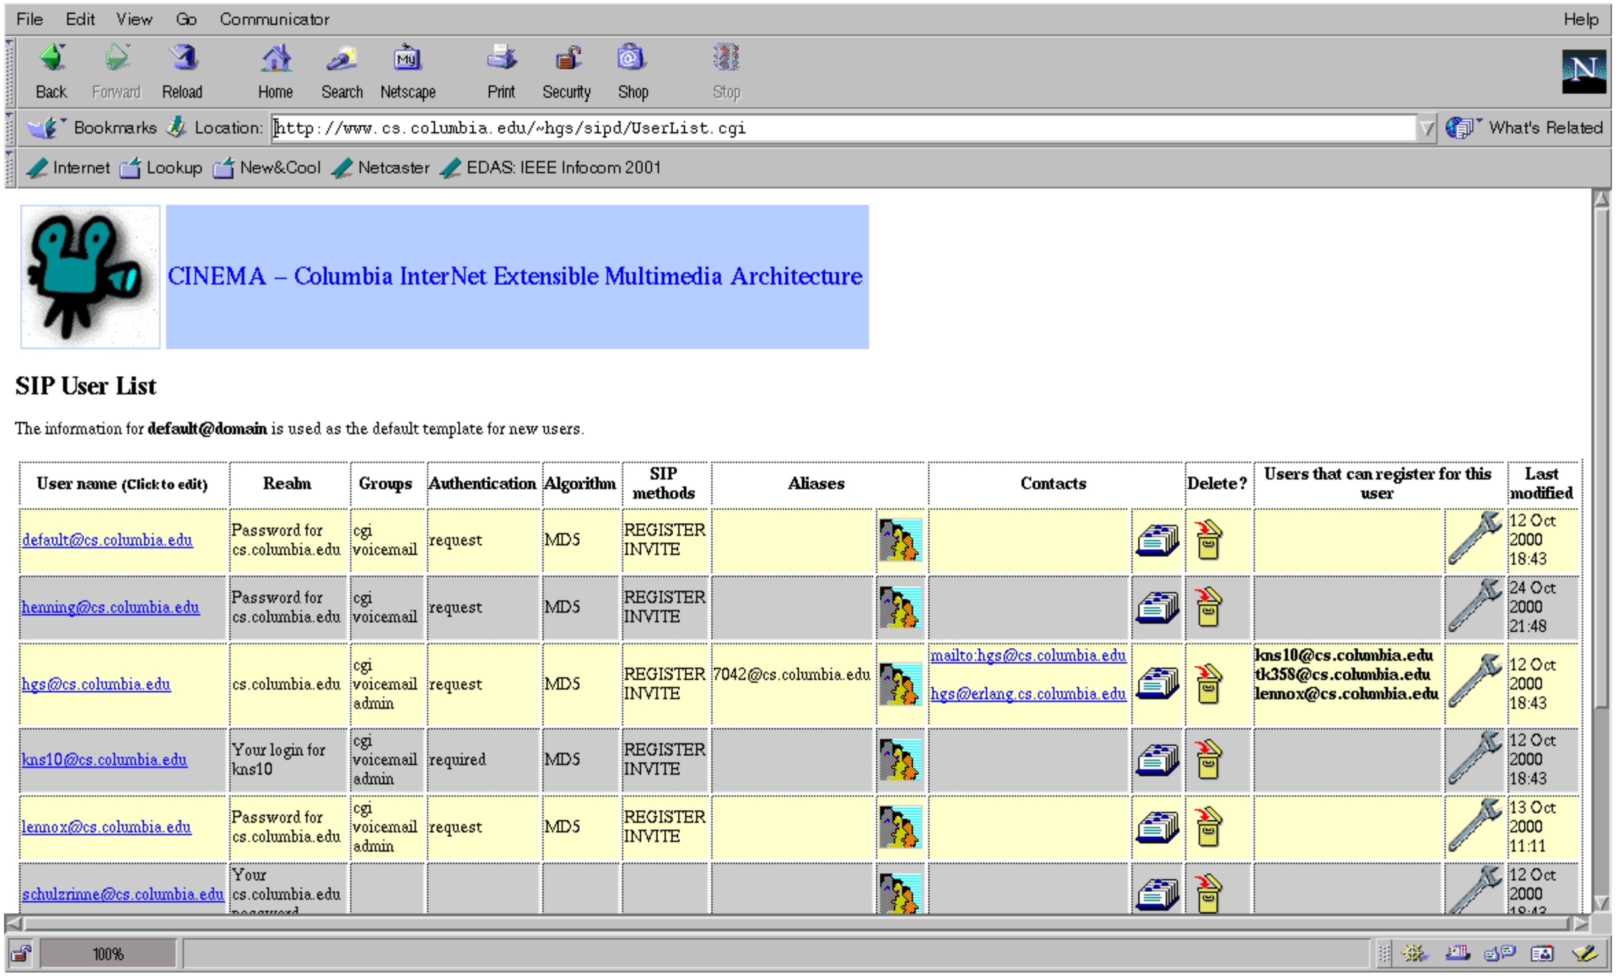

SIP User List

The information for **default@domain** is used as the default template for new users.

User name (Click to edit)	Realm	Groups	Authentication	Algorithm	SIP methods	Aliases	Contacts	Delete?	Users that can register for this user	Last modified
default@cs.columbia.edu	Password for cs.columbia.edu	cgi voicemail	request	MD5	REGISTER INVITE					12 Oct 2000 18:43
henning@cs.columbia.edu	Password for cs.columbia.edu	cgi voicemail	request	MD5	REGISTER INVITE					24 Oct 2000 21:48
hgs@cs.columbia.edu	cs.columbia.edu	cgi voicemail admin	request	MD5	REGISTER INVITE	7042@cs.columbia.edu	mailto:hgs@cs.columbia.edu hgs@erlang.cs.columbia.edu		**kns10@cs.columbia.edu** **tk358@cs.columbia.edu** **lennox@cs.columbia.edu**	12 Oct 2000 18:43
kns10@cs.columbia.edu	Your login for kns10	cgi voicemail admin	required	MD5	REGISTER INVITE					12 Oct 2000 18:43
lennox@cs.columbia.edu	Password for cs.columbia.edu	cgi voicemail admin	request	MD5	REGISTER INVITE					13 Oct 2000 11:11
schulzrinne@cs.columbia.edu	Your cs.columbia.edu									12 Oct 2000

CINEMA Policies

- third-party registration: “Anne and Bob are allowed to register for me”
- execution of scripts
- services: voicemail, conferencing, . . .

Network Management for SIP Servers

- SIP MIB, `draft-ietf-sip-mib-01`
- configuration description (outbound proxy, . . .)
- request statistics, method statistics
- current transactions
- working on initial implementation in sipd server

Mobility in an IP environment

Roaming users: logging in away from home network: hotel, home office

Terminal mobility: terminal moves between subnets

Personal mobility: different terminals, same address

Service mobility: keep same services while mobile

Session mobility: move active session between terminals

Simple Mobility: Roaming Users

- users visit other networks: laptop, PDA, hotel phone, ...
- want to maintain external identity
- usually, just pass IP address to home registrar
- difficult if firewalls and NATs
  - requests need to use local proxy
  - thus, need to register locally
- also may want to use home services while traveling

Roaming Users

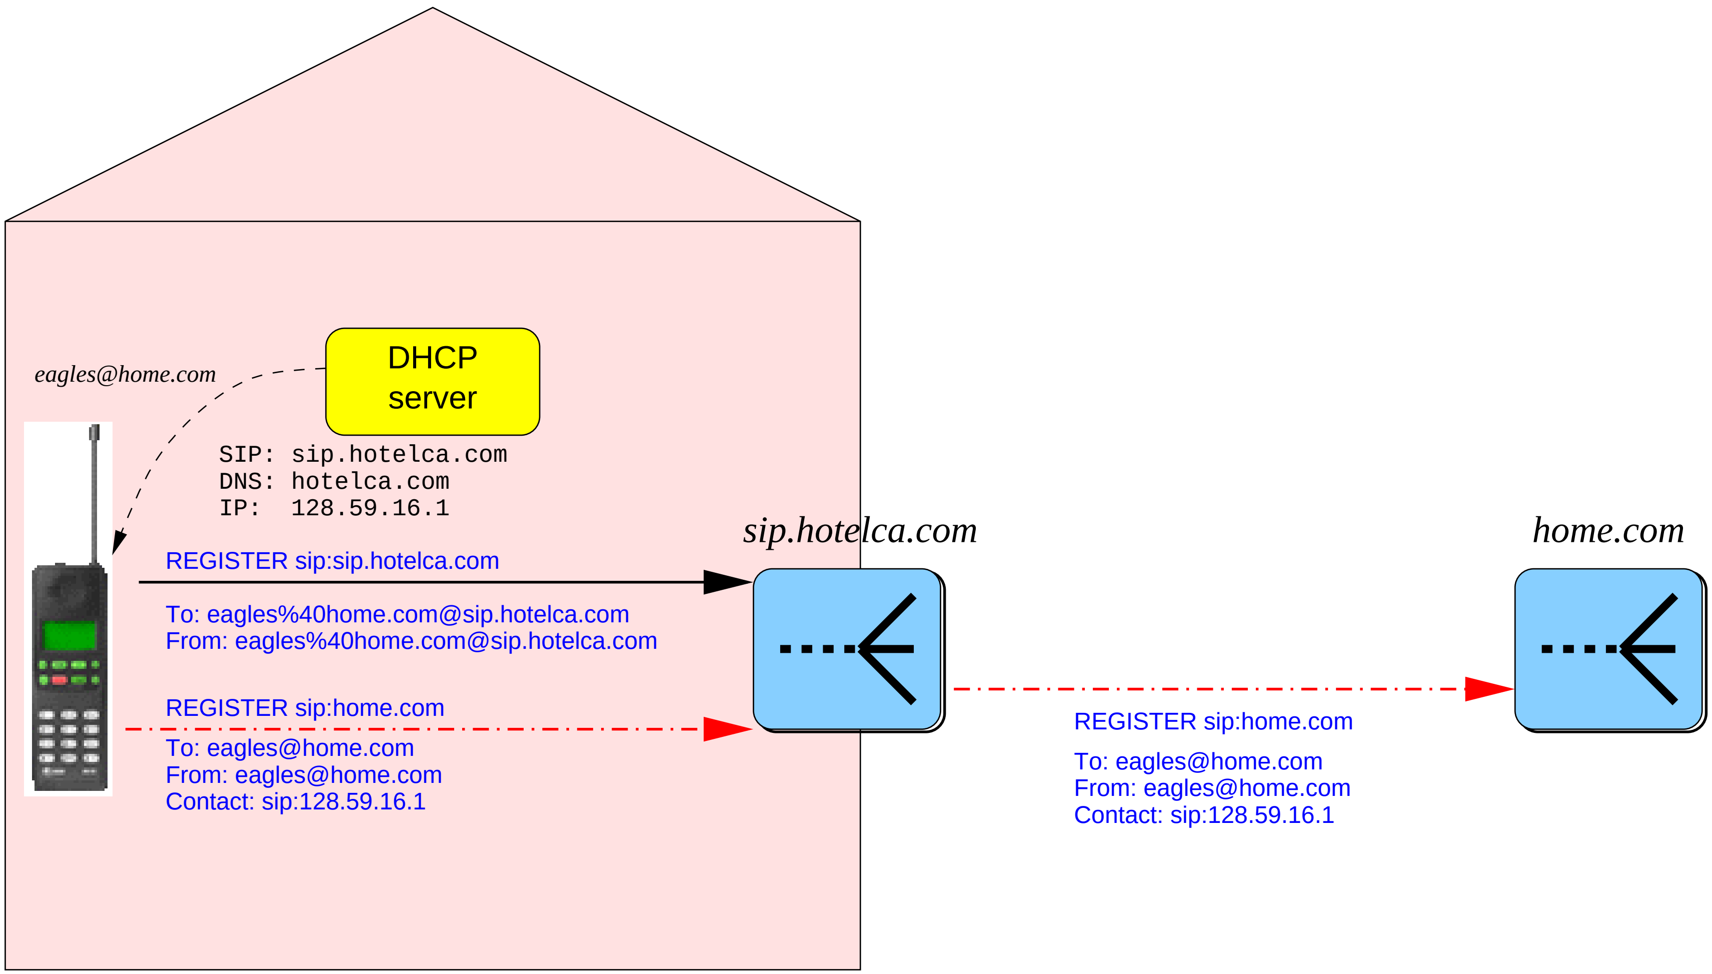

Terminal mobility

- 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

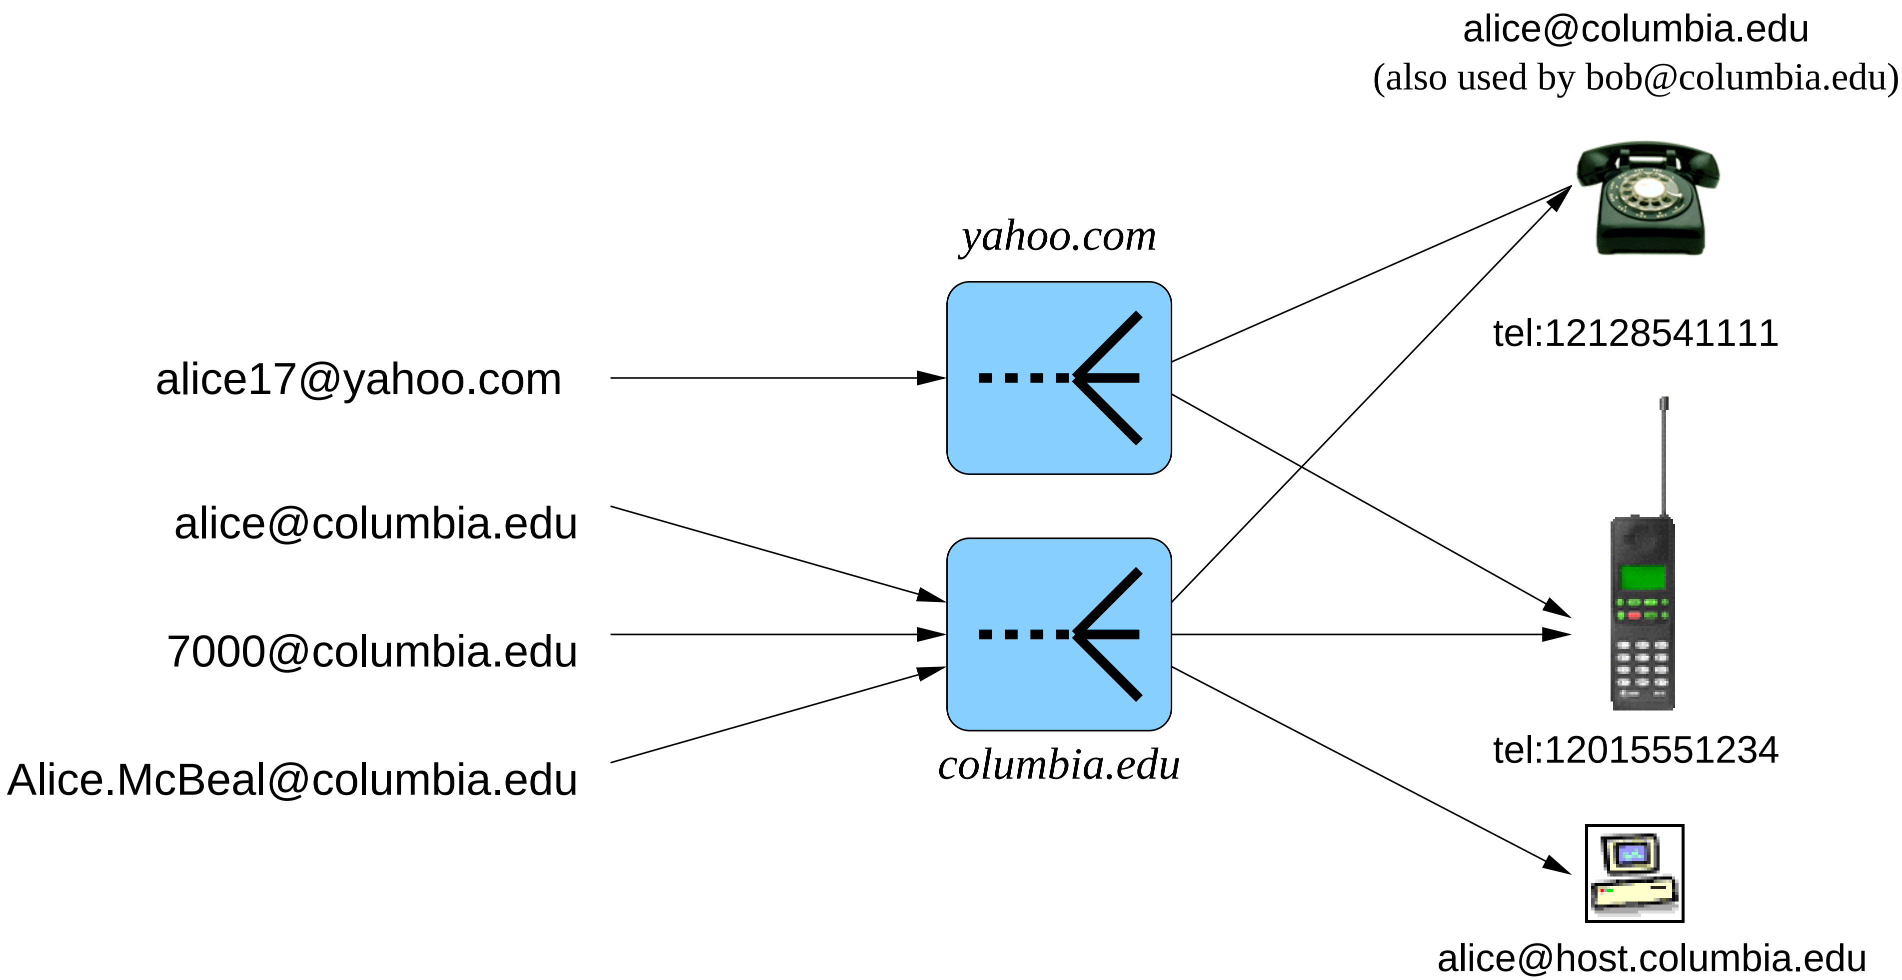

Personal mobility

- 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

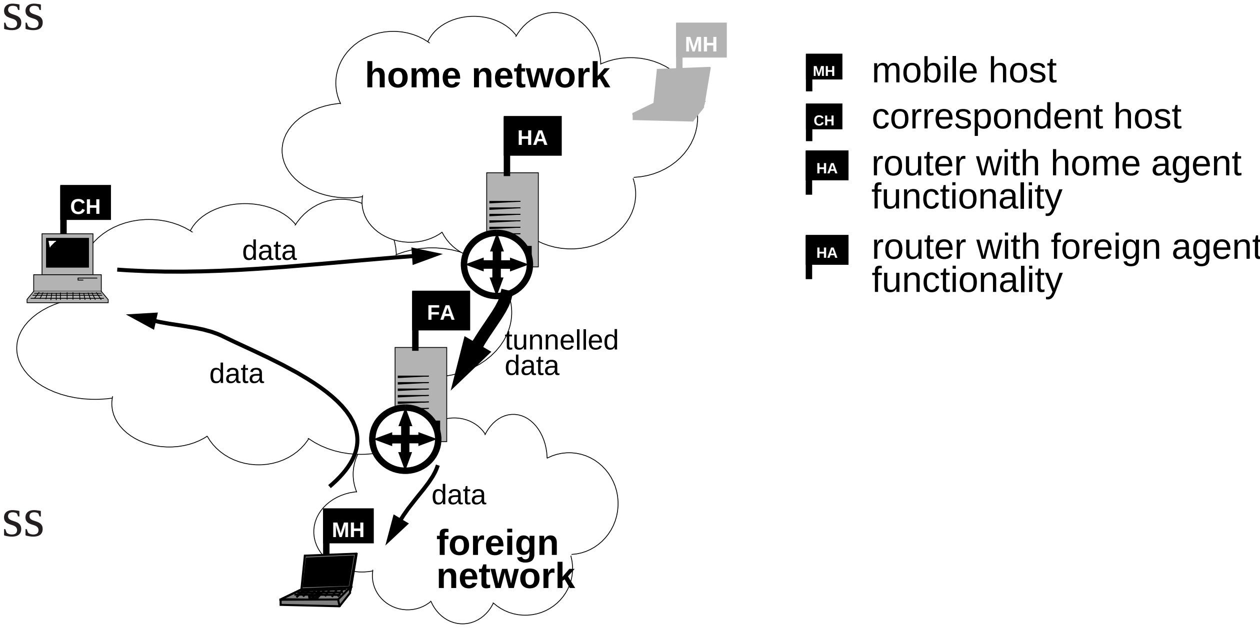

SIP terminal mobility overview

- 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)

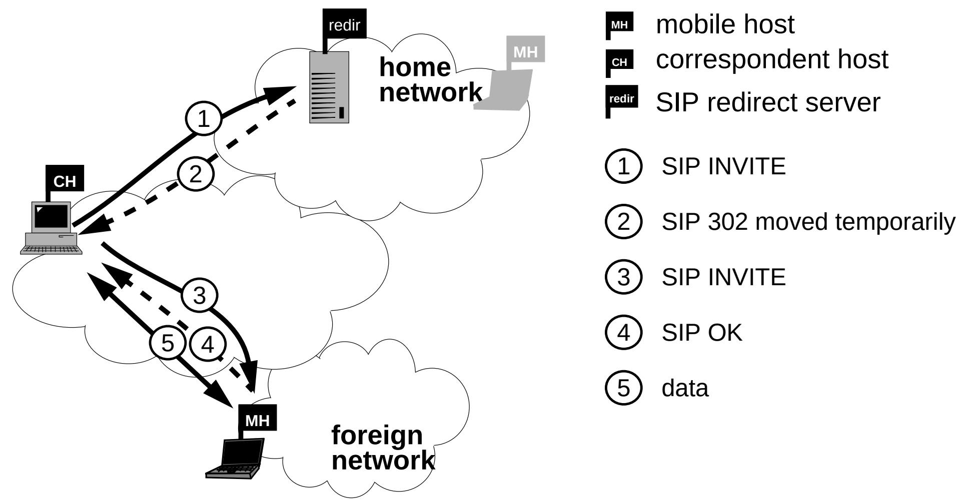

SIP terminal mobility: mid-call

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

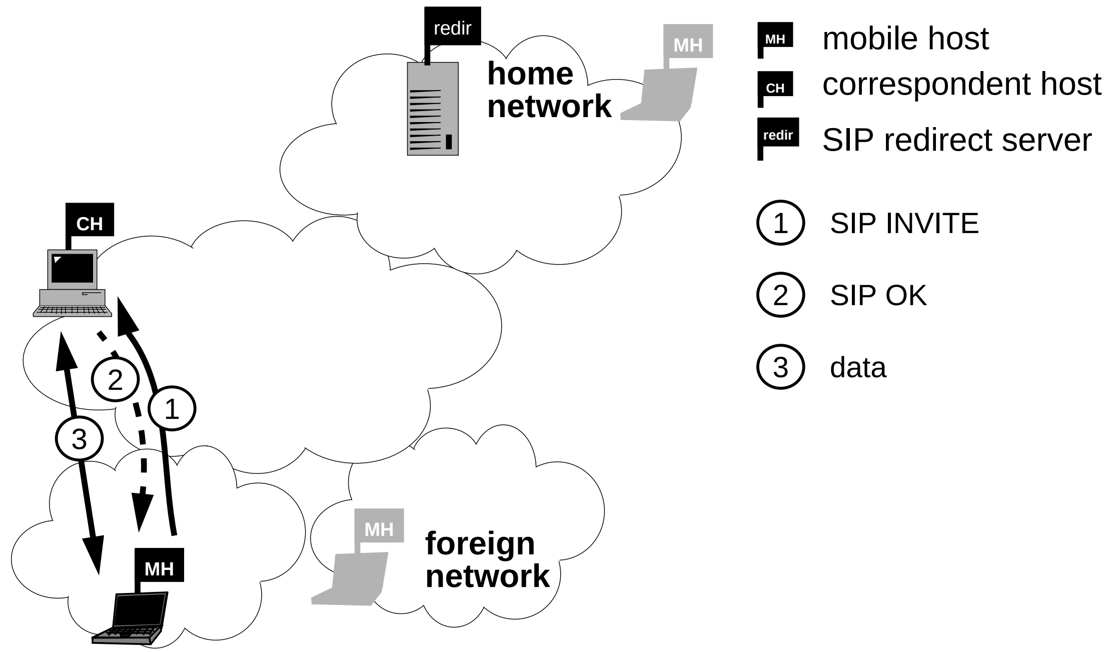

SIP terminal mobility: multi-stage registration

Don't want to bother home registrar with each move

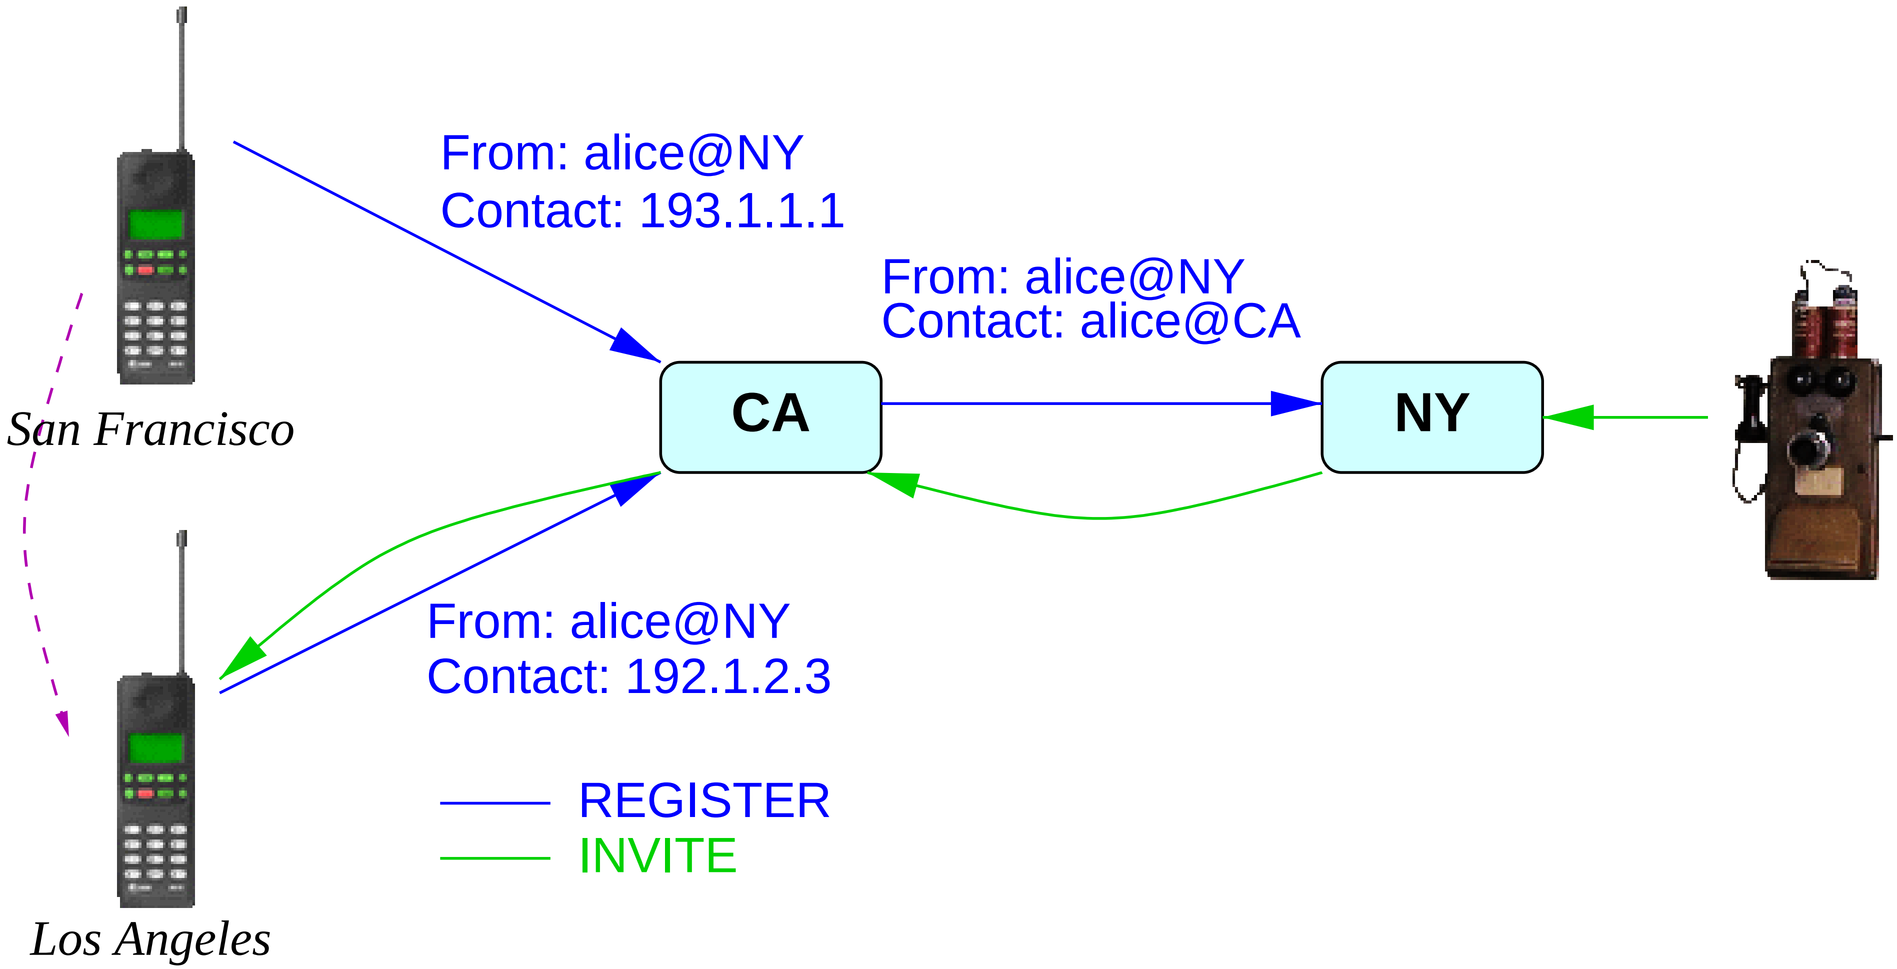

SIP and mobility: issues

- 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

Scaling & Reliability: Open Issues

- performance of real servers
- design alternatives: thread models, `select()`, etc.
- external server access models vs. in-memory databases
- impact of security
- single sign-on
- cryptographic certificates
- fail-over, state recovery

Mobility: Open Issues

- hand-off performance
- simultaneous moves
- address hiding?
- co-existence with mobile IP
  - hand-off to non-MIP networks
  - avoiding IPv4 dog-legged routing for multimedia

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