Signaling for Internet Telephony, Conferencing and Media-on-Demand

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

AT&T Labs

January 20, 1997

Overview

- why Internet telephony?
- modes
- signaling functionality
- conference initiation: SIP
- stream control: RTSP'
- open issues

Switching costs

switching method	capacity (Gb/s)	cents/kb/s
10BaseT Ethernet hub	0.12	0.5
10BaseT Ethernet switch	0.24	2.0
100BaseTX Ethernet switch	0.80	1.0
router	2.1	16.0
local ATM switch	2.48	1.6
PBX (256 lines)	0.02	218.
AT&T 4ESS toll (100k lines)	6.40	7.8
AT&T 5ESS local (107k lines)	6.85	156.

Major cost is cabling + interface (both telephone and optical), not switching ➠ single infrastructure

Internet telephony advantages

- need LAN/WAN anyway
- (currently) not encumbered by regulatory burdens (“the FCC bypass”)
- second phone ➠ reliability not as crucial
- no 7c/min access charge
- flat-rate billing
- cost of TAT transatlantic line is only $0.03/hour

Internet telephony features

- adaptive compression: trade $ for quality
- flexible source identification (“caller id”)
- easier computer-telephony integration
- multiple ccall presences ➠ fewer service interactions
- more than audio
- easy integration with email, WWW

Problems with Internet telephony

- long-distance capacity: AT&T: 150 mio. calls/day, 3 min. each (?) ➠ 20 Gb/s?
- “more than half is data”?
- predictable sound quality
- not suitable for dial-up users: signaling by email, telephone
- no advanced services
- 640 mio. phone lines $\leftrightarrow$ 180 mio. computers ➠ need interworking

Internet telephony modes

Connectivity:

- packets end-to-end: Mbone, Vocaltec, … ➠ requires continuous Internet connectivity
- tail-end hop off: use PSTN from POP

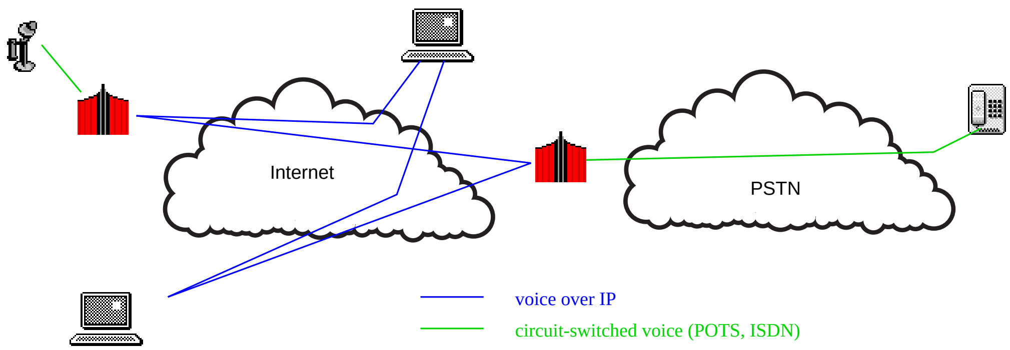

Bundling:

"Retail": individual packets per connection

"Wholesale" (trunking): interconnect PBXs or gateways

Signaling needs

Internet: separate protocols for separate functions ➠

routing OSPF, BGP, … $\surd$

quality of service: RSVP $\surd$

user location: mapping identifier or name to location (IP address)

call set-up/teardown:

- call forwarding: one, several, to answering machine
- call rejection
- parameter negotiation: media types, encodings, …

Internet multimedia stack

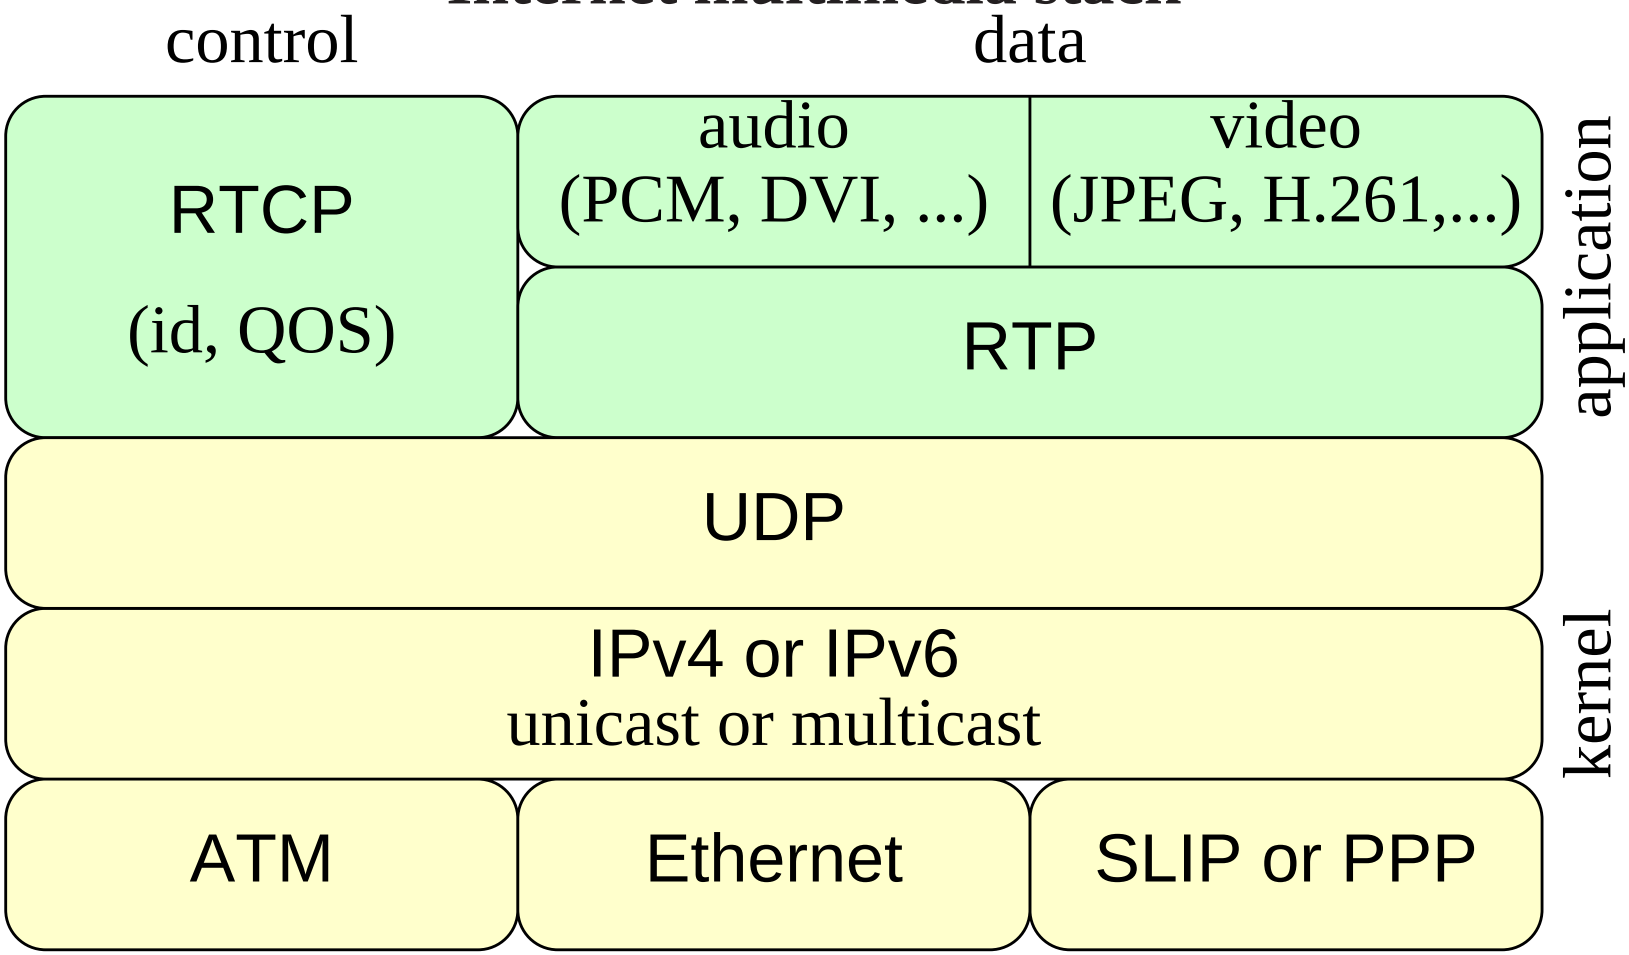

Internet conference model

- no central server(s) or multipoint control units
- hard to maintain consistent view of membership without central authority
- global view usually not necessary ➠ mostly cooperative participants
- "admission control" through encryption
- member $\in$ audio/video/…session $\in$ conference
- sessions $\equiv$ multicast group and port
- member $\equiv$ RTP CNAME (user@host) ➠ identification across sessions

Internet session model

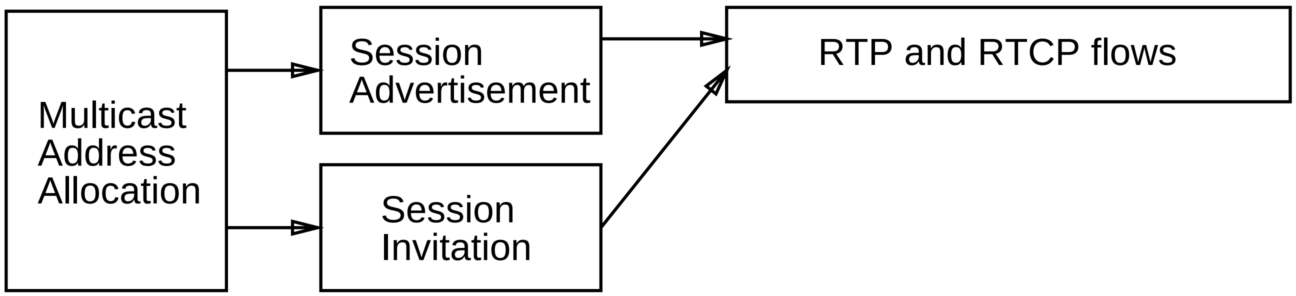

Lightweight Session Lifecycle

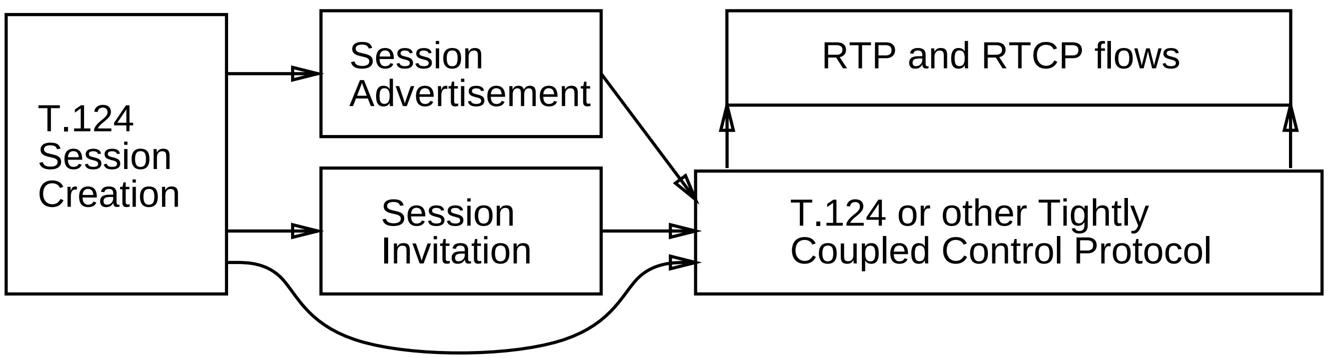

Tightly Coupled Session Lifecycle

Scenarios

1. call up another user ➠ “classical” telephony
2. invite new member to multicast conference
3. media-on-demand (unicast)
4. join on-going “broadcast”
5. invite media server or recorder to conference
6. device control: pan camera, adjust volume, . . .

H.323

- H.323 = Q.931 (ISDN signaling) + H.245 (capabilities) over TCP
- very complex (200+ pages)
- no multicast signaling
- no multicasst conferences (➠ MCUs)
- no redirection
- but: capability negotiation (H.245)
- limited media instances (one video, audio stream)
- no media servers
- but: industry support…

Conference invitation protocol

- don't need to know IP address of user's host
- many dial-up users do not have a permanent IP address
- invite to new and on-going conferences
- allow for manual and automatic forwarding
- personal mobility (complements data link/IP mobility) ➠ change of terminal, location
- reach first (load distribution) or reach all (department conference)
- possibly change permissible encodings during conference
- integrate media-on-demand servers and recorders
- separate issue: QoS ➠ RSVP (but: provide parameters)

Naming

feature	phone #	email
mnemonic	no (except 1-800-Flowers)	name + org.
multiple	no	easy
characters	≈ 12	22
location-independent	1-700	yes: j.doe@ieee.org
carrier $\neq$ naming	maybe	yes
directory	411, 1-555, switchboard.com	haphazard

Could also use URL or URN, but no established user naming

Universal Communication Identifier

email address: RFC 822 ➠ president@whitehouse.gov or Bill Clinton <president@whitehouse.gov> ➠ needs to be translated

user@host: user name at physical host

telephone number: E.164 ➠ +1.202.The.Pres or +1.212.555.1212@gateway.net2phone.com

Basic Operation

Example: call `schulzrinne@cs.columbia.edu`

- may use directory or user location service (e.g., LDAP)
- caller resolves `cs.columbia.edu` to actual host using DNS SRV or MX records
- send protocol (SIP) request to server
- server either . . .

1. accepts/rejects/forwards call, maybe after asking user (invisible to caller)
2. bridges to PSTN
3. locates user ➠ forwards call if not local ➠ redirect mode ➠ client tries again
4. acts as client ➠ proxy mode

SIP: Interaction with User Location + Email

- user location is a local matter (but: could use SIP to register)
- if all else fails: use email EXPN or VRFY to get address
- local server may use `.forward` information
- no SIP server found ➠ send MIME `application/sip` email

Operation in proxy mode

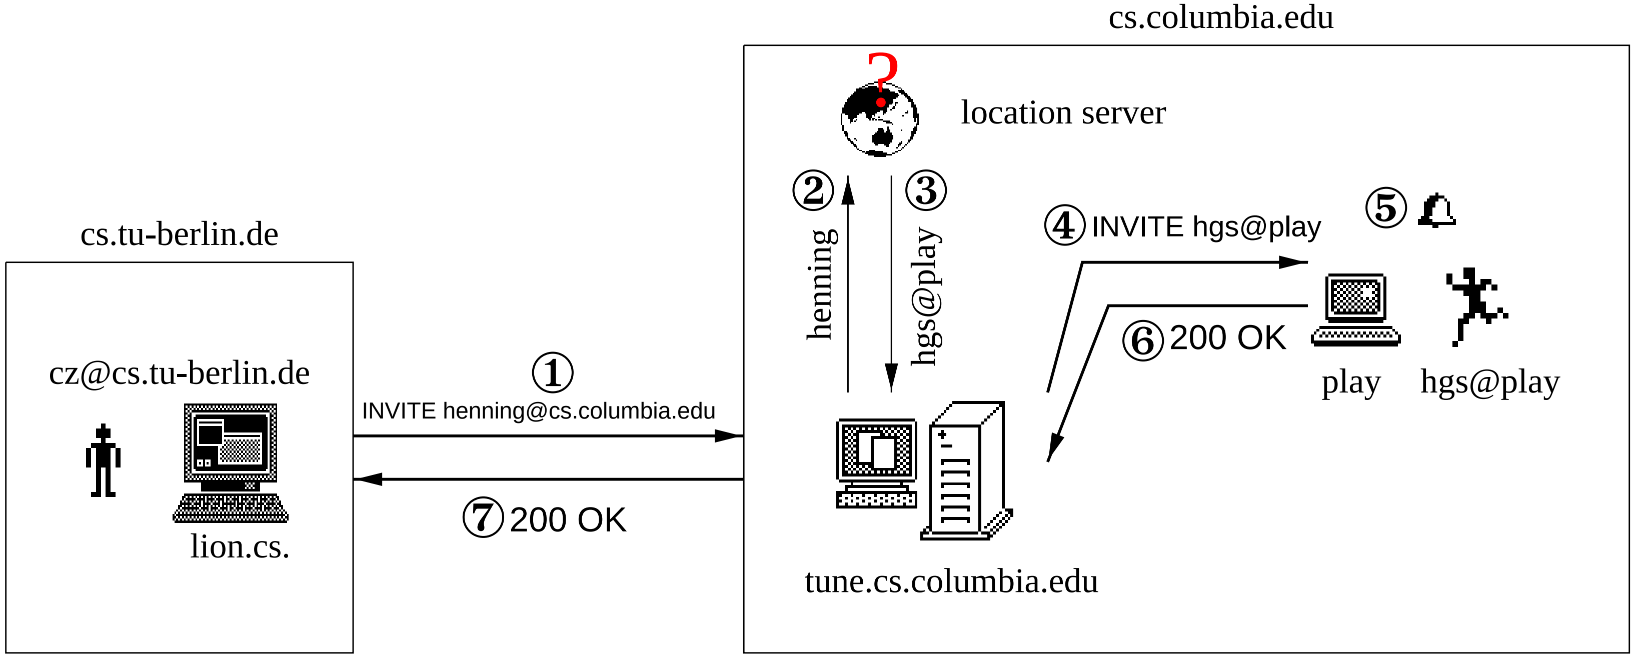

Operation in redirect mode

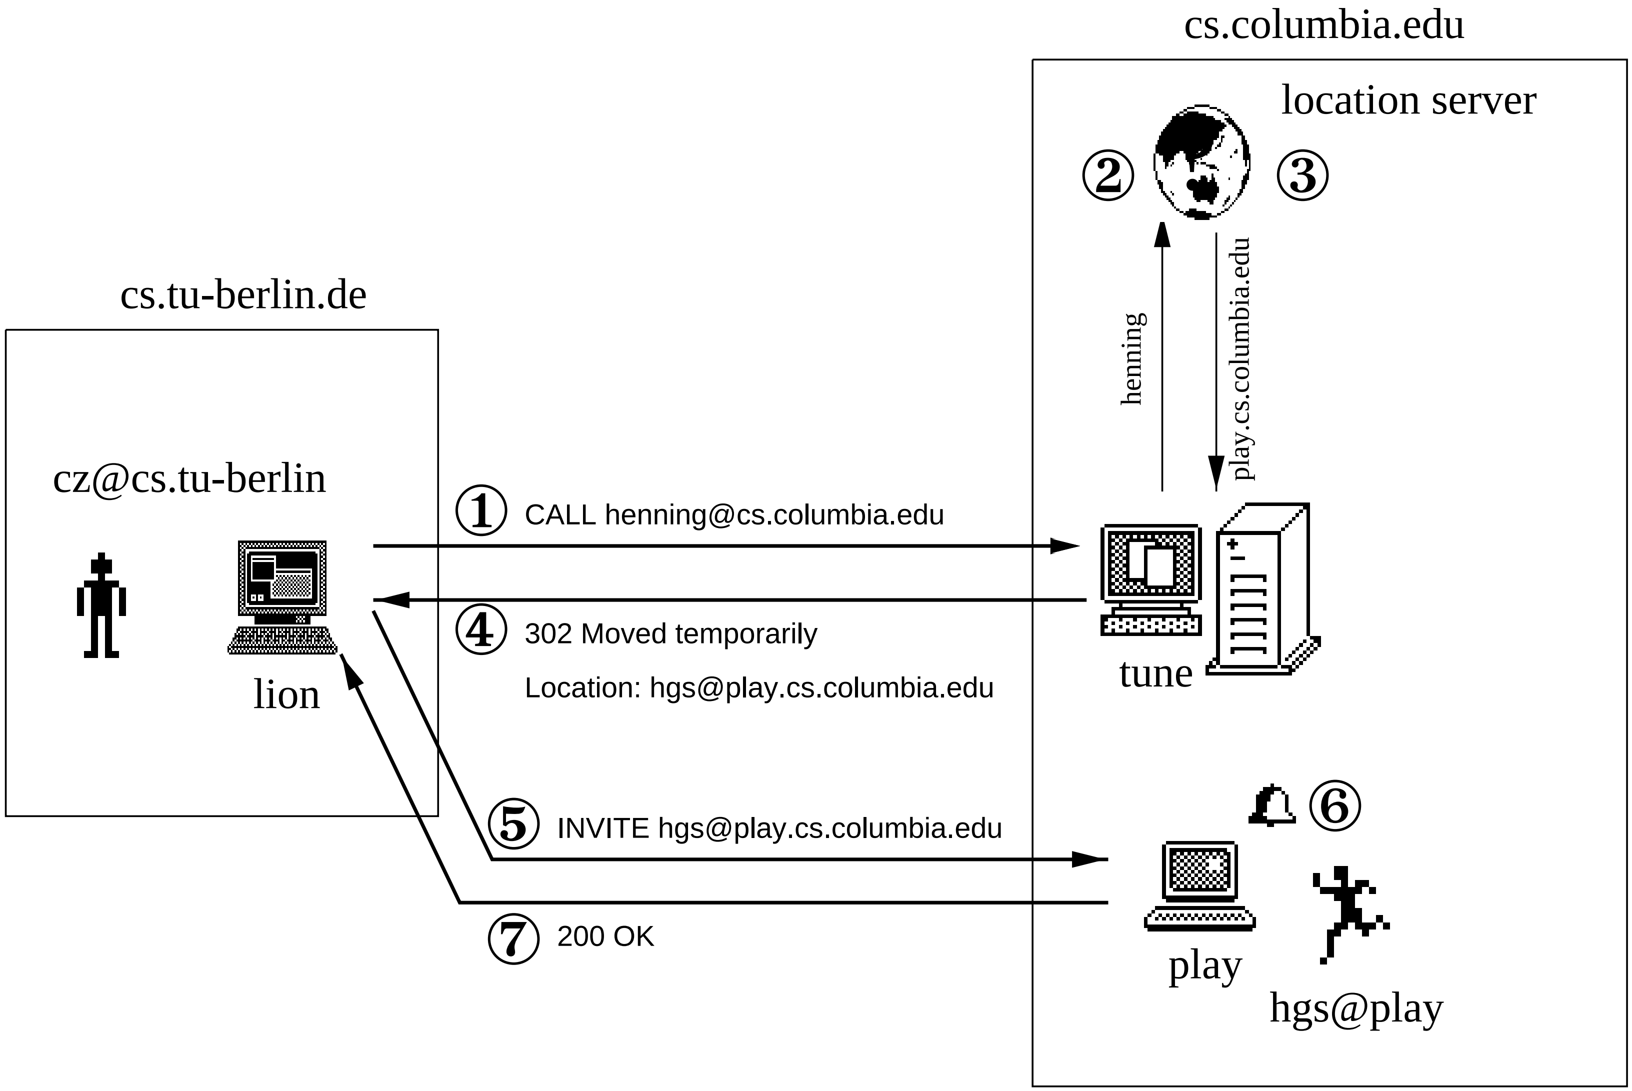

Protocol

- similar to HTTP or SMTP:
  - TCP-based (but also UDP)
  - text-based ➠ extendable, scripting languages
  - client/server
  - HTTP error codes, redirect
  - re-use authentication, payment protocol, PICS, ...
- can use HTTP server ➠ performance, stability, management, logging, ...
- several requests on same TCP connection (like HTTP/1.1) or new TCP connection for each
- disconnect through RTP BYE or timeout (soft-state)

SIP and UDP

- multicast ➠ parallel local searches
- tighter timing control on retransmission
- no connection state in proxies
- but: firewalls, TSL…
- length limitation

Protocol request

```
INVITE 199612061103.AA1528@cloud9.cs.tu-berlin.de SIP/2.0
From: Christian Zahl <cz@cs.tu-berlin.de>
To: Henning Schulzrinne <schulzrinne@cs.columbia.edu>
Via: SIP/2.0/UDP 131.215.131.131, SIP/2.0 foo.com
Content-Type: application/sdf
Content-Length: 187
Subject: New error codes
```

session description

Response

- accept or reject call, manually or automatically
- redirect temporarily or permanently

Redirection ➠ indicate several possible locations

```
SIP/2.0 100 Trying to find user

SIP/2.0 150 Ringing

SIP/2.0 302 Callee has moved temporarily
Location: jones@salt.lab3.company.com
Location: jones@pepper.lab3.company.com
Location: +1.212.939.7042
```

Session description

- currently: SDP – not extensible, not structured
- alternatives: INRIA, Quicktime with pointers (?), HyTime
- (yet another) Session Description Format (SDF):
  - live or stored media sessions
  - *not scripting*
  - structure:
    - sequence: slides
    - parallel: start at the same time
    - alternative: choose one
  - LISPish format

Implementation

- use TSL for security instead of separate signaling network
- implemented as PBX replacement
- reuse `.forward`, alias, mailcap mechanism
- part of PMM (conference bus) system
- configure media agents via WWW mailcap mechanism

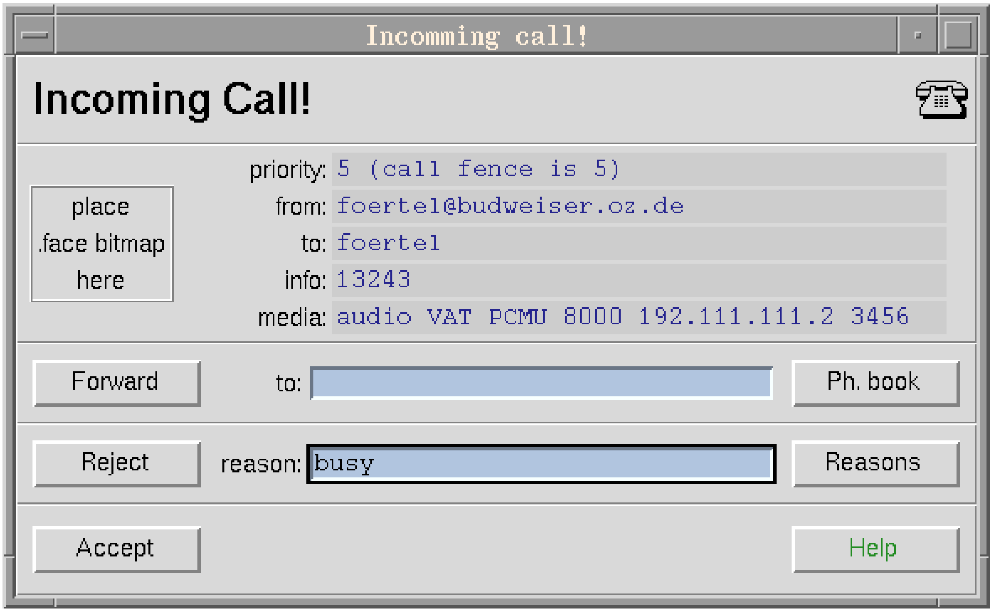

Call transfer

- multiple simultaneous call presences easy ($\leftrightarrow$ reservations)
- transfers independent of local (PBX) or Internet

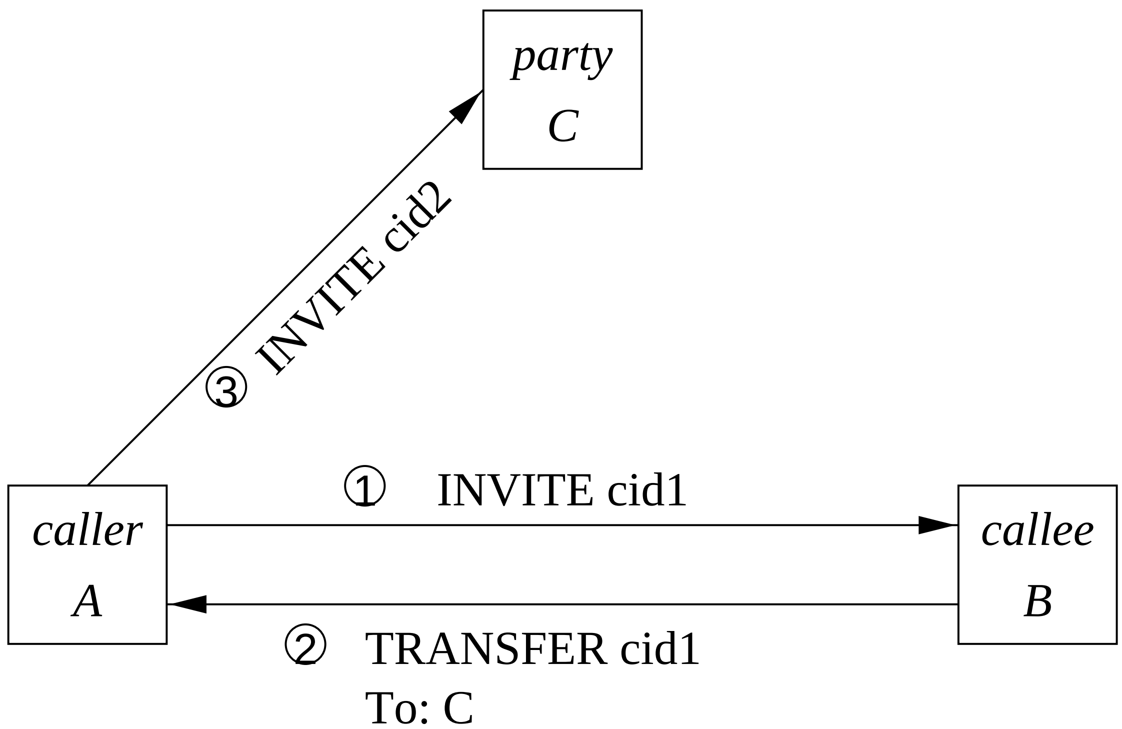

Charging and payments

- model: flat rate for best effort, per-kb minute for reserved
- RSVP: include payment authorization in RESV packet
- with RSVP, all receiving parties pay (not the same?)
- prevent kickback schemes?
- include billing authorization for particular PATH msg id?

800/888 service

- SIP has mapping capability ➠ 800/888 not needed for call distribution
- opportunity for getting paid to call somebody

900 service

- no need for phone company billing?
- callee: `402 payment required`, indicates amount
- callee reconnects with ecash token
- if $/min, keep connection open?

Trunk signaling

- PBX $\leftrightarrow$ PBX, gateway $\leftrightarrow$ gateway
- aggregate RTP streams into single payload with named slots

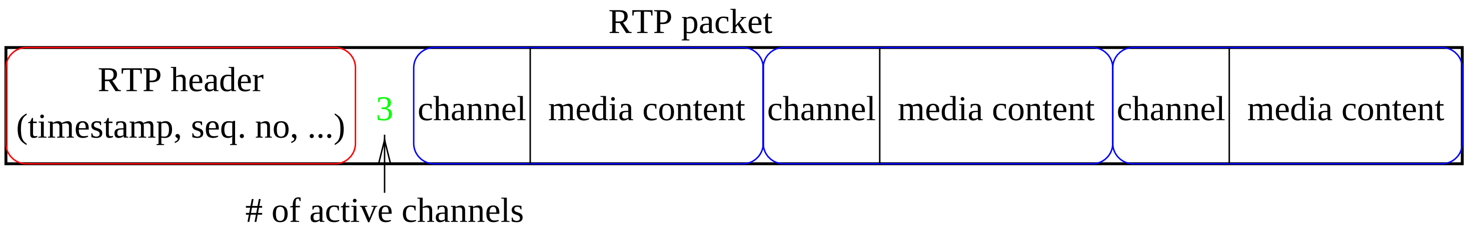

- signal mapping of slots to destinations (UCI)

Infrastructure

- push packets as close to user as possible
- network termination: packet-to-two-wire
- FTTN or HFC
- ISDN problems ➠ need power ➠ local battery?
- use (switched) Ethernet (not ATM):
  - easy local distribution (CAT3, coax)
  - variety of speed, media
  - cheap host interfaces

Future work

- trunk signaling details
- session + capability specification
- call handling language
- interaction with lower-layer mobility
- integration with RSVP?

RTSP' – Real-Time Stream Protocol

RTSP’

- control server
- motivated by RTSP (Netscape, Progressive Audio)
- similar design as SIP (TCP + UDP, HTTP, . . .)
- TCP connection $\neq$ RTSP’ association
- single session, multiple hosts
- connection/stream, connection/command, UDP
- session by identifier
- may embed data stream
- supports any session description

RTSP: media-on-demand

1. client: GET named session description from server (HTTP, RTSP)
2. client or server: SET_PARAMETER for port, blocksize, . . .
3. client: PLAY to start session or component(s)
4. server may use SESSION to update description

RTSP: play into a conference

1. member: invite server via SIP ➠ media, multicast addresses
2. member: GET named session description, indicate conference
3. …same as before …
4. issue: mapping stored contents to streams?

Summary

- price, features ➠ Internet telephony
- signaling = naming + mapping + connection
- Internet signaling ➠ improved security and call handling
- media control and signaling closely related
- simple end-to-end signaling may replace CCS#7