Internet Telephony – Towards the Integrated Services Internet

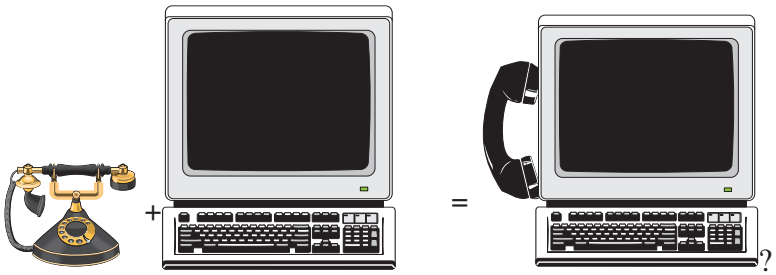

Henning Scbulzrinne
GMD Fokus, Berlin
schulzrinne@fokus.gmd.de

Workshop on Internet Telephony, Utrecht, The Netherlands; February 12, 1996

Overview

- Transport: RTP
- Adaptive applications
- End-to-end delay measurements
- Finding sessions and partners
- When can we unplug the telephone?

RTP

RTP as part of protocol stack

RTP: RFC 1889, 1890 (profile for audio and video)

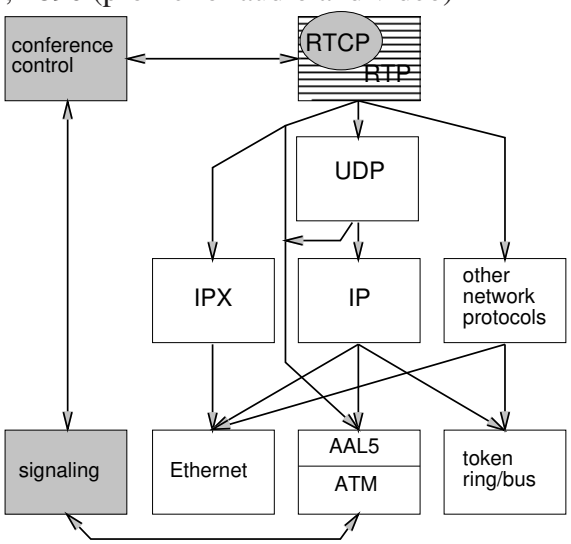

RTP functions

- 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)
- quality-of-service feedback and rate adaptation
- source identification

RTP mixers and translators

mixer:

- several media stream ➠ one new stream (new encoding)
- reduced bandwidth networks (dial-up)
- appears as new source

translator:

- operates on individual media streams
- *may* convert encoding
- protocol translation, firewall

RTP mixers and translators

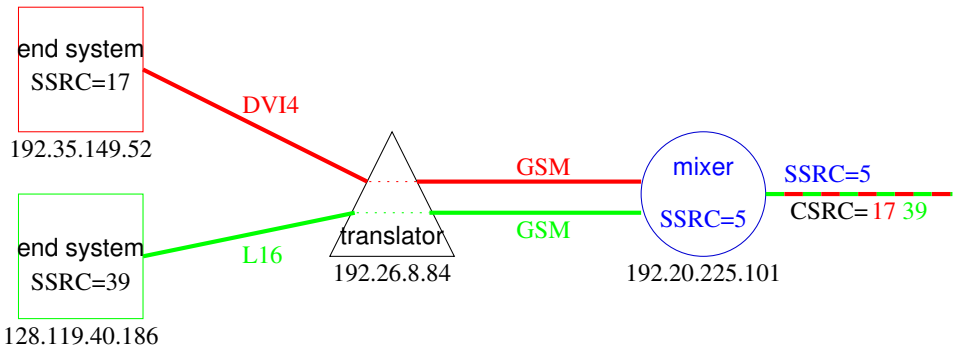

RTP packet format

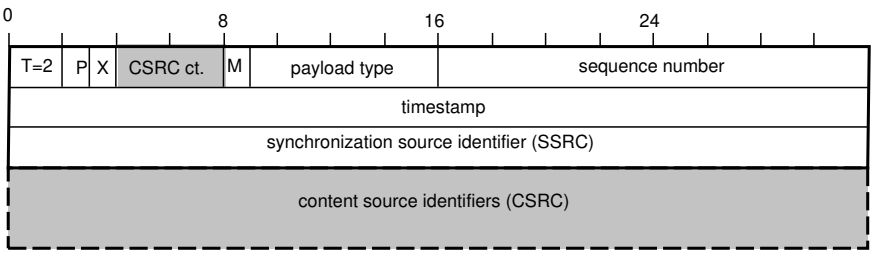

P: padding (for encryption)

M: marker bit; indicates frame, talkspurt

CC: content source count

Payload type: audio, video encoding method (static and dynamic)

SSRC: synchronization source – random 32-bit identifier

CSRC: list of contributing sources (mixer)

RTP control protocol – algorithm

Goals:

- estimate current number of participants – dynamic
- participant information ➠ talker indication
- quality-of-service feedback ➠ adjust sender rate
- side effect: connectivity indication
- scale to $O(1000)$ participants, small fraction of data bandwidth

➠ randomized response with rate $\downarrow$ as members $\uparrow$

- limited by tolerable age of status
- gives active senders more bandwidth

RTP control protocol – types

stackable packets, similar to data packets

sender report (SR): bytes send ➠ estimate rate;
timestamp ➠ synchronization

reception reports (RR): number of packets sent and expected ➠ loss,
interarrival jitter; round-trip delay

source description (SDES): name, email, location, …

explicit leave (BYE): in addition to time-out

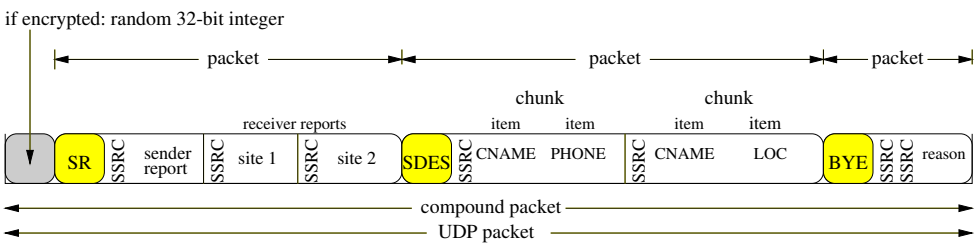

Adaptive applications

Adaptive applications

Multimedia applications can adjust their data rates:
Audio: encoding parameters (MPEG L3), encoding, sampling rate, mono/stereo

encoding	sampling rate	bit rate
LPC	8,000	5,600
GSM	8,000	13,200
DVI4	8,000	32,000
μ-law	8,000	64,000
DVI4	16,000	64,000
a range of DVI4 and MPEG L3		
L16 stereo	44,100	1,411,200

Adaptive applications

Video: frame rate, quantization, image resolution, encoding

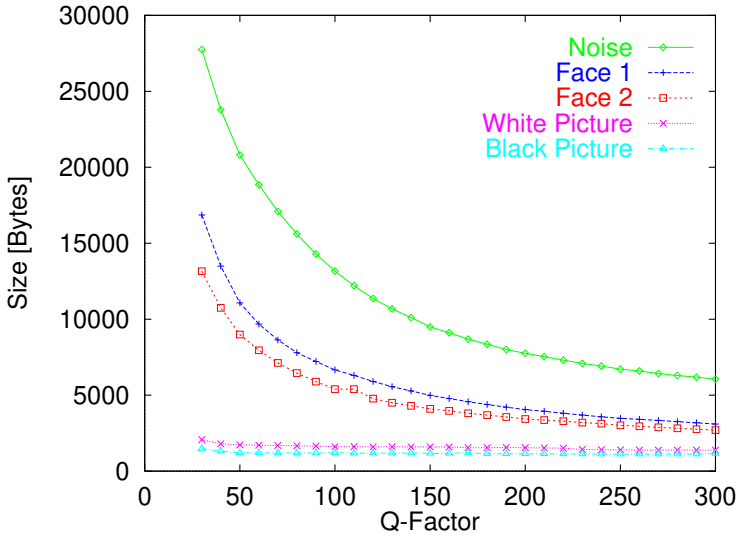

Application control

- networks with QoS guarantees:
  - QoS negotiation at call set-up time, network guarantees this quality
  - long call durations ➠ network load may change significantly
  - "wrong" initial allocation ➠ many rejected calls or low quality
- networks without QoS guarantees (= current Internet) or shared reserved link:
  - adapt application to available bandwidth
  - share bandwidth fairly with TCP?
  - lowest common demoninator ➠ combine with mixers, translators

End-to-End feedback control mechanism

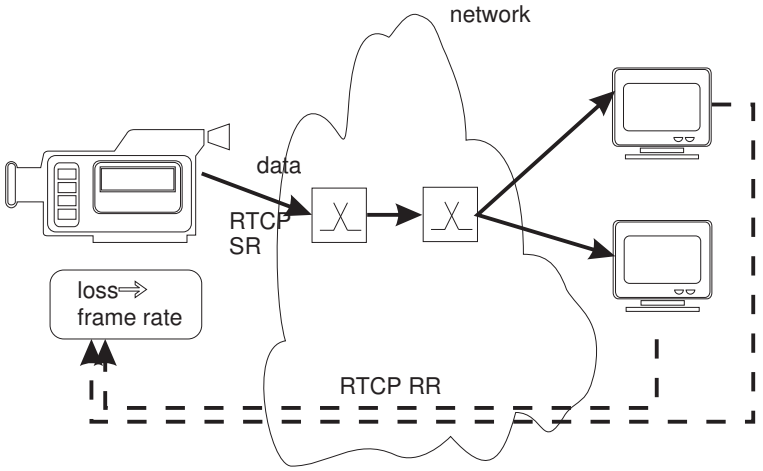

video application bandwidth is based on network feedback:
low losses ➠ slow bandwidth increase ➠ higher framerate
high losses ➠ bandwidth decrease ➠ lower framerate

Network state estimation and bandwidth adjustment

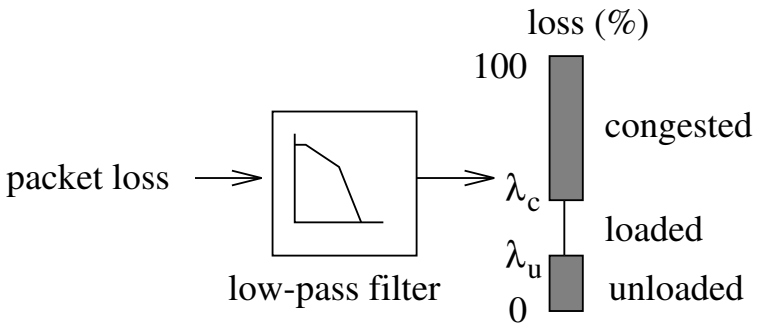

- loss information is filtered: $\lambda \leftarrow (1-\alpha)\lambda + \alpha b,\ \alpha \leq 1$
- linear regulator with deadzone
- multiplicative decrease if network is congested:
  $b_a \leftarrow \max\{b_a * \mu, b_{min}\},\ \mu < 1$
- additive increase if network is unloaded: $b_a \leftarrow \min\{b_a + \nu, b_{max}\}$

Multicast scalability

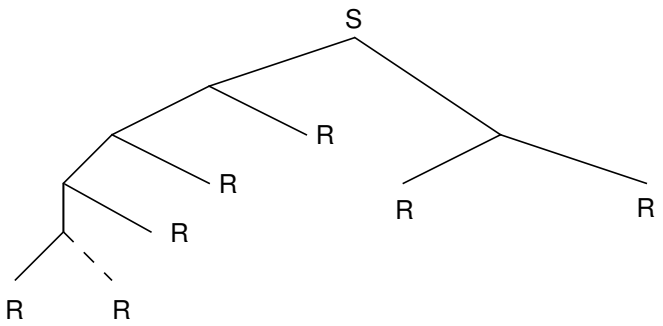

algorithm 1: adjust according to the worst-positioned receiver

algorithm 2: allow a fraction of the receivers to be congested

other solutions: video gateways, layered encodings

Internet scenario

- Measurements made on the 2 Mbit/s X.25 link between GMD Fokus and TU Berlin (5 hops distance)
- deadzone between 5% and 10%

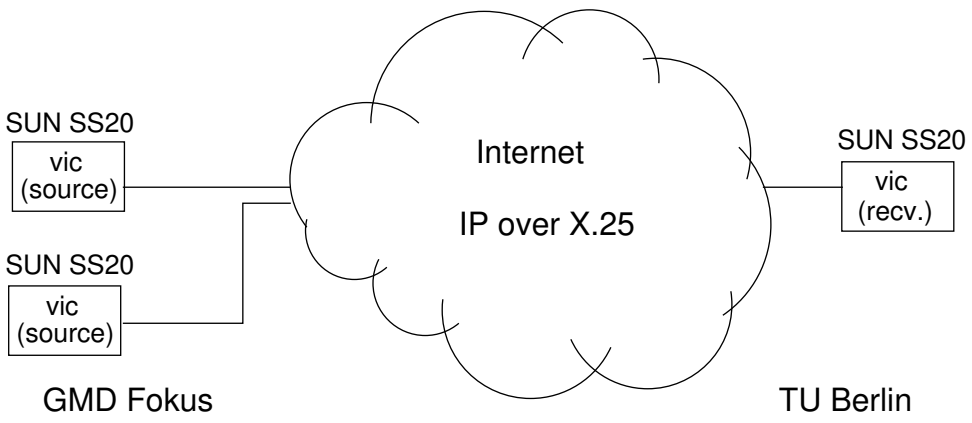

Internet measurement

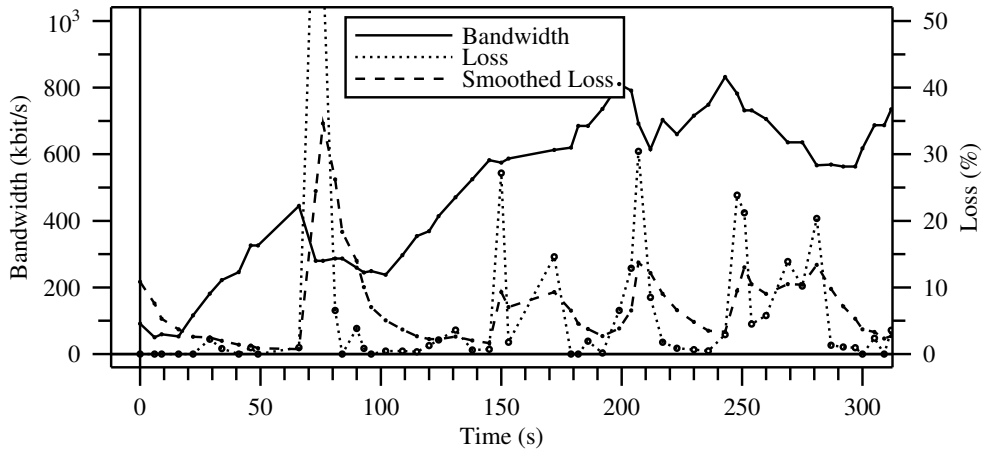

ATM scenario

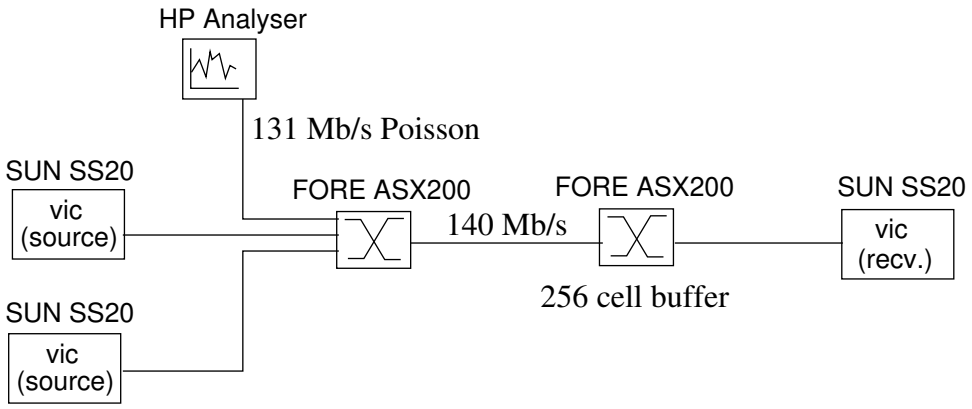

ATM measurements

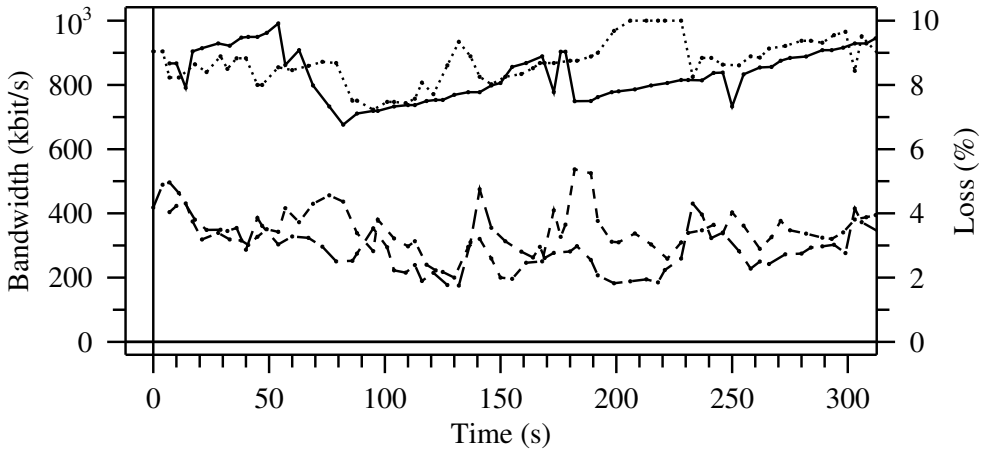

Observations

- jitter as loss predictor does not work well
- coexistence with controlled data applications
- use in ABR-like services?

Measuring end-to-end delays

Estimation of end-to-end delay

- clocks are not synchronized
- unknown delays in operating system, hardware codecs
- ➠ network delay measurement using cross correlation of audio signals

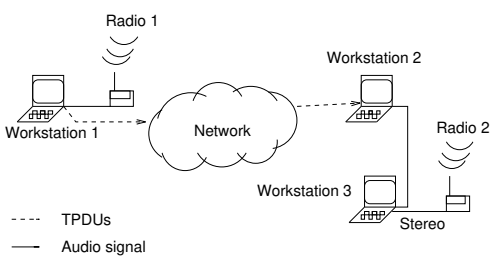

Estimation of end-end delay

Cross correlation of a recorded stereo signal

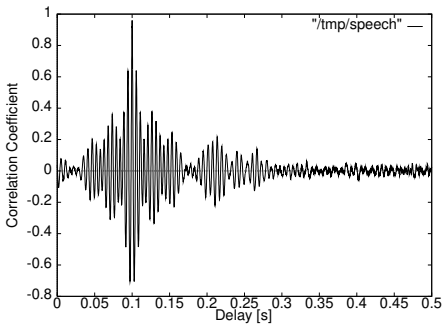

Finding sessions and partners

Finding sessions and partners

Need to convey

- multicast or unicast address and port
- preferred/possible encodings, dynamic payload mappings
- encryption keys
- conference policy (floor control, …)

Member discovery: participant searches for conference:

- (global) multicast directory: SD; `sdr` ➠ periodic events
- WWW pages
- netnews, MIME (active) mail, …

Invitation: real-time, directed to a particular user

Also need multicast address allocation (done by SD)

Internet conference model

- hard to maintain consistent view of membership without central authority
- global view usually not necessary
- “admission control” through encryption
- member ∈ audio/video/…session ∈ conference
- sessions ≡ multicast group and port
- member ≡ RTP CNAME (user@host) ➠ identification across sessions
- conference ≡ name?

The conference invitation protocol

- reach one or more subscribers by their normal email address (rather than current host name)
- new and on-going conferences
- allow for manual and automatic forwarding
- personal mobility, complements data link/IP mobility
- reach first (load distribution, e.g., `support@company.com`) or reach all (department conference)
- possibly change permissible encodings during conference

fokus

Conference invitation protocol: reusing email infrastructure

- friendly naming: `john.doe`, `j.doe`, …
- DNS MX records for organization addresses, reliability
- `.forward` mechanism
- email aliases (“life-time phone number”: `foo@ieee.org`)
- mailing lists
- email directories (X.500, whois++, …)
- but: real-time, with fallback to email delivery

fokus

Resolving an email address

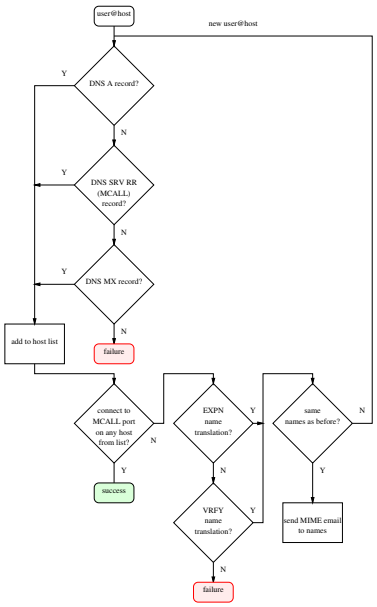

Operation in proxy mode

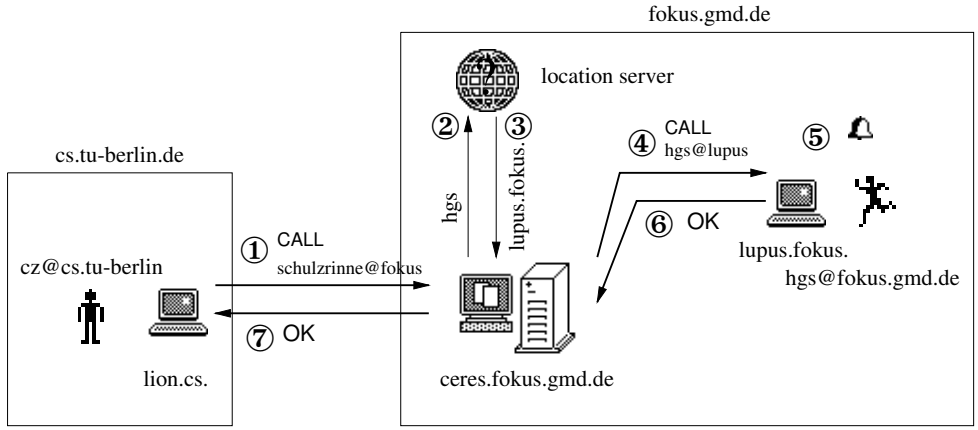

Operation in redirect mode

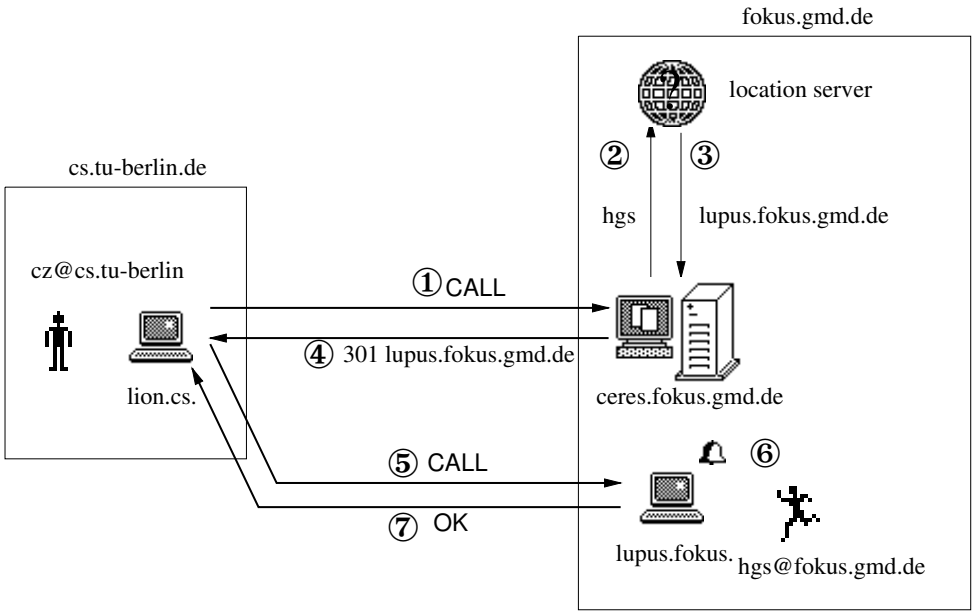

Protocol

- similar to HTTP or SMTP: TCP-based, ASCII, client/server
- connection may stay open for change requests
- HTTP error codes, redirect

Protocol request

```
CALL hgs@lupus.fokus.gmd.de MUCS/1.0
User-Agent: isc/1.0
From: Christian Zahl <cz@cs.tu-berlin.de>
To: Henning Schulzrinne <schulzrinne@fokus.gmd.de>
Call-Id: 9510021900.AA07734@lion.cs.tu-berlin.de
Referer: ceres.fokus.gmd.de
Expires: Mon, 02 Oct 1995 18:44:11 GMT
Required: fc99cb08 audio/pcmu; port=3456; transport=RTP;
  rate=16000; channels=1; pt=97; net=224.2.0.1; ttl=128,
  audio/gsm; port=3456; transport=RTP; rate=8000; channels=1,
  audio/lpc; port=3456; transport=RTP; rate=8000; channels=1
Required: 83ae5290 video/h261;port=4134;transport=RTP;rate=128,
  video/nv;port=4136;transport=RTP;rate=250,
  application/x-wb;port=1236
Accept: 56af7e9c application/editor;port=3500
Phone: +1 413 555 1212
Email: Christian Zahl <cz@cs.tu-berlin.de>
Location: Technical University Berlin; tz=MET;
  loc=52 32 00 N 13 25 00 E
Reimburse: ecash
Priority: urgent
Reach: first
Key: C7 48 90 F4 27 7B A1 CF
Subject: New MUCS error codes
```


Response

- accept or reject call
- redirect temporarily or permanently
- indicate several possible locations

```
MUCS/1.0 302 Callee has moved temporarily
Location: jones@salt.lab3.company.com
Location: jones@pepper.lab3.company.com
```

Protocol extensions

- freephone (1-800) facilities (ecash, reimburse provider, …)
- interoperation with ISDN, POTS, H.320
- transition from two-party to multicast?
- only usable with permanently connected system (intranet, CATV Internet appliance)

When can we unplug the telephone?

Audio issues

- X-window-like libraries usually just make life difficult
- need: media agent, for both remote and local audio (➠ *one* playout delay adaptation)
- audio interfaces need: low latency, independent input/output, single copy
- CPU's with efficient support for DSP operations (multiply-accumulate)
- timestamped audio samples to allow correlation, echo cancellation
- system libraries: coding, AGC, rate conversion, silence detection, …
- Internet terminals with built-in packet audio?
- need to avoid push-to-talk feel ➠ better silence detection

Network support

- ATM
  - connection management promises pain ➠ receiver joins help, but still need to locate senders
  - but: may be a lot simpler than IETF sparse-mode multicast for small groups
- charging for congestion-induced, variable quality?

Pros

- integration with computer applications (easier than CTI)
- price advantage for hardware and service (but: may be temporary)
- support of large groups (↔ MCUs)
- "advanced" services are much easier than on SS7:
  - directory services
  - speaker indication
  - source selection at receiver
- flexible mix of audio, video, shared applications

Cons

- quality currently unacceptable/unpredictable for "commercial" use
- hands-free speaking in infancy
- high overhead (per 160 byte packet):

	IPv4	IPv6
RTP	12	12
UDP	8	8
IP	20	40
sum	40	60
LLC/SNAP	(8)	(8)
AAL5	8	8
ATM	25	25
sum	81	101
utilization (%)	66	61

but: save half with silence suppression

- use too complicated (TTL, multicast/unicast, encodings, …)
- phone calls need to be pre-arranged (➠ intranets, Internet on CATV)
- network/application reliability (GPF when calling the fire department?)
- Internet handy? Internet pay phone?

Perspectives

- toy for computer hobbyists, long-distance lovers, CB crowd
- distribution of conferences, radio ➠ delay does not matter, no ready alternative
- “second phone” for international calls
- long-term replacement for telephone – at least for business use?