

INF3190 – Data Communication Multimedia Protocols

Carsten Griwodz

Email: griff@ifi.uio.no

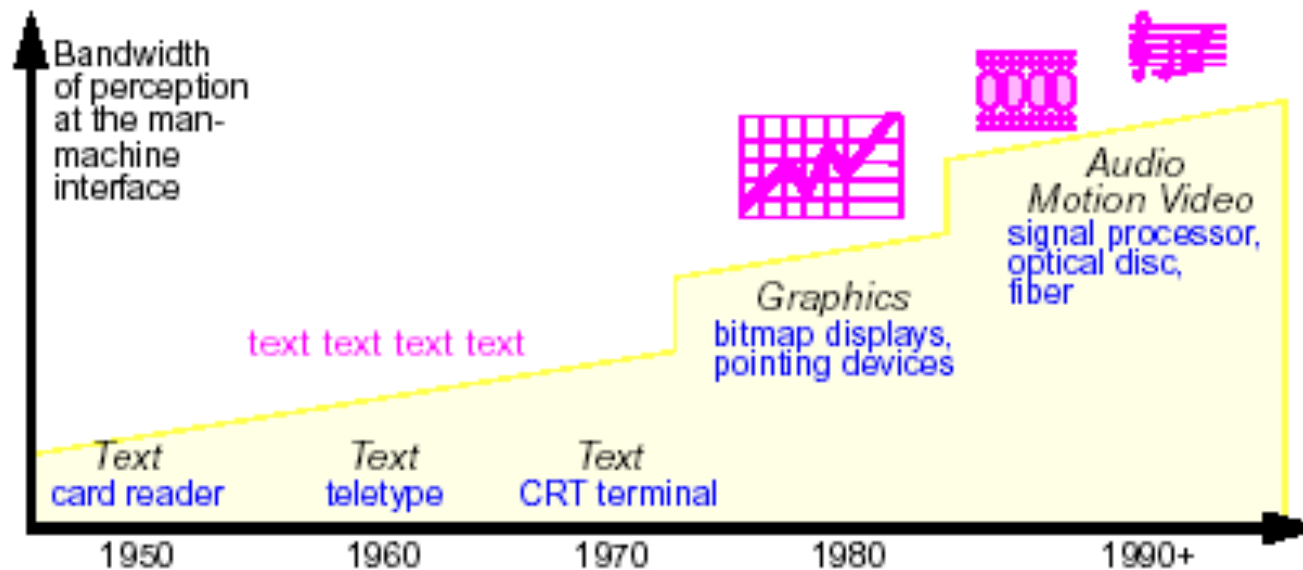


Media

Medium: "Thing in the middle"

- here: means to distribute and present information

Media affect human computer interaction



The mantra of multimedia users

- Speaking is faster than writing
- Listening is easier than reading
- Showing is easier than describing

Dependence of Media

- Time-independent media
 - Text
 - Graphics
 - *Discrete* media
 - Time-dependent media
 - Audio
 - Video
 - Animation
 - Multiplayer games
 - *Continuous* media
 - Interdependant media
 - *Multi*media
- "Continuous" refers to the user's impression of the data, not necessarily to its representation
 - Combined video and audio is multimedia - relations must be specified

Continuous Media

Fundamental characteristics

- Typically **delay sensitive**
- Often **loss tolerant**: infrequent losses cause minor glitches that can be concealed
- Antithesis of discrete media (programs, banking info, etc.), which are loss intolerant but delay tolerant

Classes of MM applications

- Streaming stored audio and video
- Streaming live audio and video
- Interactive real-time audio and video
- Interactive real-time event-driven applications

Multimedia in networks

Streaming stored MM

- Clients request audio/video files from servers and pipeline reception over the network and display
- Interactive: user can control operation (pause, resume, fast forward, rewind, etc.)
- Delay: from client request until display start can be 1 to 10 seconds

Unidirectional Real-Time

- similar to existing TV and radio stations, but delivery over the Internet
- Non-interactive, just listen/view

Interactive Real-Time

Phone or video conference


- More stringent delay requirement than Streaming & Unidirectional because of real-time nature
- Audio: < 150 msec good, < 400 msec acceptable
- Video: < 150 msec acceptable
[Note: higher delays are feasible, but usage patterns change *(!)*]

Games *(but also high-speed trading)*

- Role playing games: < 500 msec
- First person shooter (FPS) games: < 100 msec *(may be too high)*
- Cloud gaming FPS: < 40 msec *(estimated)*

Slides by Prof. Ramesh Sitaraman, UMass, Amherst (shown with permission)
"Video Stream Quality Impacts Viewer Behavior: Inferring Causality using Quasi-Experimental Designs", S. S. Krishnan and R. Sitaraman, ACM Internet Measurement Conference (IMC), Boston, MA, Nov 2012

Viewers with better connectivity have less patience for startup delay and abandon sooner.



Quality of service - QoS

A term that is used in all kinds of contexts.

Be careful what it means when you hear it.

In this lecture: 3 *classical*

parameters of **network QoS**:

- end-to-end delay
- packet loss
- jitter

end-to-end delay

- transmission time
- Σ propagation time on link l
sum of propagation times over all links l
- Σ queueing time on router r
sum of queueing times at all routers' queues r

packet loss

- probability of a packet to get lost
- $1 - (\prod (P(\text{queue at } r \text{ not full})))$
1 - product of probabilities for all r that queue at r is not full

jitter

- variance of end-to-end delay
- estimated for several packets
- reasons
 - link layer retransmissions
 - queue length variation

Multimedia Networking

Internet without network QoS support

- Internet applications must cope with networking problems
 - Application itself or middleware
 - "Cope with" means either "*adapt to*" or "*don't care about*"
 - "Adapt to" must deal with TCP-like service variations
 - "Don't care about" approach is considered "unfair"
 - "Don't care about" approach cannot work with TCP

Internet with network QoS support

- Application must specify their needs
- Internet infrastructure must change – negotiation of QoS parameters
- Routers need more features
 - Keep QoS-related information
 - Identify packets as QoS-worthy or not
 - Treat packets differently keep routing consistent

- approach seemed "dead" for many years
- revival with recent Software Defined Networking (SDN) idea
- not yet mainstream again

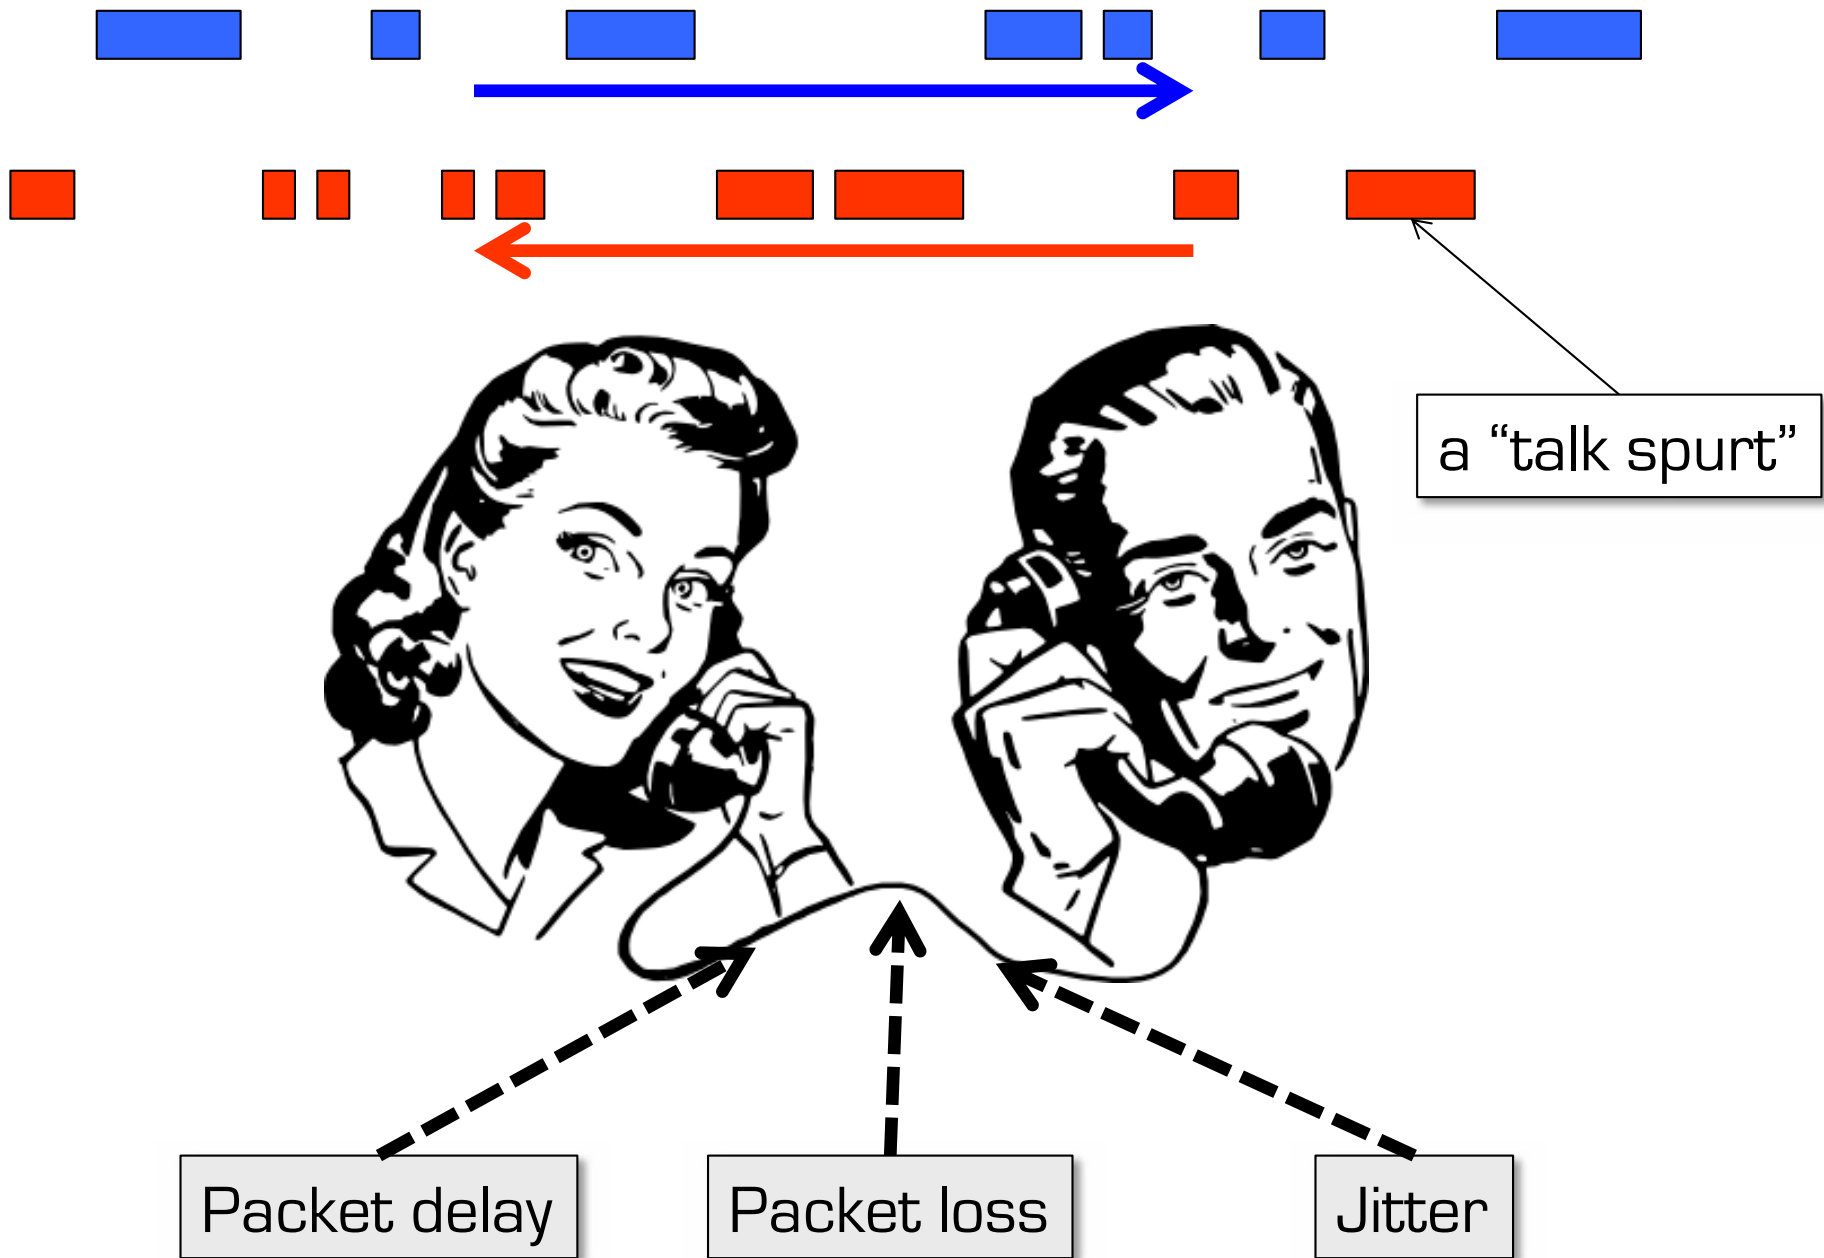


Non-QoS Multimedia Networking

Basics



Streaming over best-effort networks: audio conferencing



Streaming over best-effort networks: audio conferencing

end-to-end delay

- end-to-end delay can seriously hinder interactivity
- smaller is always better? not true for cooperative music making!

packet loss

- UDP segment is encapsulated in IP datagram
- datagram may overflow a router queue
- TCP can eliminate loss, but
 - retransmissions add delay
 - TCP congestion control limits transmission rate
- redundant packets can help

delay jitter


- consider two consecutive packets in talk spurt
- initial spacing is 20 msec, but spacing at receiver can be more or less than 20 msec

removing jitter

- sequence numbers
- timestamps
- delaying playout

Delay compensation

All techniques rely on *Prediction*



no delay compensation
in this example

Teleconferencing

- no known technique: cannot predict what people will say

For on-demand

- usually content is consumed linearly, prefetching is easy, limited only by resources and legal constraints

For event-based multimedia

- predict future movement
- perform audiovisual rendering based on prediction
- compensate for errors in next prediction
- used in computer games and other distributed simulations, head- and gesture tracking, mouse or joystick inputs

Jitter compensation

Receiver attempts to playout each chunk at exactly q msec after the chunk is generated

- If chunk is time stamped t , receiver plays out chunk at $t+q$
- If chunk arrives after time $t+q$, receiver discards it

Sequence numbers not necessary

Strategy allows for lost packets

Tradeoff for q :

- large q : less packet drop/loss (better audio quality)
- small q : better interactive experience

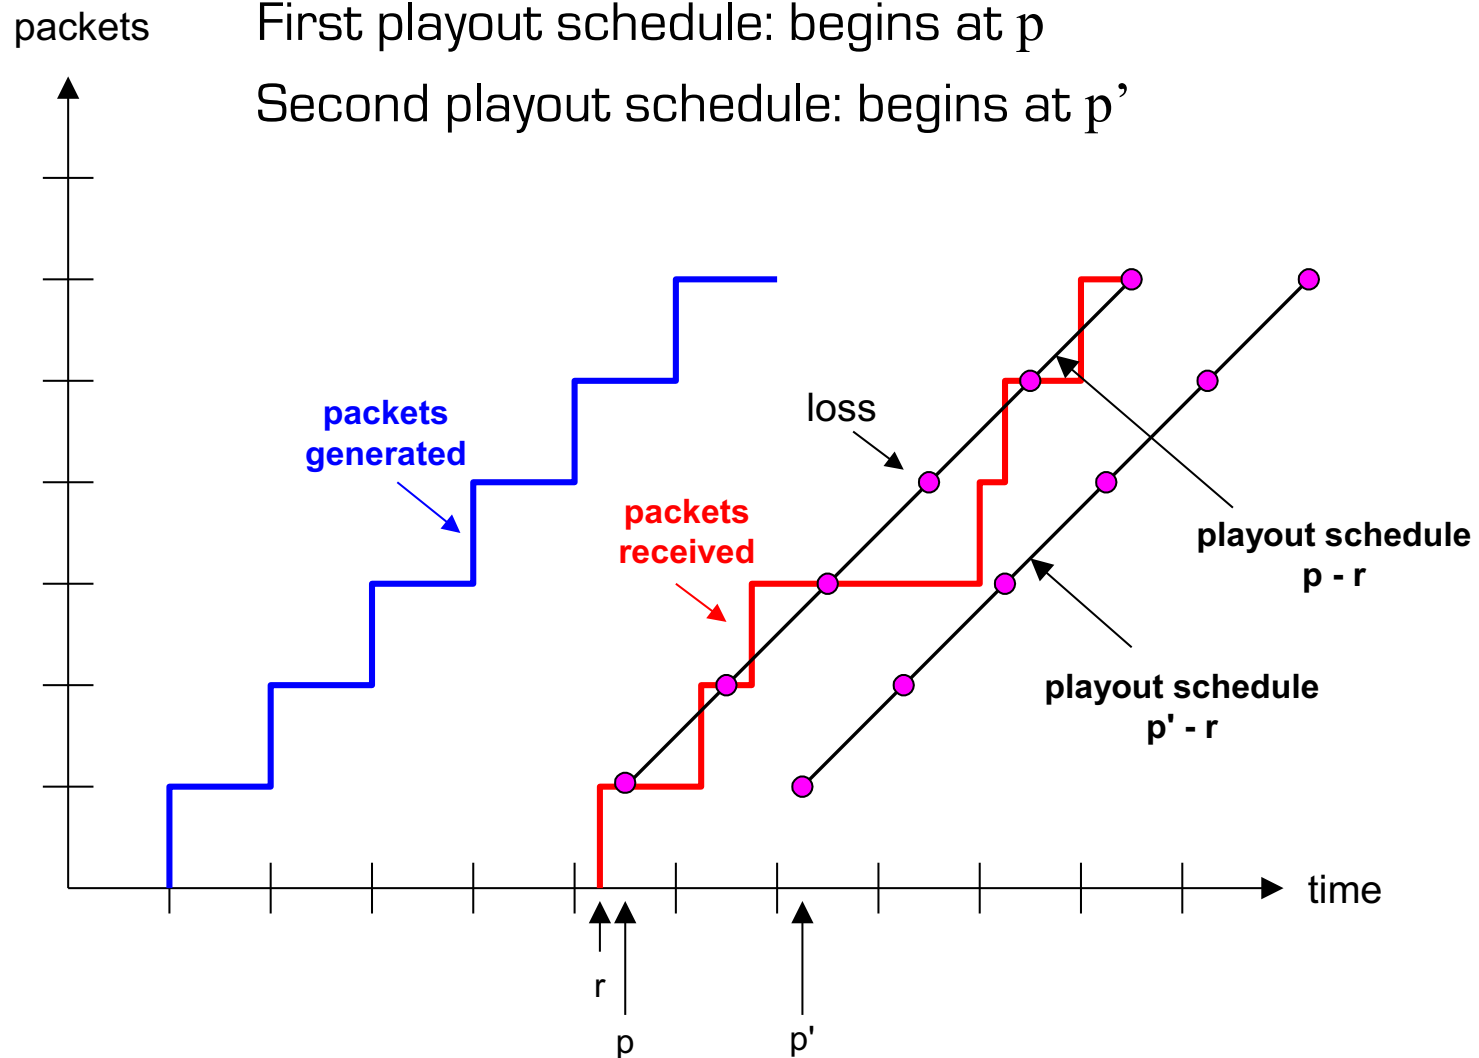
Jitter compensation

Sender generates packets every 20 msec during talk spurt

First packet received at time r

First playout schedule: begins at p

Second playout schedule: begins at p'



Jitter compensation: Adaptive playout delay

Estimate network delay and adjust playout delay at the beginning of each talk spurt

Silent periods are compressed and elongated as needed

Chunks *still* played out every 20 msec during talk spurt

t_i = timestamp of the i th packet

r_i = the time packet i is received by receiver

p_i = the time packet i is played at receiver

$r_i - t_i$ = network delay for i th packet

d_i = estimate of average network delay after receiving i th packet

Dynamic estimate of average delay at receiver:

$$d_i = (1 - u)d_{i-1} + u(r_i - t_i)$$

where u is a fixed constant [e.g., $u = .01$]

Jitter compensation: Adaptive playout delay

Also useful to estimate the average deviation of the delay, v_j :

$$v_i = (1 - u)v_{i-1} + u |r_i - t_i - d_i|$$

Deviation: How strongly does the queue length change?

The estimates d_i and v_i are calculated for every received packet, although they are only used at the beginning of a talk spurt

For first packet in talk spurt, playout time is:

$$p_i = t_i + d_i + Kv_i$$

application chooses the safety margin Kv_i

where K is a positive constant

Playout delay is $q_i = p_i - t_i = d_i + Kv_i$

for this and **all other** packets in **this** talk spurt

Jitter compensation: Adaptive playout delay

How to determine whether a packet is the first in a talkspurt?

- If there were never loss, receiver could simply look at the successive time stamps
 - Difference of successive stamps > 20 msec, talk spurt begins
- But because loss is possible, receiver must look at both time stamps and sequence numbers
 - Difference of successive stamps > 20 msec and sequence numbers without gaps, talk spurt begins



Loss compensation

Basic assumption

- we have very little time to loose in audio conferencing
- every packet carries dozens of samples
- adding several packets delay for complex schemes is not viable

forward error correction (FEC): simple scheme

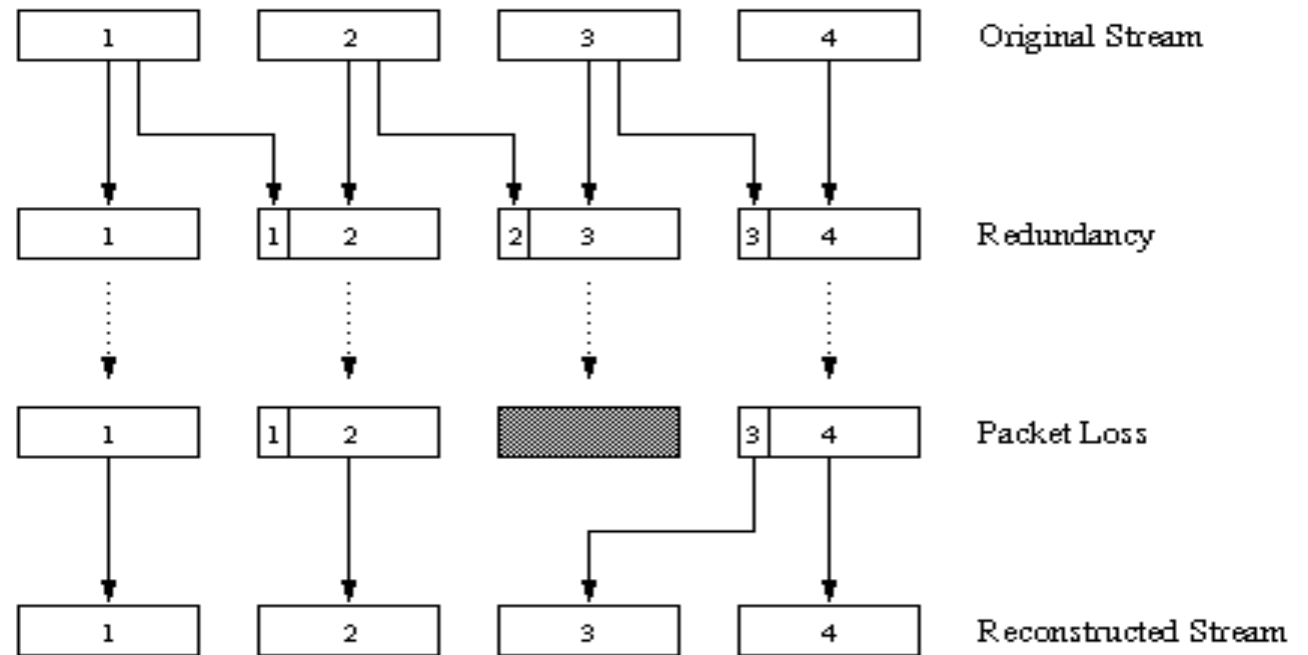
- for every group of n chunks create a redundant chunk by exclusive OR-ing the n original chunks
- send out $n+1$ chunks, increasing the bandwidth by factor $1/n$.
- can reconstruct the original n chunks if there is at most one lost chunk from the $n+1$ chunks
- Playout of first packet has to wait for arrival of $(n+1)^{\text{st}}$ packet
- Playout delay needs to be fixed to the time to receive all $n+1$ packets
- Tradeoff:
 - increase n , less bandwidth waste
 - increase n , longer playout delay
 - increase n , higher probability that 2 or more chunks will be lost



Loss compensation

2nd FEC scheme

- “piggyback lower quality stream”
- send lower resolution audio stream as the redundant information
- for example, nominal stream PCM at 64 kbps and redundant stream GSM at 13 kbps.

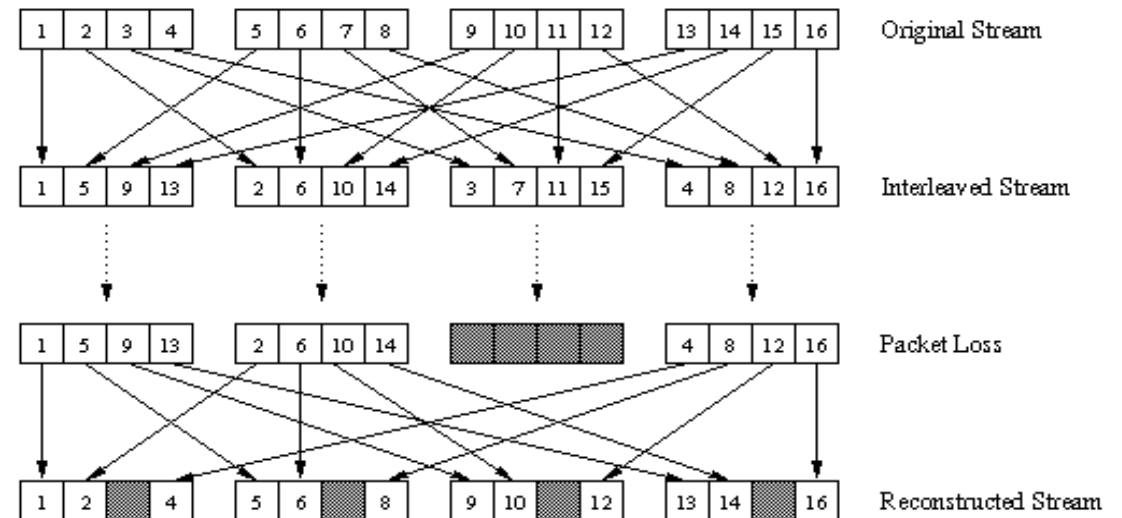


- Sender creates packet by taking the n th chunk from nominal stream and appending to it the $(n-1)$ st chunk from redundant stream.
- Whenever there is non-consecutive loss, the receiver can conceal the loss.
- Only two packets need to be received before playback
- Can also append $(n-1)$ st and $(n-2)$ nd low-bit rate chunk

Loss compensation

Interleaving

- chunks are broken up into smaller units
- for example, 4 5 msec units per chunk
- interleave the chunks as shown in diagram
- packet now contains small units from different chunks



- Reassemble chunks at receiver
- if one packet is lost, still have most of every chunk

Loss compensation

Receiver-based repair of damaged audio streams

- produce a replacement for a lost packet that is similar to the original
- can give good performance for low loss rates and small packets (4-40 msec)
- simplest: repetition
- more complicated: interpolation

Making the best of best effort

Mitigating the impact of “best-effort” in the Internet

Use UDP to avoid TCP and its slow-start phase

... but TCP is changing (removing slow start, larger initial windows)

Buffer content at client and control playback to remedy jitter

... but not for event-based multimedia (games)

We can timestamp packets, so that receiver knows when the packets should be played back

... but applications may ignore this and look for timestamps in content (typical in MPEG video)

Adapt compression level to available bandwidth

... but not for event-based multimedia

We can send redundant packets to mitigate the effects of packet loss

... but retransmission and TCP may be more efficient

Non-QoS Multimedia Networking

Application Layer Framing &
Integrated Layer Processing



Multimedia Content Processing

- Problem: optimize transport of multimedia content

- It is application-dependent and specific
 - Application-layer processing has high overhead
 - Application processes data as it arrives from the network

- Impact of lost and mis-ordered data
 - either:** Transport layer tries to recover from error
 - Prevents delivery of data to application
 - Prevents immediate processing as data arrives
 - Application must stop processing
 - or:** Transport layer ignores error
 - Application experiences processing failures
 - Application must stop processing

Application Level Framing

[Clark/Tennenhouse 1990]

Give application more control

- Application understands meaning of data
- Application should have the option of dealing with a lost data
 - Reconstitute the lost data (recompute/buffer by applications)
 - Ignore the lost data

Application level framing

- Application breaks the data into suitable aggregates
 - Application Data Units (ADUs)
- Lower layers preserve the ADU frame boundaries
- ADU takes place of packet as the unit of manipulation

ALF: Application Data Units

ADUs become the unit of error recovery

- Should be upper bounded
 - loss of large ADUs is more difficult to fix
- Lower bounded
 - application semantics define smallest sensible unit
 - small ADUs mean larger protocol overhead
- Segmentation/reassembly
 - try to avoid
 - multi-TPDU ADU is wasted when one packet is loss

ADU “name”


- Sender computes a name for each ADU (e.g. sequence number)
- Receiver uses name to understand its place in the sequence of ADUs
- Receiver can process ADUs out of order

- ADUs are an old concept
- The term “*ADU*” is mostly used for RTP specifications
- The MPEG, H.264, H.265 term for ADU is **NAL unit** (network abstraction layer unit)

Integrated Layer Processing


Layered engineering is not fundamental

- Assignment of functions to layers in OSI is not following fundamental principles
- Specific application may work better with different layering of functions or no layering at all
- Sequential processing through each layer
 - Not an efficient engineering
 - Processing all functions at once saves computing power



Integrated Layer Processing

- Vertical integration
- Performing all the manipulation steps in one or two integrated processing loops, instead of serially



Integrated Layer Processing

- Ordering constraint
 - Data manipulation can only be done after specific control steps
 - Data manipulation can only be done once the data unit is in order
 - Layered multiplexing (extract the data before it can be demultiplexed)
- Minimize inter-layer ordering constraints imposed on implementors
 - Implementors know best which data must be ordered
- Drawback: complex design due to fully customized implementation

Non-QoS Multimedia Networking

RTP – Real-Time Transfer Protocol

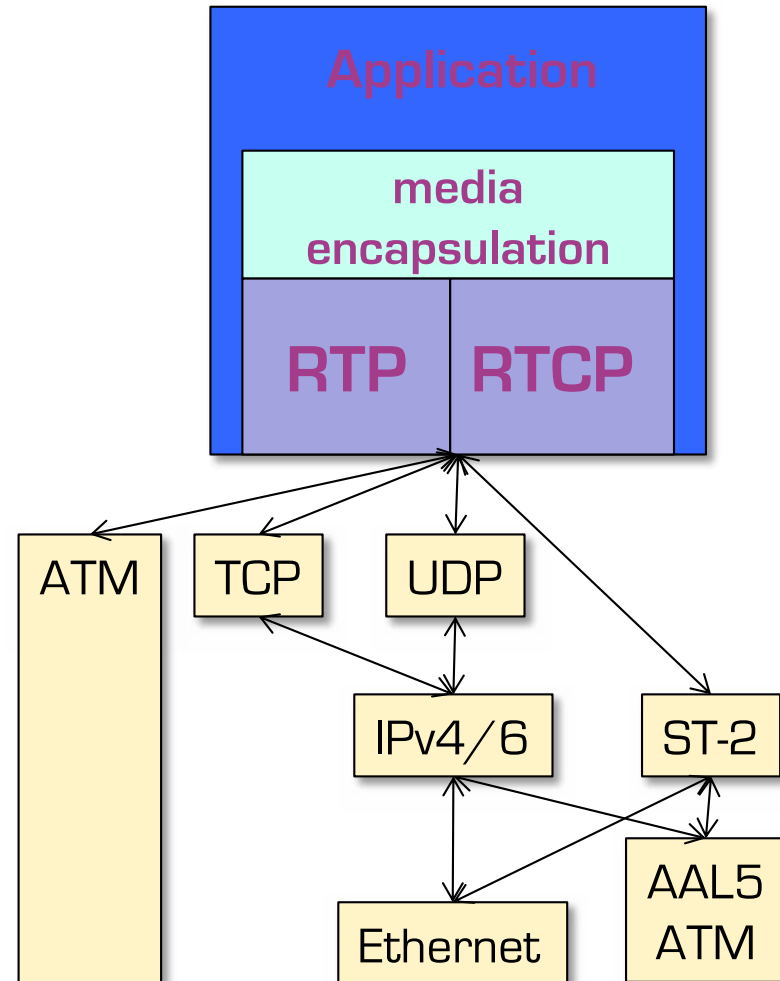


Real-time Transport Protocol (RTP)

- Real-time Transport Protocol (RTP)
 - RFC 3550 (replaces RFC 1889)
 - Designed for requirements of real-time data transport
 - **NOT** real-time
 - **NOT** a transport protocol
- Two Components
 - RTP Data Transfer Protocol (RTP)
 - RTP Control Protocol (RTCP)
- Provides end-to-end transport functions
 - Scalable in multicast scenarios
 - Media independent
 - Mixer and translator support
 - RTCP for QoS feedback and session information

Real-time Transport Protocol (RTP)

- No premise on underlying resources
 - layered above transport protocol
 - no reservation / guarantees
- Integrated with applications
- RTP follows principles of
 - Application Level Framing and
 - Integrated Layer Processing



WebRTC / rtcweb

In the last 5 years,
RTP was nearly killed by HTTP Adaptive Streaming (HAS)
but Google brought it back

WebRTC

- free, open project
- adopted by Google, later Mozilla Foundation, Opera, ...
- Real-Time Communications (RTC) for browsers and mobile devices through HTML5 and JavaScript APIs

rtcweb

- Real Time Collaboration on the World Wide Web
- standardize infrastructure for real-time communication in Web browsers
- IETF: formats and protocols
- W3C: APIs for control



RTP

- RTP services are
 - sequencing
 - synchronization
 - payload identification
 - QoS feedback and session information
- RTP supports
 - multicast in a scalable way
 - generic real-time media and changing codecs on the fly
 - mixers and translators to adapt to bandwidth limitations
 - encryption
- RTP is **not** designed for
 - reliable delivery
 - QoS provision or reservation

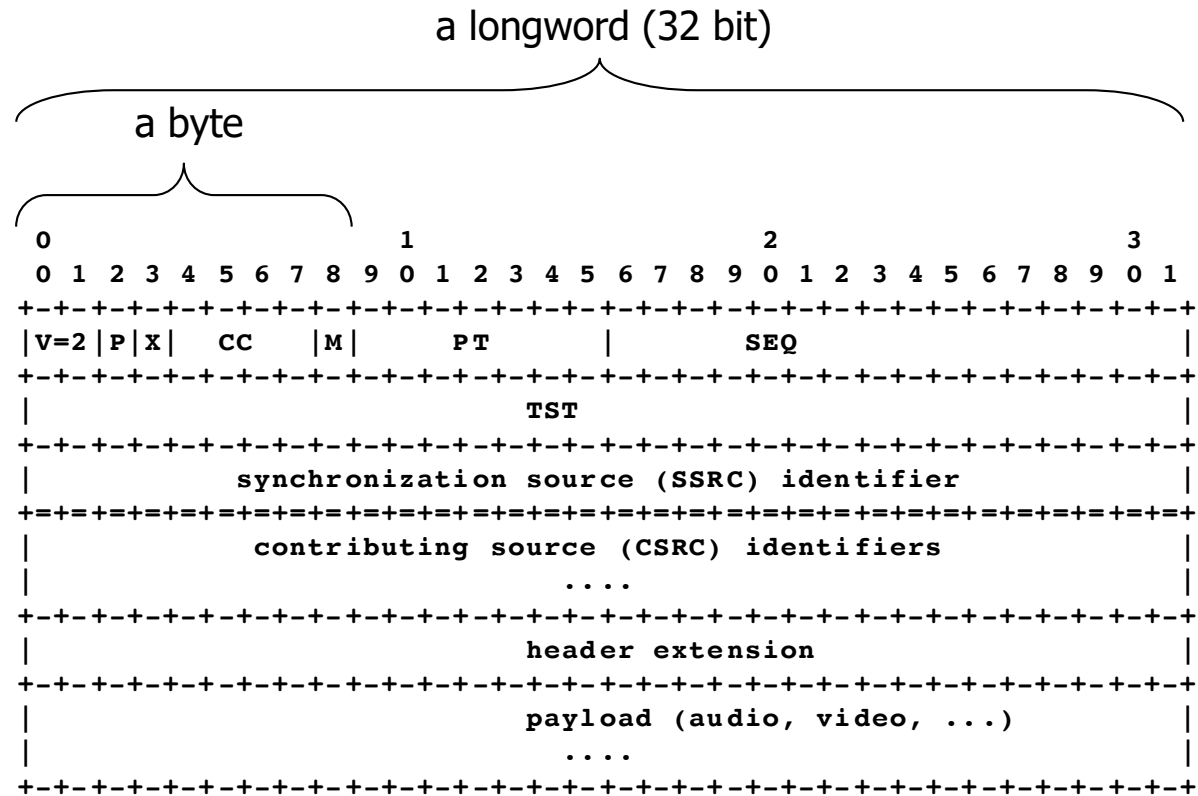
RTP Functions

- RTP with RTCP provides
 - support for transmission of real-time data
 - over multicast or unicast network services

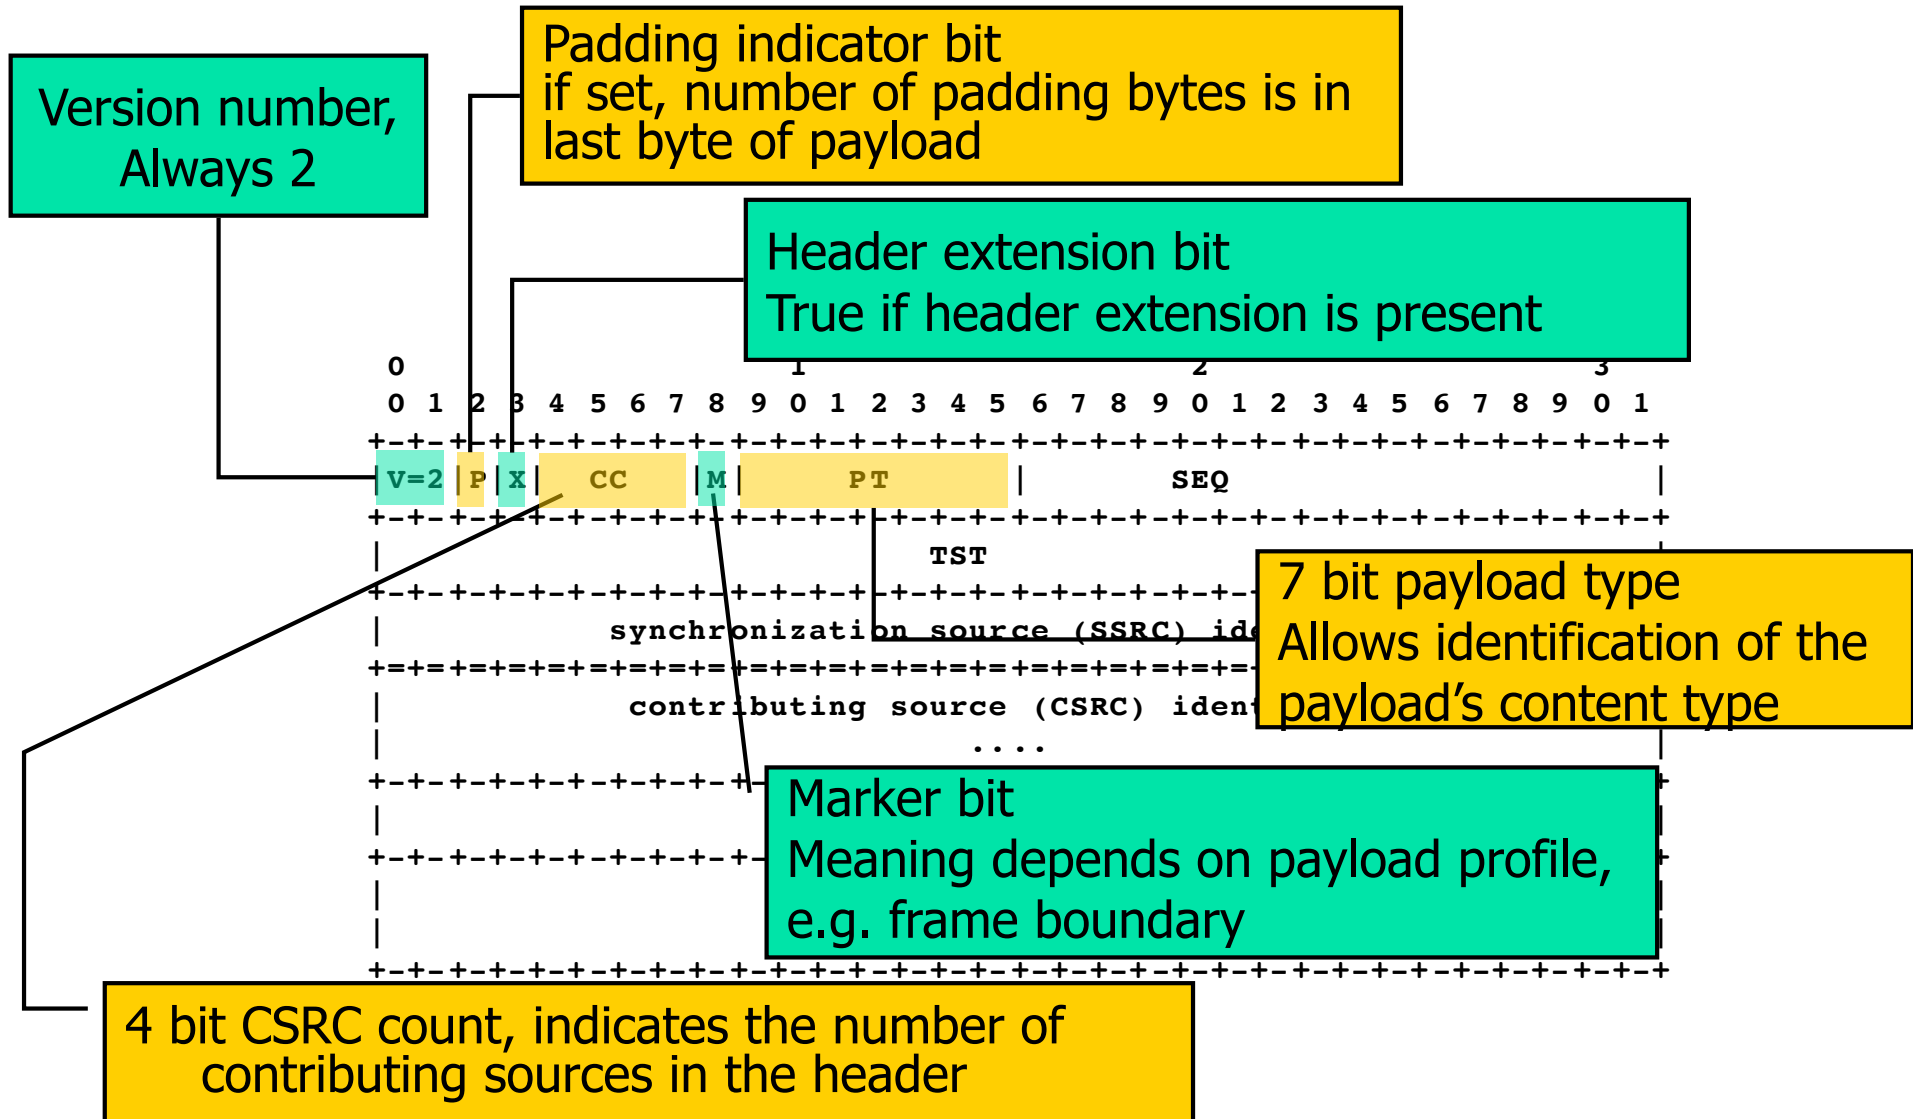
- Functional basis for this
 - Loss detection – sequence numbering
 - Determination of media encoding
 - Synchronization – timing
 - Framing - “guidelines” in payload format definitions
 - Encryption
 - Unicast and multicast support
 - Support for stream “translation” and “mixing” (SSRC; CSRC)

RTP Packet Format

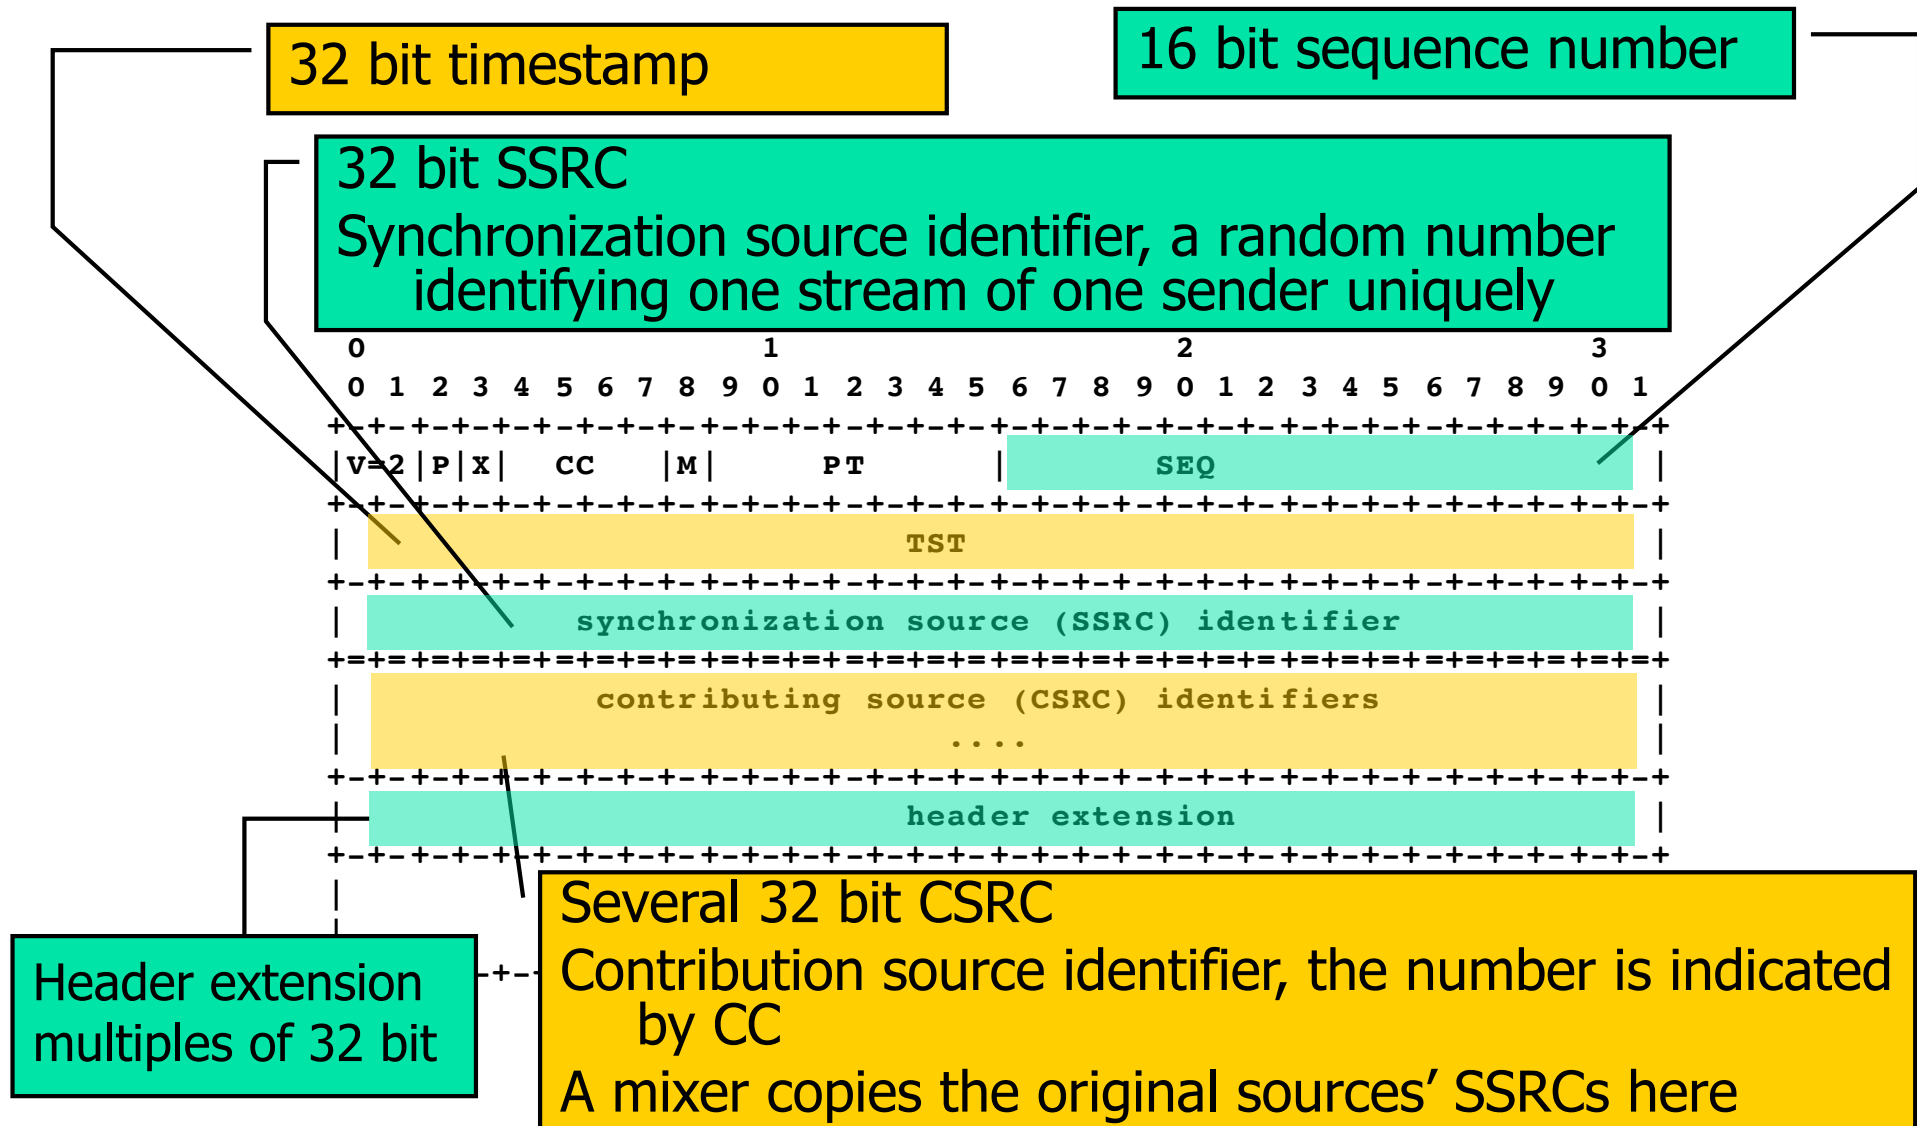
Typical IETF RFC bit-exact representation



RTP Packet Format



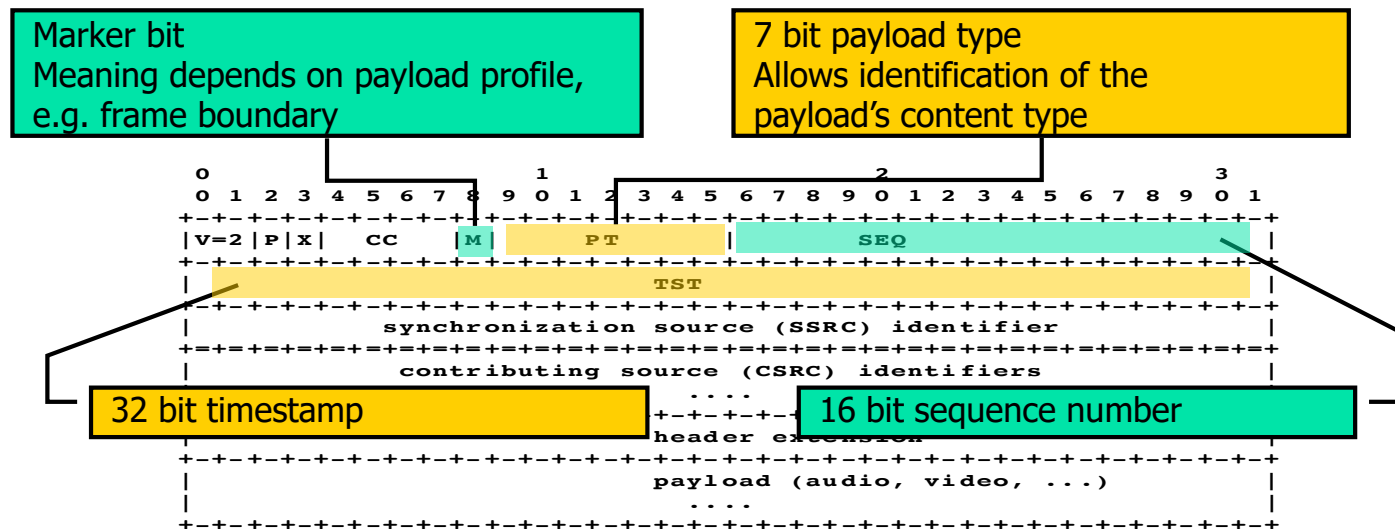
RTP Packet Format



RTP Architecture Concepts

Integrated Layer Processing

- Typical for layered processing
 - Data units sequentially processed by each layer
- Integrated layer processing
 - Adjacent layers tightly coupled
- Therefore, RTP is not complete by itself: requires application-layer functionality/information in header



RTP Packet Format

- Relatively long header (>40 bytes)
 - overhead carrying possibly small payload
 - header compression
 - other means to reduce bandwidth (e.g. silence suppression, or RFC 3389 “RTP Payload for Comfort Noise”)

- No length field
 - Exactly one RTP packet carried in one UDP packet
 - When you use RTP with TCP or SCTP or RTSP or ATM AAL5:
 - do-it-yourself packaging


- Header extensions for payload specific fields possible
 - Specific codecs
 - Error recovery mechanisms

RTP Profile for MPEG-1 Video Payload

International Standard: Moving Pictures Expert Group


- Compression of audio and video for playback (1.5 Mbit/s)
- Real-time decoding

Sequence of I-, P-, and B-Frames



RTP Profile for MPEG-1 Video Payload

Note: MPEG-4 profile for RTP exists, but is much more complex due to H.264's 16-way dependencies.




RTP Profile for MPEG-1 Video Payload

- Fragmentation rules
 - Video sequence header
 - if present, starts at the beginning of an RTP packet
 - GOP sequence header
 - Either at beginning of RTP packet
 - Or following video sequence header
 - Picture header
 - Either at beginning of RTP packet
 - Following GOP header
 - No header can span packets

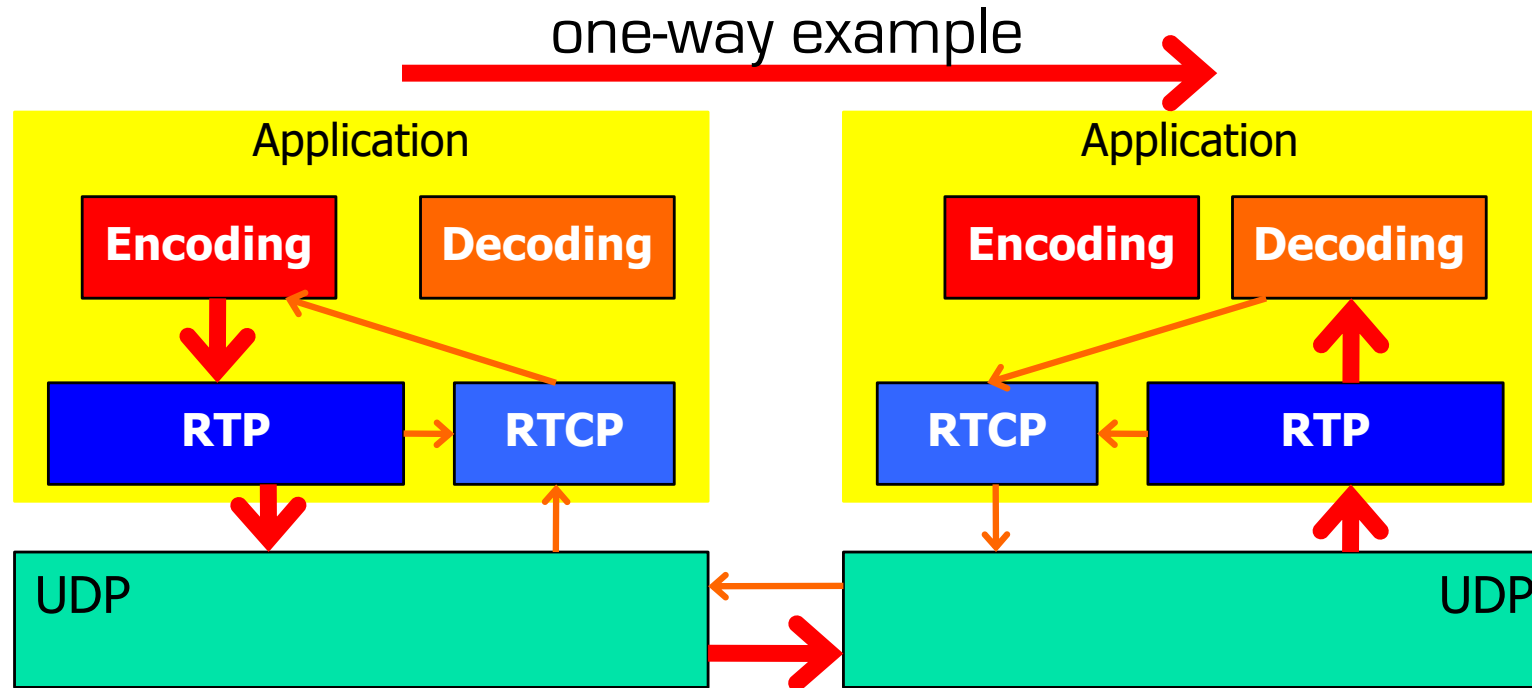
- Marker Bit
 - Set to 1 if packet is end of picture

RTP Profile for MPEG-1 Video Payload



- MPEG-1 Video specific payload header
- TR
 - Temporal reference
 - The same number for all packets of one frame
 - For ordering inside an MPEG GOP
- MBZ
 - Must be zero
- S
 - 1 if sequence header is in this packet
- B
 - 1 if payload starts with new slice
- E
 - 1 if last byte of payload is end of slice
- P
 - 3 bits that indicate picture type (I, P, B or D)
- FBV, BFC, FFV, FFC
 - Indicate how a P or B frame is related to other I and P frames (copied from last frame header)

RTP-enabled Quality Adaptation



- Component interoperations for control of quality
- Evaluation of sender and receiver reports
- Modification of encoding schemes and parameters
- Adaptation of transmission rates
- Hook for possible retransmissions (outside RTP)

RTP Control Protocol (RTCP)


Companion protocol to RTP (tight integration with RTP)

- Monitoring
 - of QoS
 - of application performance
- Feedback to members of a group about delivery quality, loss, etc.
 - Sources may adjust data rate
 - Receivers can determine if QoS problems are local or network-wide
- Loose session control
 - Convey information about participants
 - Convey information about session relationships
- Automatic adjustment to overhead
 - report frequency based on participant count

Typically, “RTP does ...” means “RTP with RTCP does ...”



RTCP Packets



- Several RTCP packets carried in one compound packet
- RTCP Packet Structure
 - SR Sender Report (statistics from active senders: bytes sent -> estimate rate)
 - RR Receiver Report (statistics from receivers)
 - SDES Source Descriptions (sources as “chunks” with several items like canonical names, email, location,...)
 - BYE explicit leave
 - APP extensions, application specific

RTP Mixer

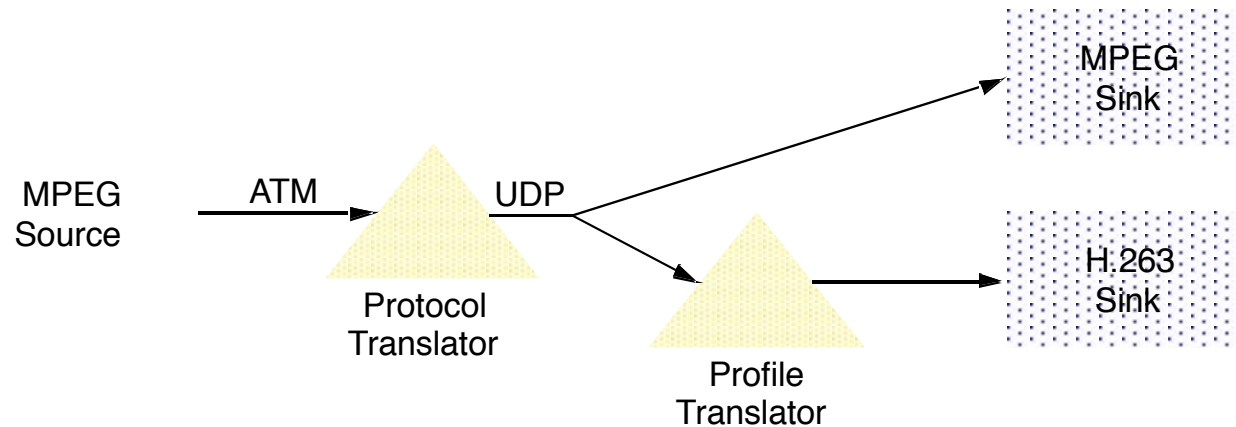
Mixer idea

- If everybody in a large conference talks at the same time, understanding anything is anyway impossible
- Implement mixing of several senders in conference bridges
- Reduce bandwidth in large conferences by mixing several speakers into one stream

Mixer tasks

- Reconstruct constant spacing generated by sender (jitter reduction)
- Translate audio encoding to a lower-bandwidth
- Mix reconstructed audio streams into a single stream
- Resynchronize incoming audio packets
 - New synchronization source value (SSRC) stored in packet
 - Incoming SSRCs are copied into the contributing synchronization source list (CSRC)
- Forward the mixed packet stream

RTP Translator



Translation between protocols

- e.g., between IP and ST-2

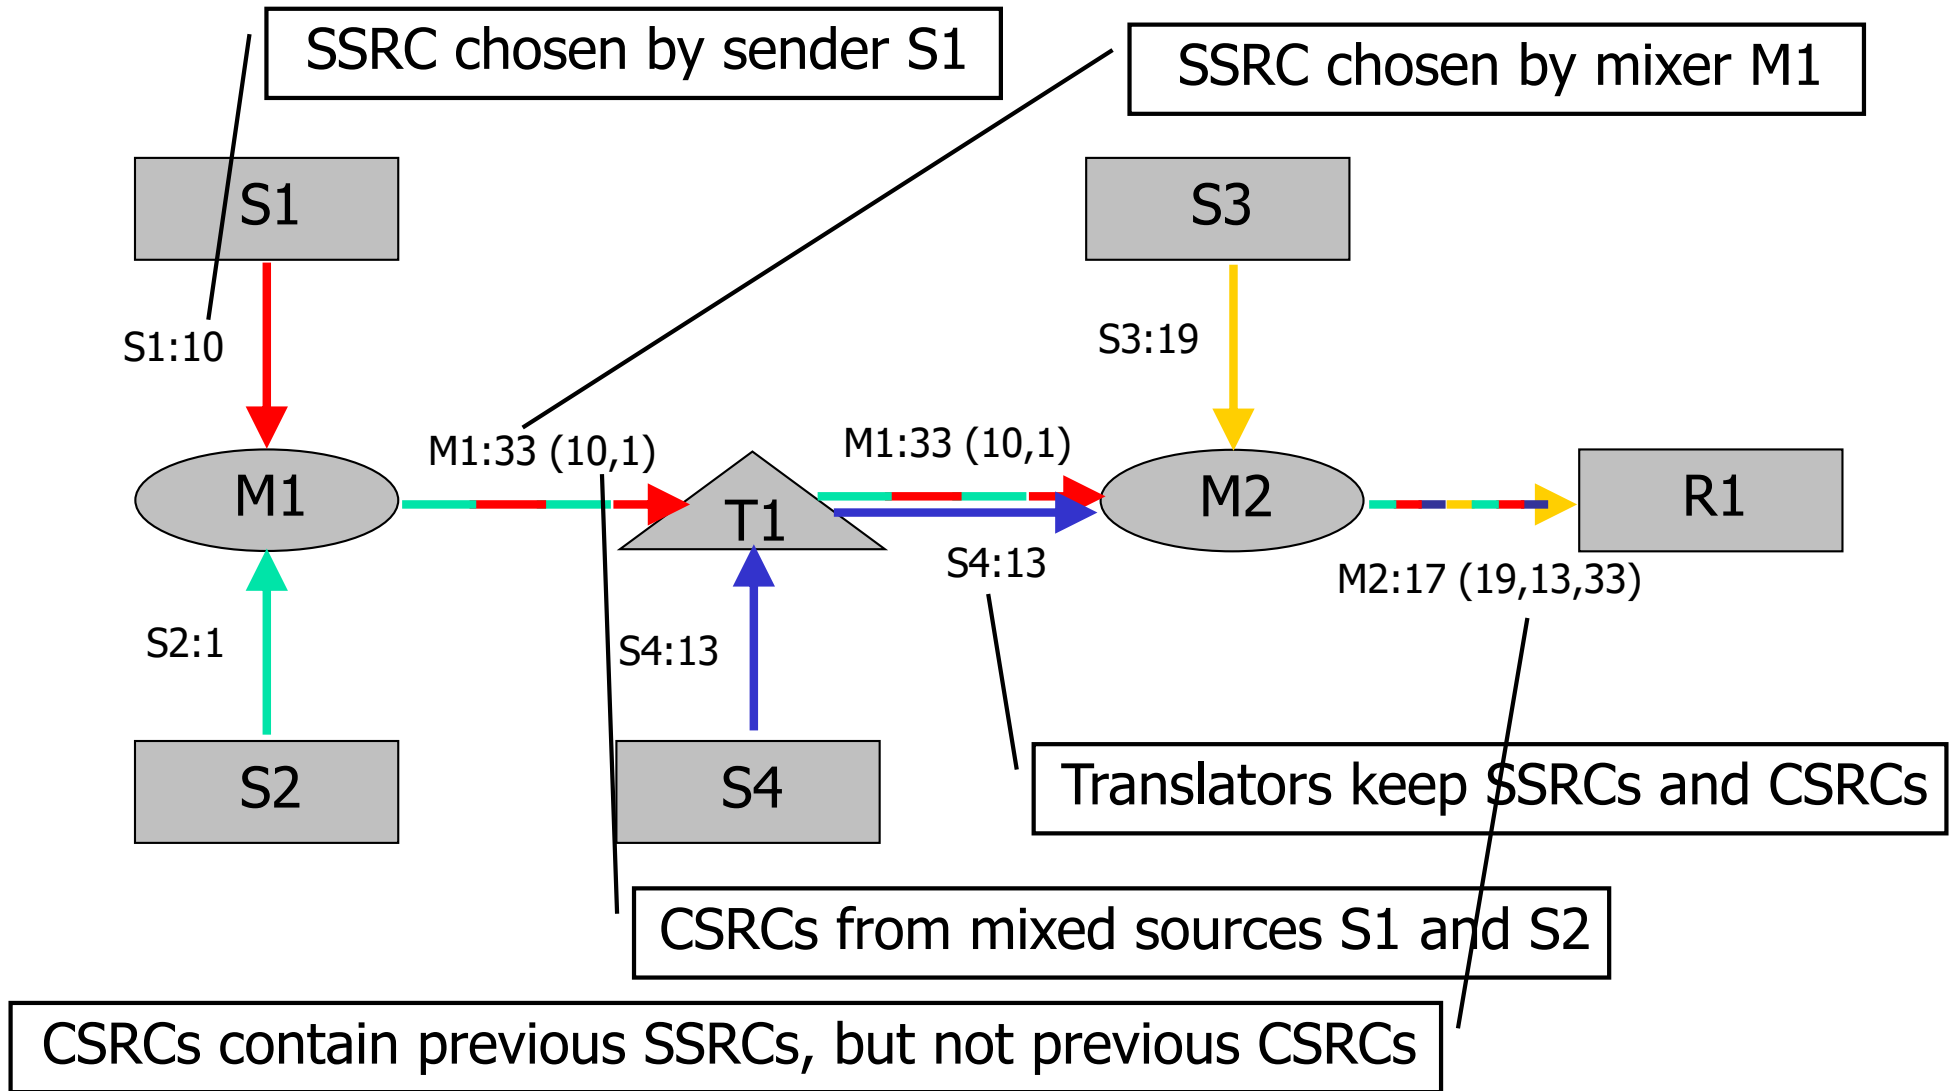
Translation between encoding of data

- e.g. H.265 to H.263
- for reduction of bandwidth without adapting sources

No resynchronization in translators

- SSRC and CSRC remain unchanged

RTP Identifiers



Protocol Development

- Changes and extensions to RTP
 - Scalability to very large multicast groups
 - Congestion Control
 - Algorithms to calculate RTCP packet rate
 - Several profile and payload formats
 - Efficient packetization of Audio / Video
 - Loss / error recovery
 - circuit breakers – worst case behaviour for RTP
 - rmcats – congestion control activity for webrtc

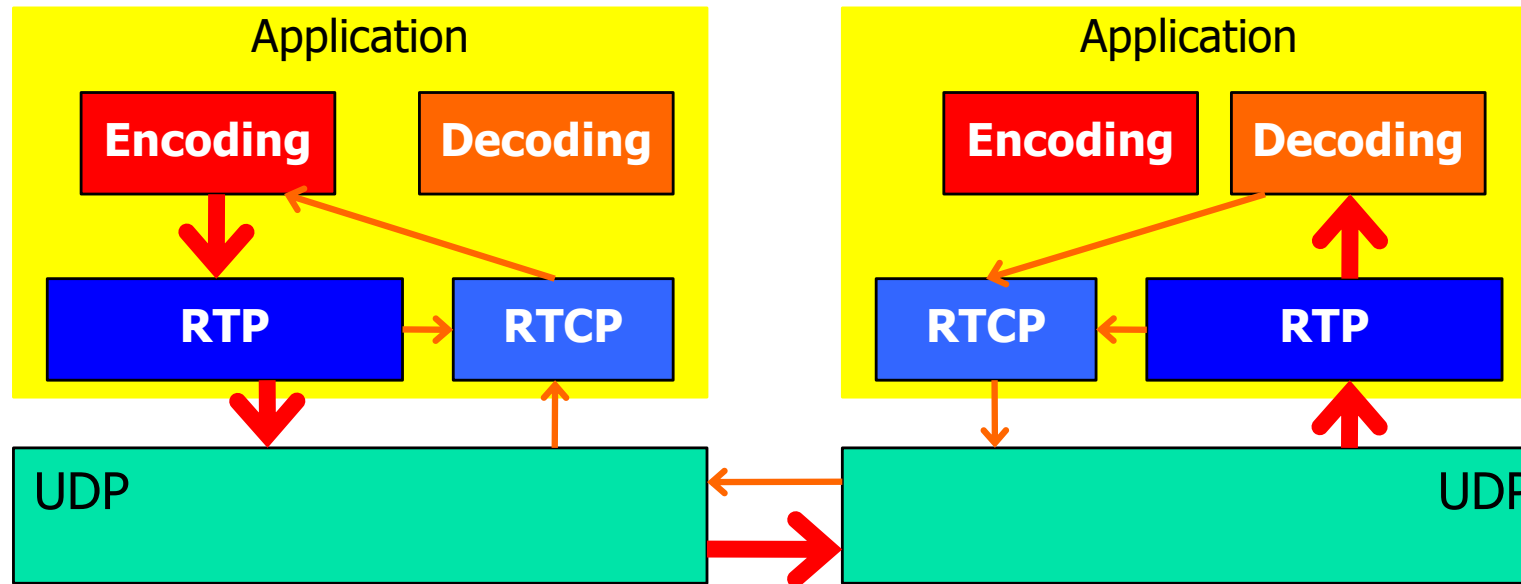


Co-existing with TCP

Adapt audiovisual quality to your
bandwidth share



RTP Quality Adaptation



Application level framing idea

- application knows best how to adapt
- protocol (i.e. RTP) provides information about the network

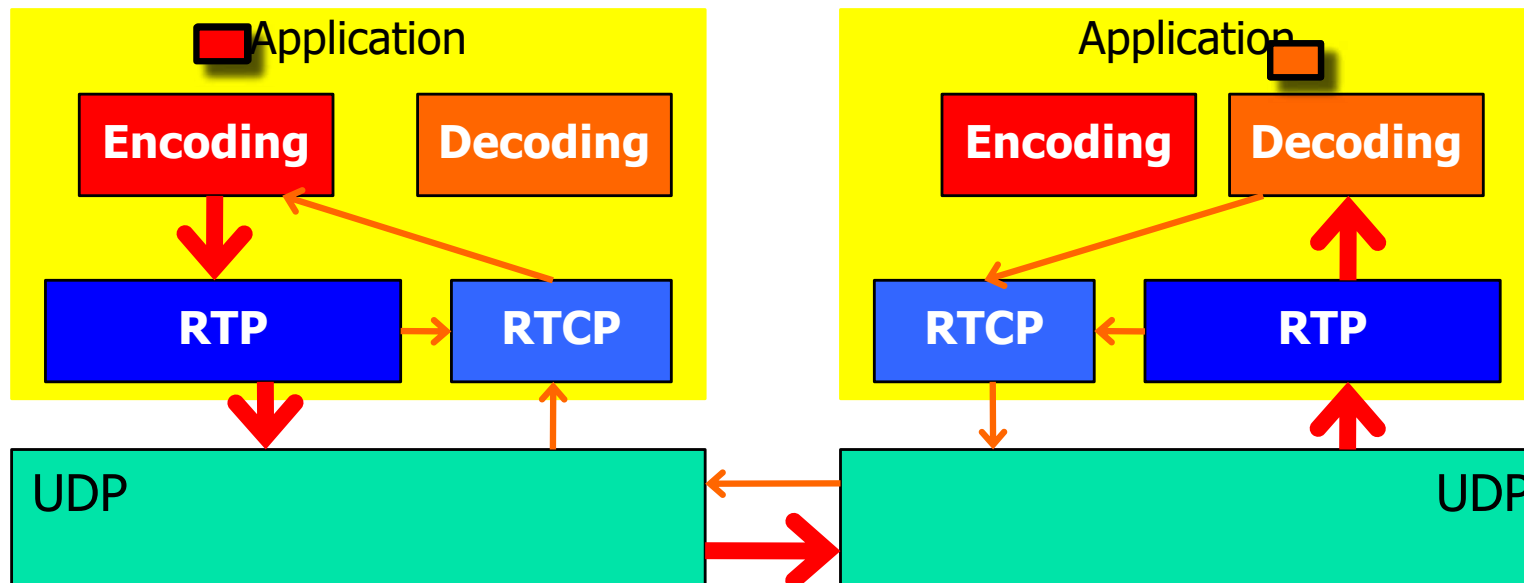
Application can

- evaluate sender and receiver reports
- modify encoding schemes and parameters
- adapt its transmission rates

Loss-Delay Adjustment Algorithm

■ LDA

- An algorithm to stream with RTP in a TCP-friendly way
- use RTCP receiver reports (RR)
 - RTCP sends RR periodically



"The Loss-Delay Based Adjustment Algorithm: A TCP-Friendly Adaptation Scheme",
D. Sisalem, H. Schulzrinne, NOSSDAV 1998

Loss-Delay Adjustment Algorithm

■ LDA

- An algorithm to stream with RTP in a TCP-friendly way
- use RTCP receiver reports (RR)
 - RTCP sends RR periodically
- works like TCP's AIMD
 - but RRs are rare
 - max 5% of RTP BW, max $\frac{3}{4}$ of this RR, equally shared among receivers
 - can't adapt every time
- step one: estimate the bottleneck bandwidth b

Sender



- use packet size and gap sizes

$$b = \frac{1}{n} \sum_{i=1}^n \frac{\text{packet size}(i)}{\text{time}(i+1) - \text{time}(i)}$$

Receiver



Loss-Delay Adjustment Algorithm

■ LDA

- An algorithm to stream with RTP in a TCP-friendly way
- use RTCP receiver reports (RR)
 - RTCP sends RR periodically
- works like TCP's AIMD
 - but RRs are rare
 - can't adapt every time
- no loss:
 - use "AIR" – additive increase *rate*
 - but never more than 1 packet/RTT
- loss:
 - RTCP counts losses,
 l is *fraction* of lost packets
 - *guess* 3 of those losses in one RTT

$$AIR_t = AIR * \left(1 - \frac{r_t}{b}\right)$$
$$r_{t+1} = r_t + AIR_t$$

$$r_{t+1} = r_t * (1 - l * 3)$$

INF3190 – Data Communication Application Layer

Carsten Griwodz

Email: griff@ifi.uio.no



Coding for adaptation


Adapt audiovisual quality to your
bandwidth share



Coding for Adaptive Streaming: MPEG-1

Frames can be dropped


- In a controlled manner
- Frame dropping does not violate dependancies
- Example: B-frame dropping in MPEG-1



Coding ...: H.264 SVC extensions

H.264: most common codec for MPEG-4


SVC: Scalable Video Codec



Simplified representation of H.264/SVC

- the H.264 motion vectors of a frame can point to 16 different frames
- motion vectors in H.264 can cross I-frames
- ...


Coding ...: H.264 SVC extensions



Simplified representation of H.264/SVC

- the H.264 motion vectors of a frame can point to 16 different frames
- motion vectors in H.264 can cross I-frames
- ...

How to change quality?



The ***BAD CHOICE: PSNR***

Peak Signal-to-Noise Ratio

- you find it everywhere
- but it is really, really bad

Example from
Prof. Touradj Ebrahimi,
ACM MM'09 keynote

Reference



PSNR = 24.9 dB





PSNR = 24.9 dB




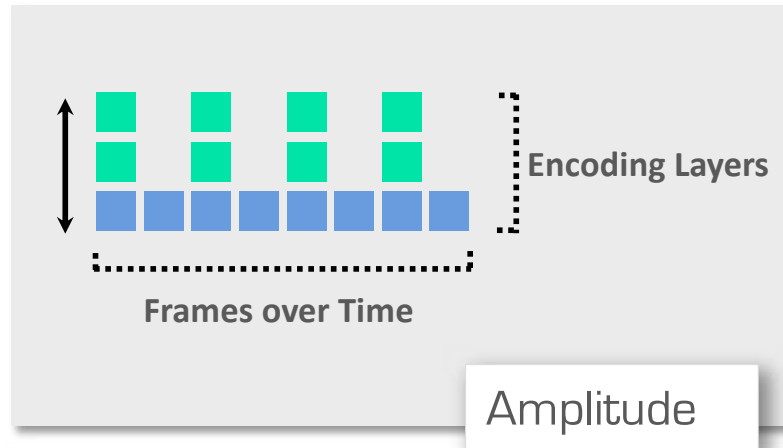
PSNR = 24.9 dB



Coding ...: hierarchical layer coding



Coding ...: perception study






Field study

- mobile devices, free seating, resolution 480x320@30fps, no sunlight, lounge chairs



Blurriness, noise and motion flicker



Blurriness, noise and motion flicker

Three influential factors

Amplitude

Most dominant effect

Flicker is

- almost undetectable for amplitudes with H.264 quantization factors < 8
- almost always detectable for larger amplitudes

Content

Minor effect – within the class

But

- content can influence flicker perception;
- low interaction for noise flicker and stronger for blur flicker

Frequency

Major effect

Acceptance thresholds compared to constant low quality video:
worse when above 1 Hz,
often better when below 0.5 Hz

Remember for later:

- change at most once a second, better every two seconds


Dynamic Adaptive Streaming over HTTP

The State-of-the-Art



Dynamic Adaptive Streaming over HTTP

Divide video into segments at recording time or convert later




Complete little movies

Choose the segment duration

Choose the number of layers

Choose the adaptation strategy

Typical segment lengths:
2-10 seconds
(2-hour movie → 3600++ small, indexed videos)



Dynamic Adaptive Streaming over HTTP

UDP-based streaming

- resists packets loss
- random loss

DASH & similar


- scales to available bandwidth
- congestion loss

Applications

- IPTV
- DVB-H
- 4G telephony
- video conferencing
- classical RTSP/RTP servers

Applications

- Commercial VoD: Netflix, Akamai, Microsoft, Apple, Comoyo, ...
- MPEG DASH
- Free VoD: Youtube, Metacafe, Dailymotion, Vimeo, Rewer, Flixya ...



Dynamic Adaptive Streaming over HTTP

UDP-based streaming

DASH & similar

Challenges to be addressed

Resilience to packet loss


Possibly resilience to bit errors

Possibly active adaptation (server-side decision)


Resilience to buffer underruns

Active adaptation (client-side decision)

Works with web caches



Fluctuating Bandwidth Problem




Adaptive Delivery: Tested Systems

- Adobe Strobe Media Playback (v1.6.328 for Flash 10.1) using HTTP Dynamic Streaming Format
- Apple's native iPad player (iOS v4.3.3) using native HLS format
- Microsoft Silverlight/IIS Smooth (v4.0.60531.0 on Win7) using native Smooth format and default desktop scheduler
- Netview Media Client (v2011-10-10) using Apple HLS format (worst case) and Netview 3G scheduler



Comparison of Existing Quality Schedulers

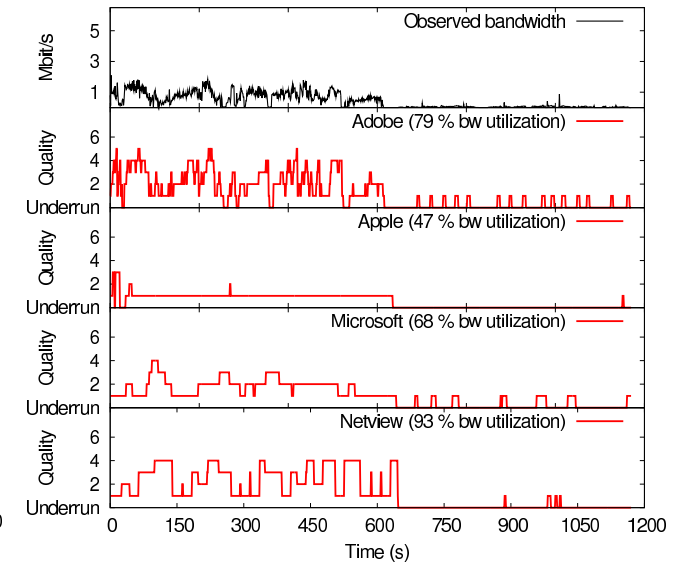


Bus:



Ferry:



Metro:



Distribution Architectures

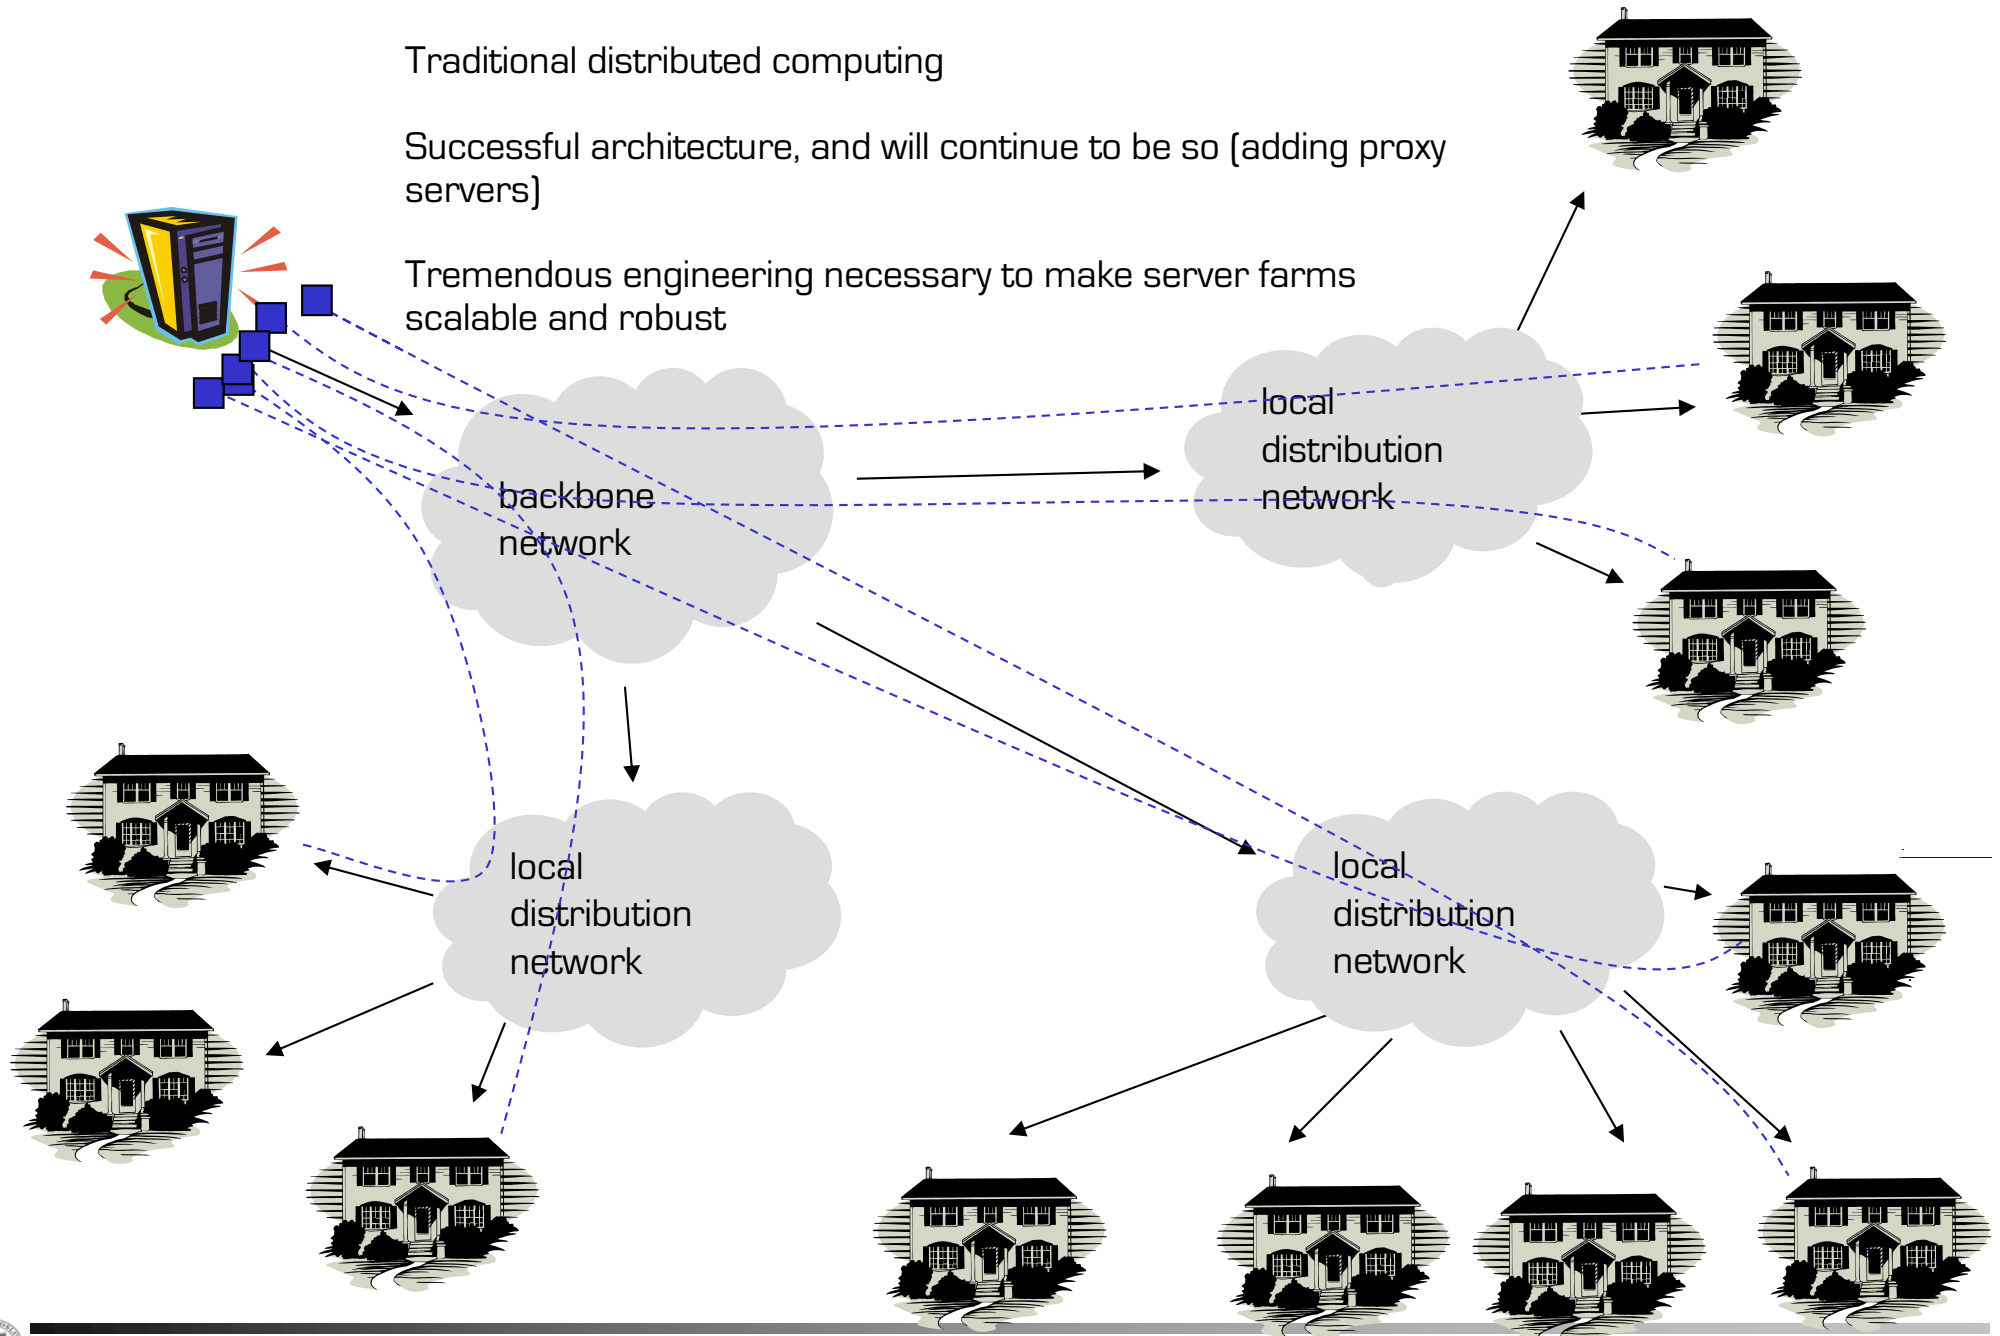


Client-Server

Traditional distributed computing

Successful architecture, and will continue to be so (adding proxy servers)

Tremendous engineering necessary to make server farms scalable and robust



Distribution with proxies


Hierarchical distribution system

- E.g. proxy caches that consider popularity

Popular data replicated more frequently and kept close to clients

Unpopular ones close to the root servers

→ Where to keep copies?



Zipf distribution and features

Popularity


- Estimate the popularity of movies (or any kind of product)
- Frequently used: Zipf distribution

$$z(i) = \frac{C}{i^s} \quad C = 1 / \sum_{n=1}^N \frac{1}{n^s}$$


DANGER

- Zipf-distribution of a process

- can only be applied while popularity doesn't change
- is only an observed property
- a subset of a Zipf-distributed dataset is no longer Zipf-distributed



Access probability distributions



Zipf-distribution


- often used to assign popularities to simulated files

Frequently observed

- Popularity over an entire log does not match a Zipf distribution

Why?

- Zipf-distribution models a snapshot in time
- Popularity of news changes daily
- **Must not** model according to Zipf-distribution with access counts for more than one interval of popularity change

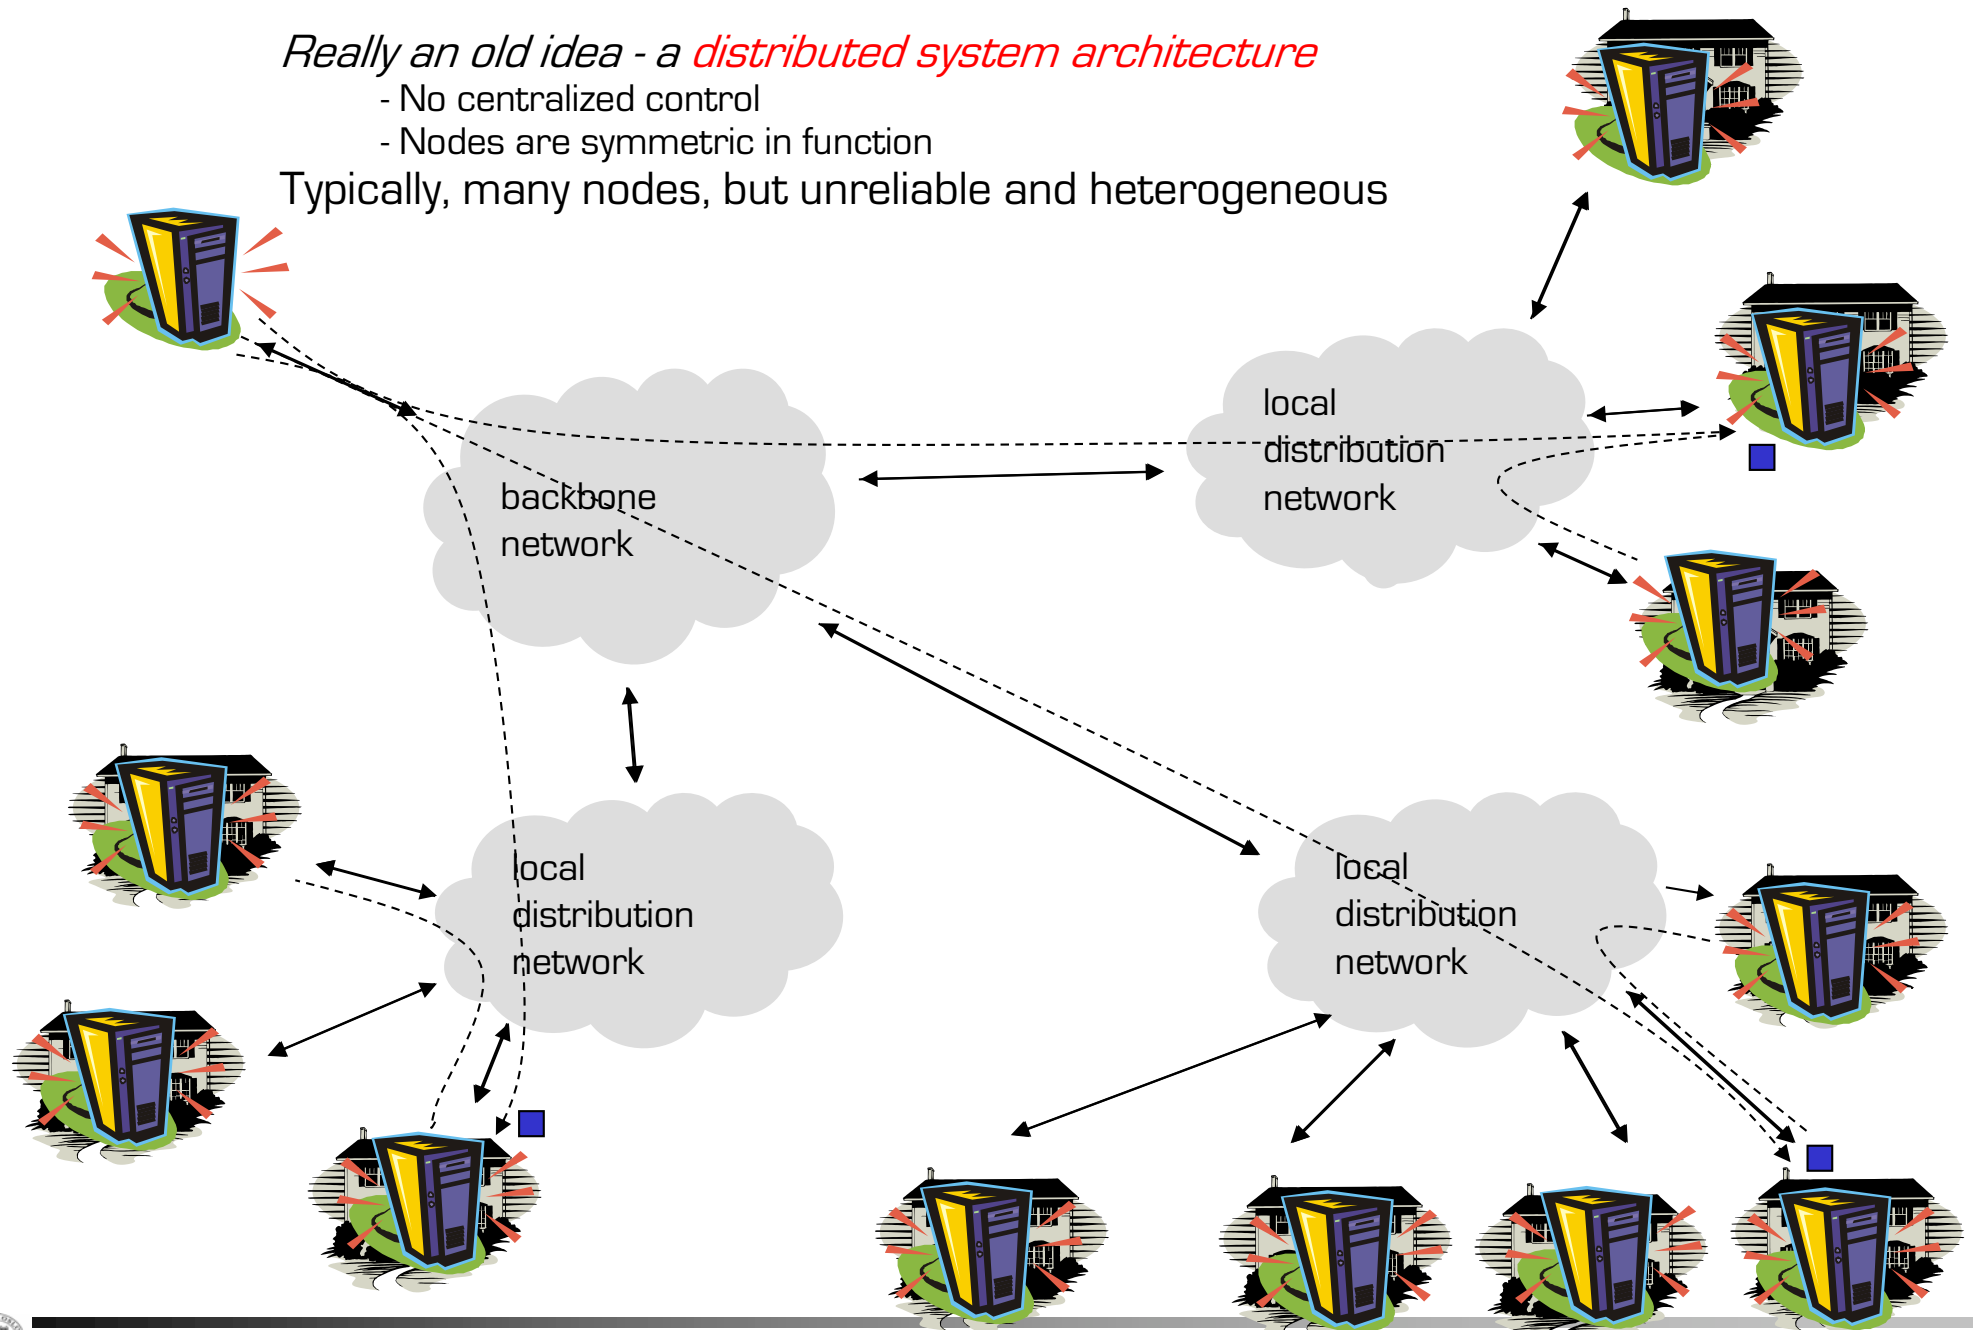


Peer-to-Peer (P2P)

Really an old idea - a *distributed system architecture*

- No centralized control
- Nodes are symmetric in function

Typically, many nodes, but unreliable and heterogeneous



P2P

Many aspects similar to proxy caches

- Nodes act as clients and servers
- Distributed storage
- Bring content closer to clients
- Storage limitation of each node
- Number of copies often related to content popularity
- Necessary to make replication and de-replication decisions
- Redirection

But

- No distinguished roles
- No generic hierarchical relationship: at most hierarchy per data item
- Clients do not know where the content is
 - May need a *discovery protocol*
- All clients may act as roots (origin servers)
- Members of the P2P network come and go (churn)



P2P

Distributed Hash Tables (DHTs)



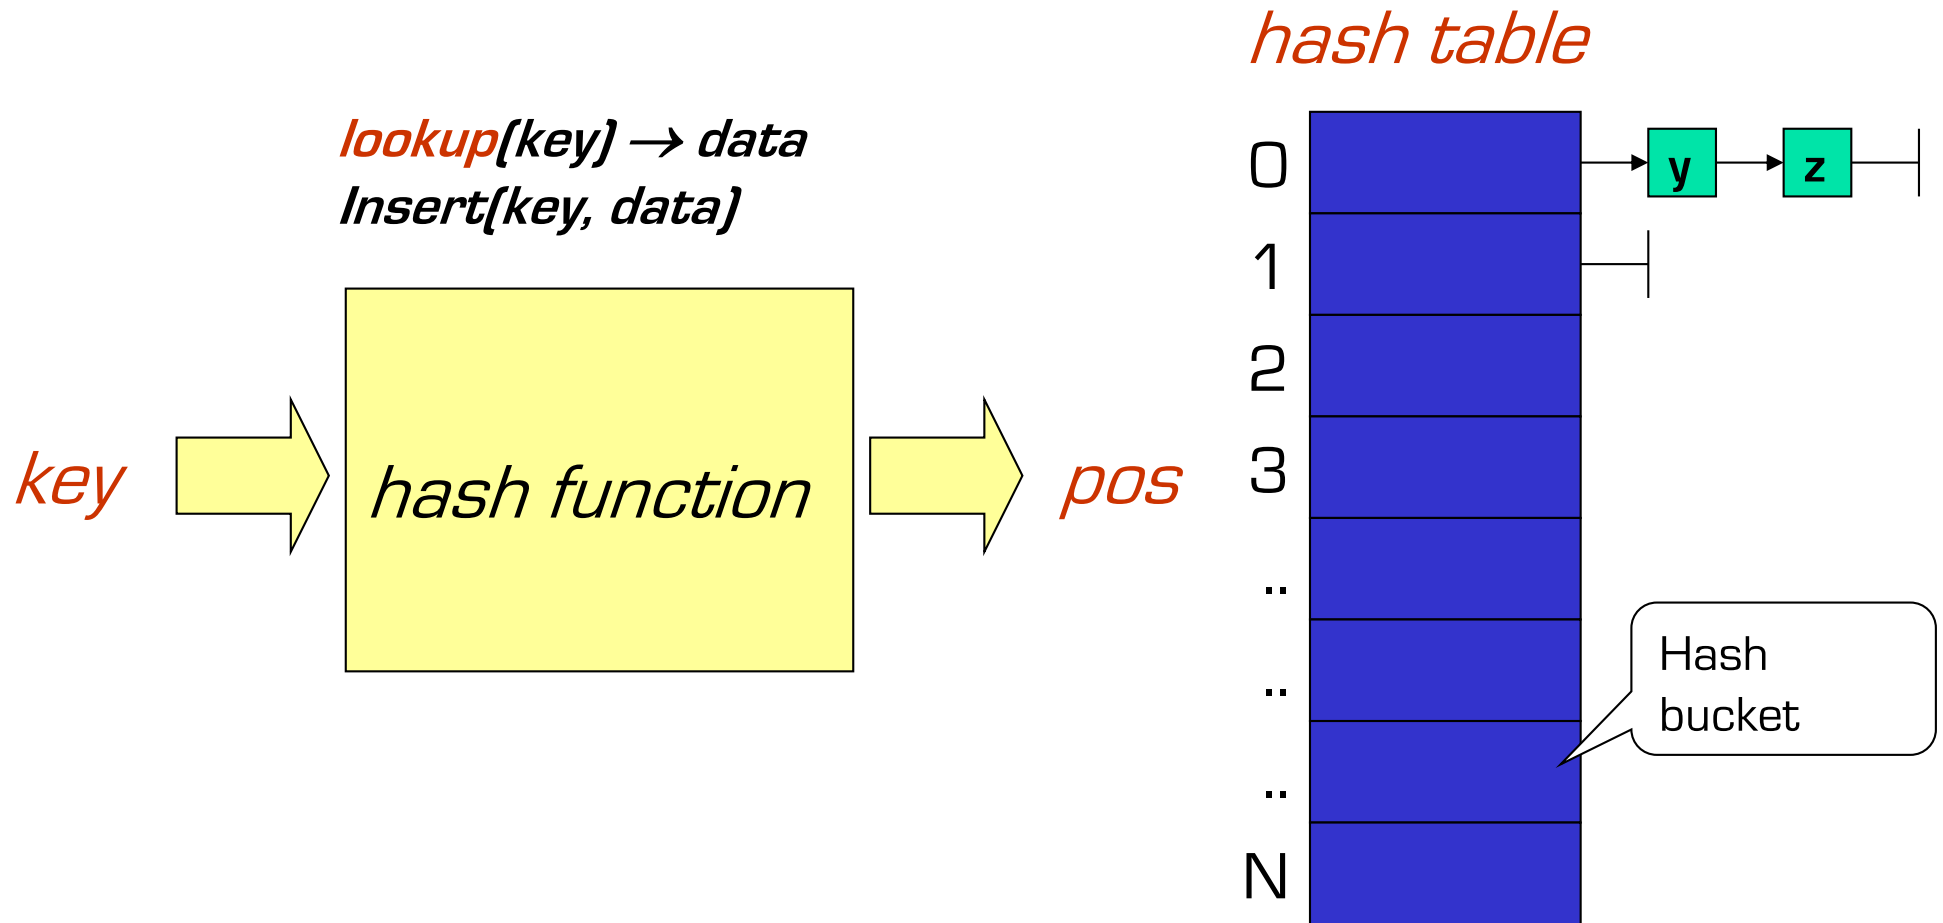
Challenge: Fast, efficient lookups

The BitTorrent tracker is a single point of failure

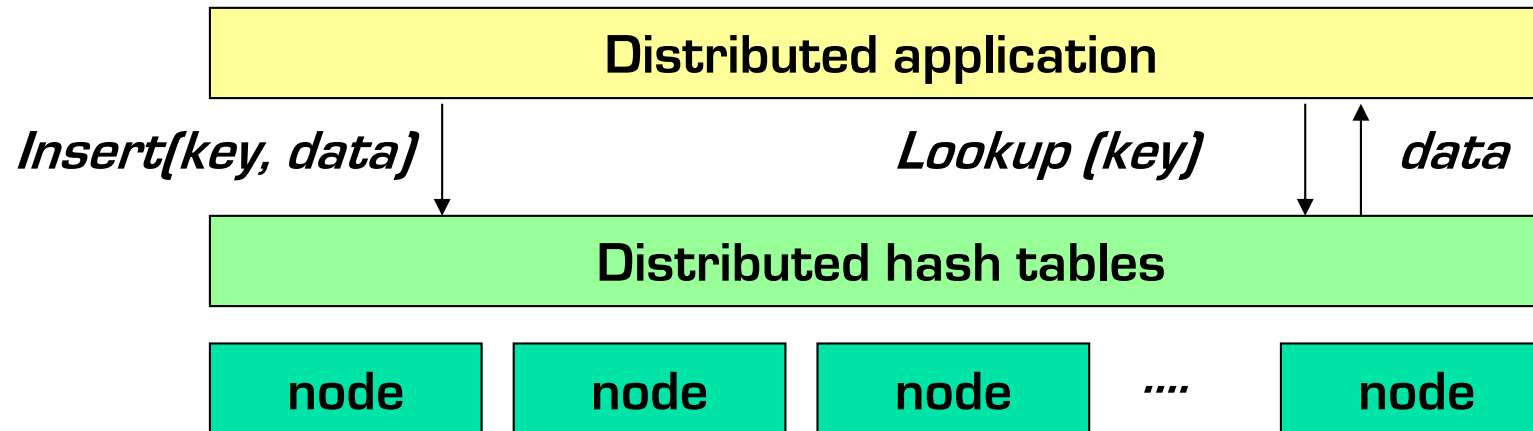
How to distribute tracker functions to many (all) machines?

→ Distributed Hash Tables (DHTs) use a more structured key based routing

Lookup Based on Hash Tables



Distributed Hash Tables (DHTs)



Key identifies data uniquely

Nodes are the hash buckets

The keyspace is partitioned

- usually a node with ID = X has elements with keys close to X
- must define a useful *key nearness metric*
- DHT should balance keys and data across nodes

Keep the hop count small

Keep the routing tables “right size”

Stay robust despite rapid changes in membership

Distributed Hash Tables

Chord



Chord

- Approach taken
 - Only concerned with efficient indexing
 - Distributed index - decentralized lookup service
 - Inspired by consistent hashing: SHA-1 hash
 - Content handling is an external problem entirely
 - No relation to content
 - No included replication or caching
- P2P aspects
 - Every node must maintain keys
 - Adaptive to membership changes
 - Client nodes act also as file servers



Chord IDs & Consistent Hashing

- m -bit identifier space for both keys and nodes
 - Key identifier = SHA-1(key)


Key="LetItBe" $\xrightarrow{\text{SHA-1}}$ ID=54

- Node identifier = SHA-1(IP address)

IP="198.10.10.1" $\xrightarrow{\text{SHA-1}}$ ID=123

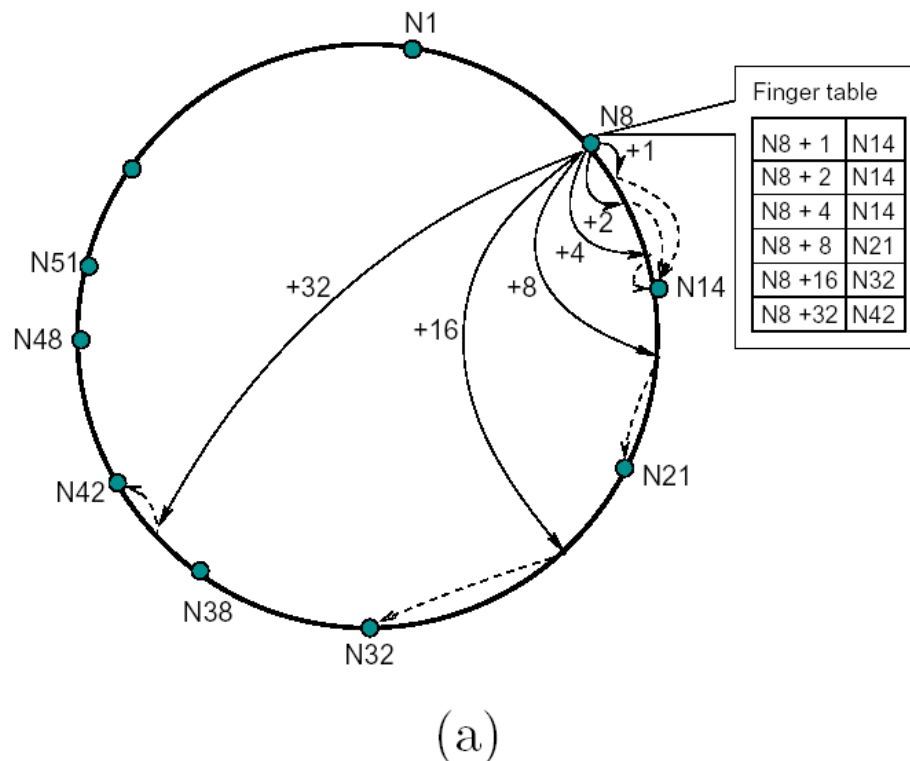
- Both keys and identifiers are **uniformly distributed** in the space of m -bit numbers

- Identifiers ordered in a circle modulo 2^m
- A key is mapped to the first node whose node $ID \geq key ID$




Routing: “Finger Tables”

- Every node knows nodes that represent m other IDs in the ring
- Increase distances between these IDs exponentially
- Finger i points to **successor** of $n+2^i$




Routing: “Finger Tables”

- Every node knows nodes that represent m other IDs in the ring
- Increase distances between these IDs exponentially
- Finger i points to **successor** of $n+2^i$
- Lookup jumps to that node in its lookup table with largest ID where $hash(search\ term) \geq ID$



(a)



(b)

Chord Assessment

Large distributed index

Scalability, fairness, load balancing

- Space complexity: routing tables are size $O(\log(\#nodes))$
- Logarithmic insert effort
- Network topology is **not** accounted for
- Quick lookup in large systems, low variation in lookup costs

Content location

- Run-time complexity: $O(\log(\#nodes))$ lookup steps
- Search by hash key: limited ways to formulate queries

No failure resilience in basic approach

- Easy fix
 - Successor lists allow use of neighbors to failed nodes
 - create several index with different has functions [$O(1)$]



P2P

BitTorrent



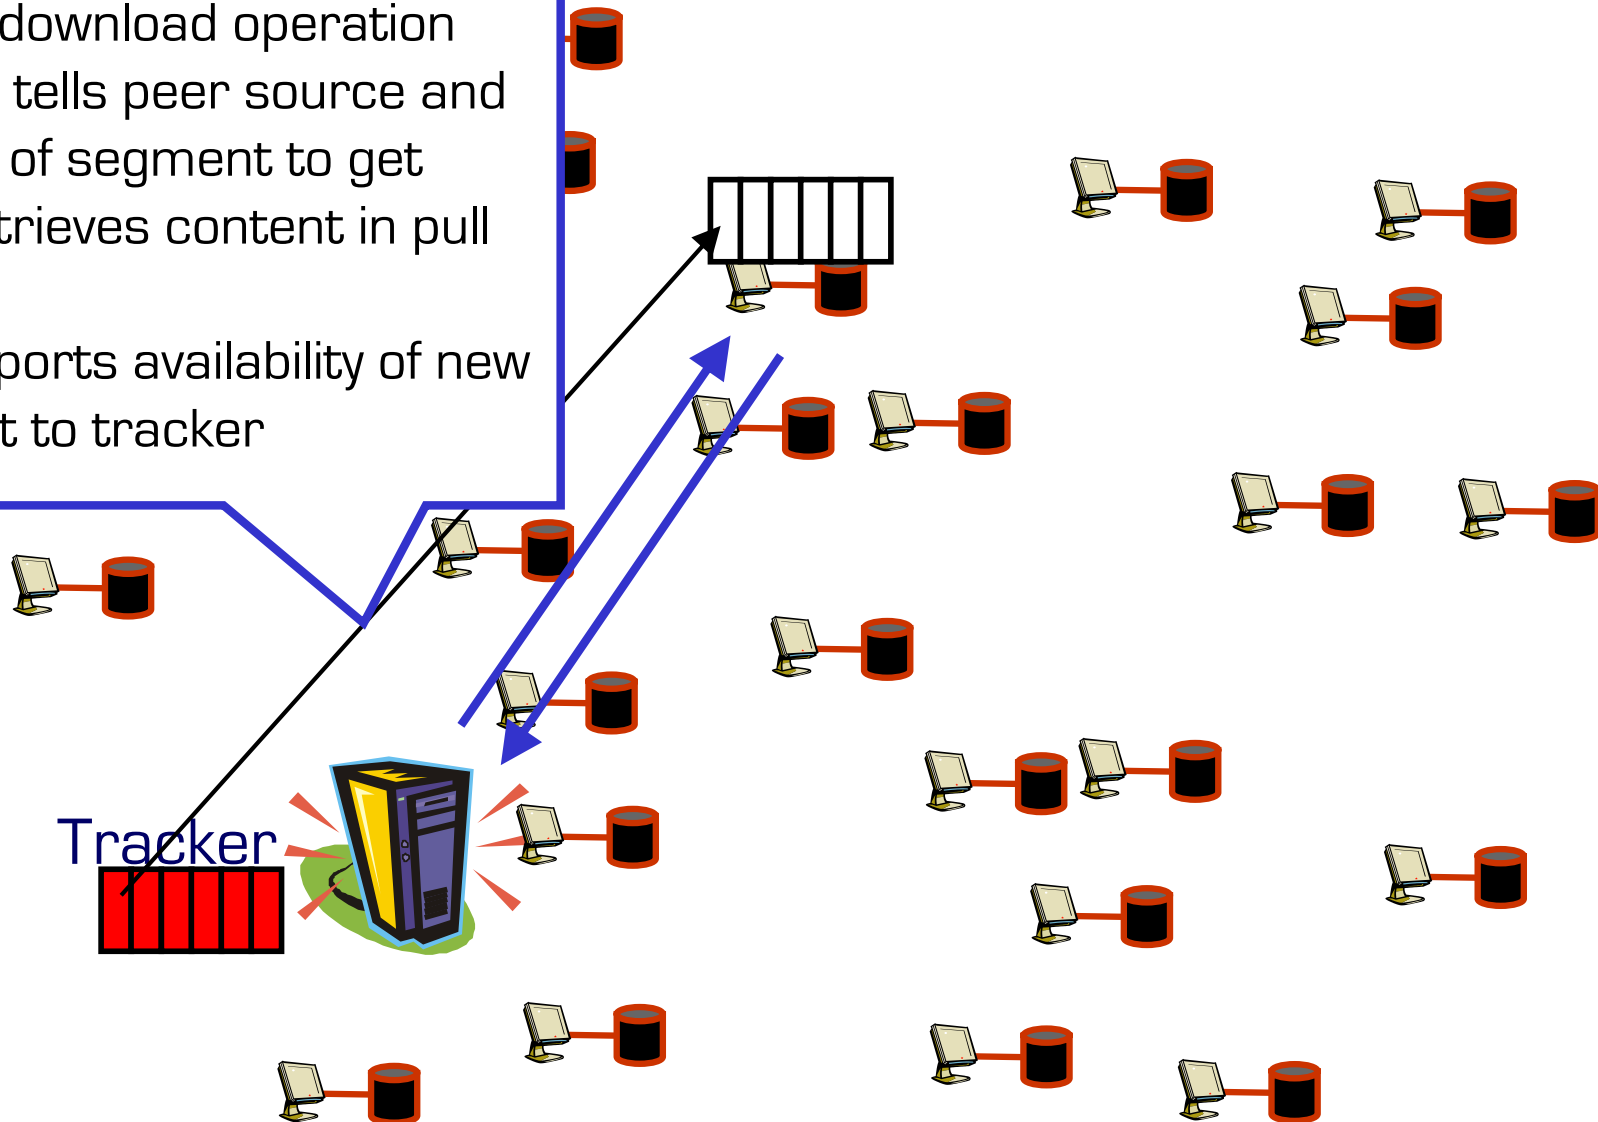
BitTorrent

- Distributed download system
- Content is distributed in segments
- Tracker
 - One central download server per content
 - Approach to fairness (tit-for-tat) per content
 - No approach for finding the tracker
- No content transfer protocol included

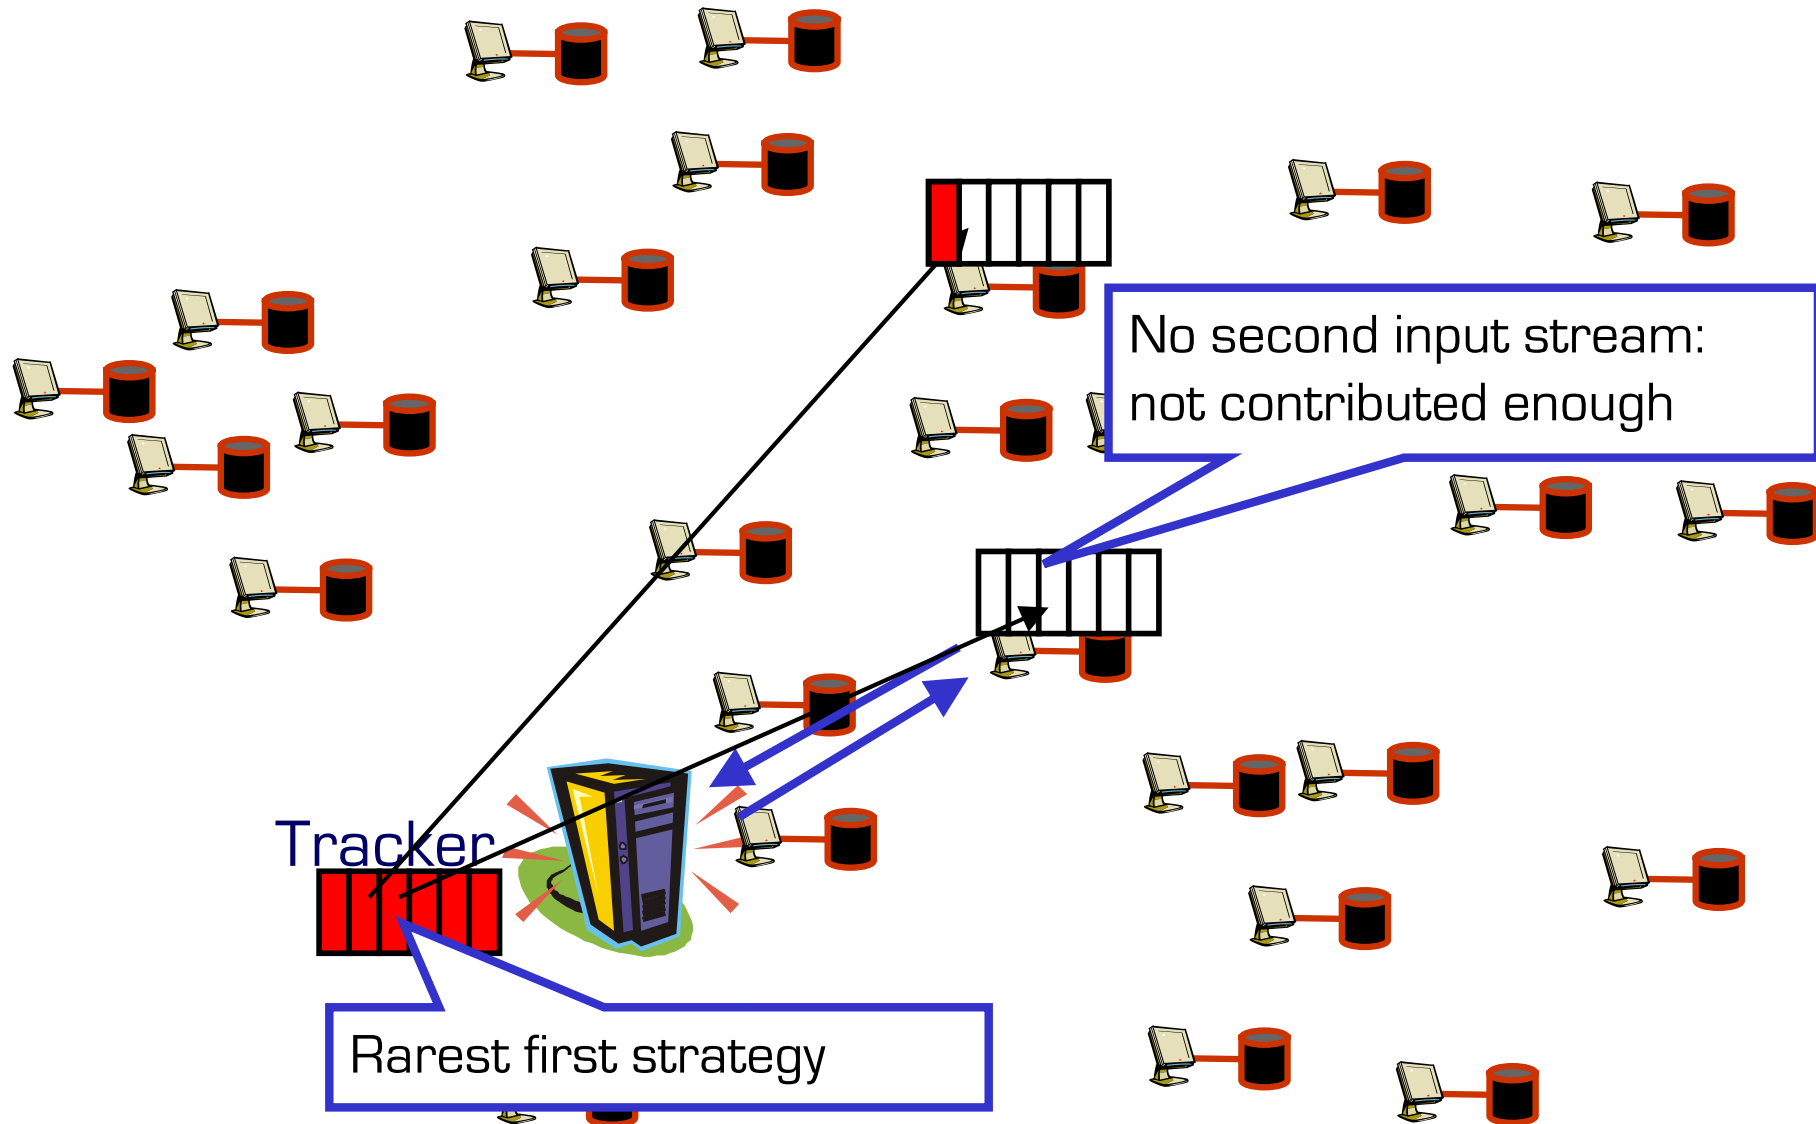
BitTorrent

Segment download operation

- Tracker tells peer source and number of segment to get
- Peer retrieves content in pull mode
- Peer reports availability of new segment to tracker



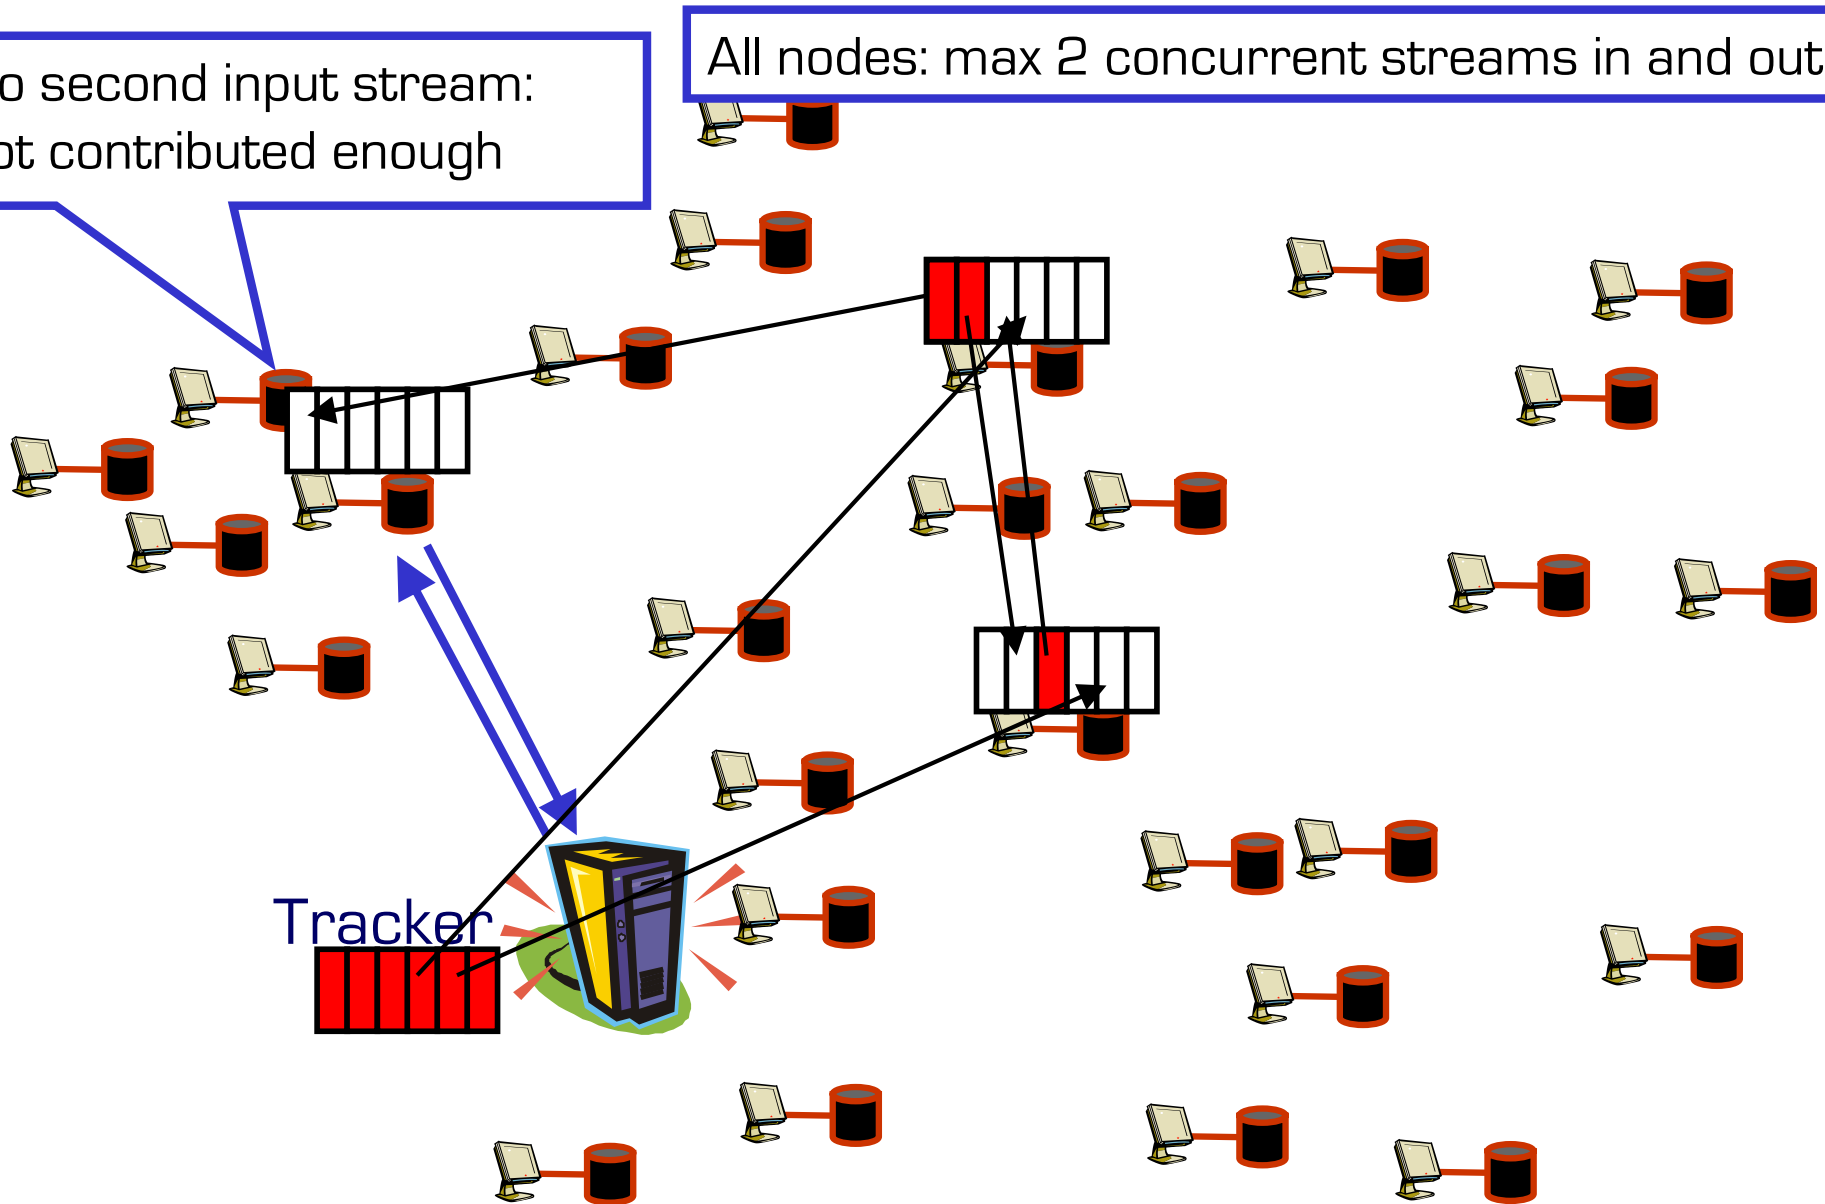
BitTorrent




BitTorrent

No second input stream:
not contributed enough


All nodes: max 2 concurrent streams in and out



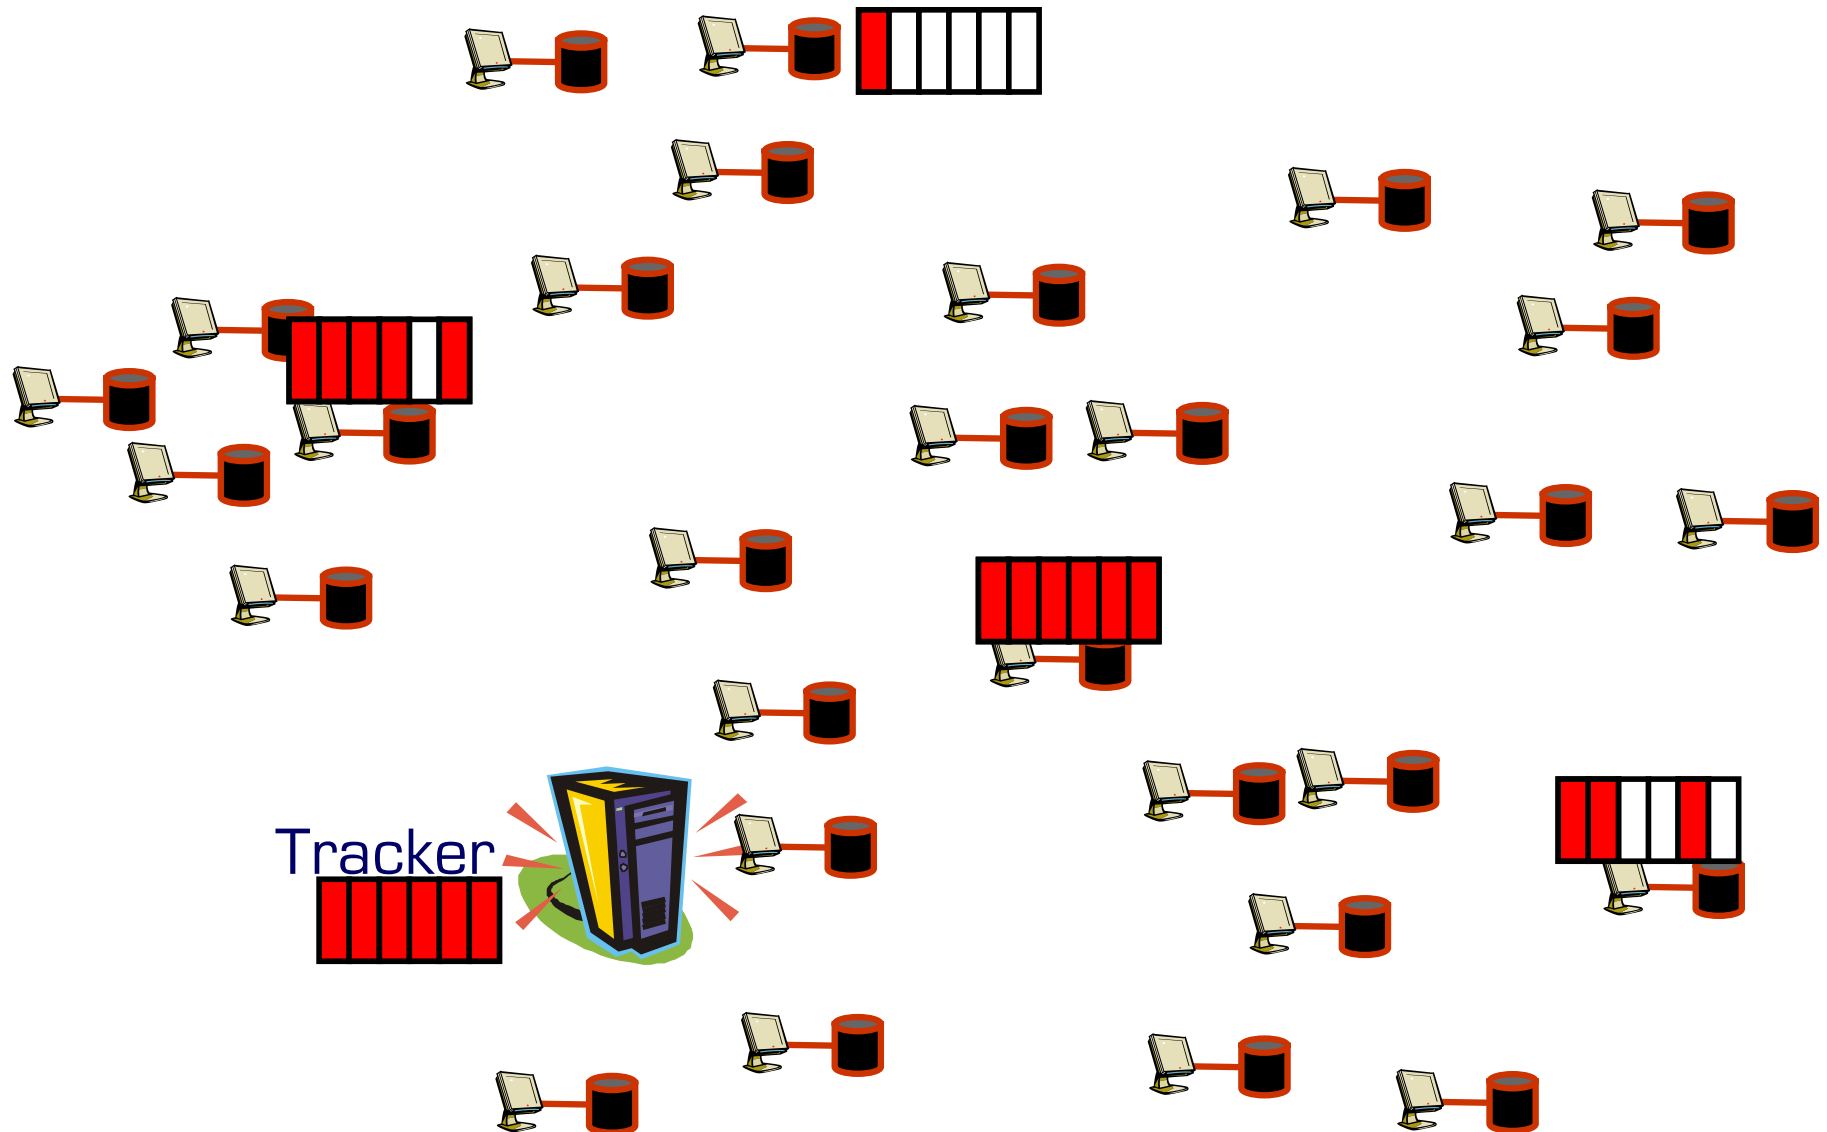
BitTorrent



BitTorrent



BitTorrent



Signalling Protocols: RTSP & SIP



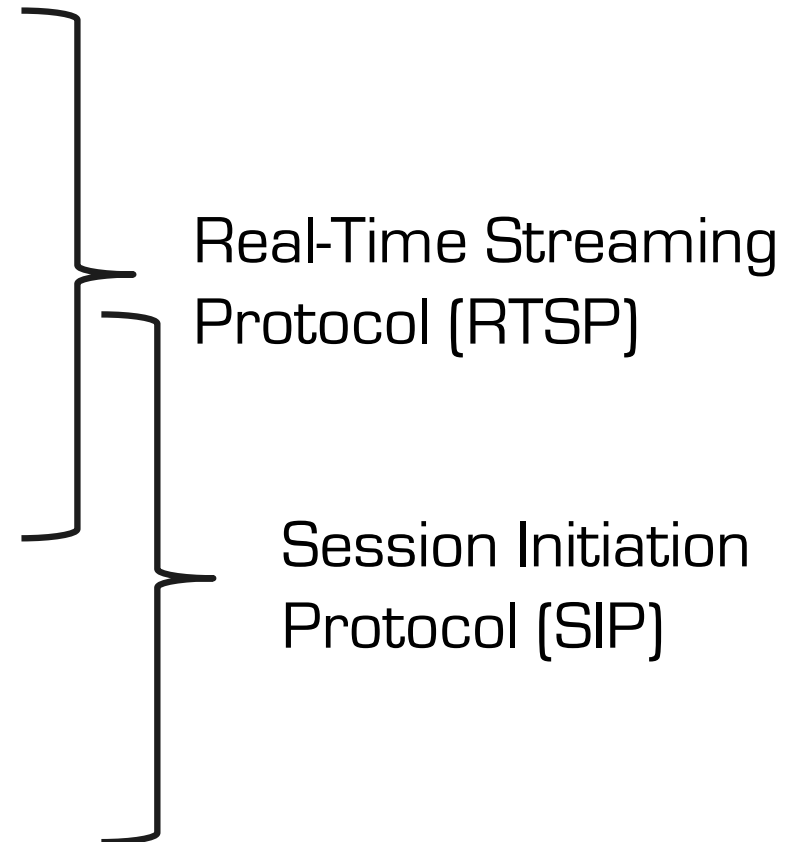
Signaling Protocols

Applications differ

- Media delivery controlled by sender or receiver
- Sender and receiver “meet” before media delivery

Signaling should reflect different needs

- Media-on-demand
 - Receiver controlled delivery of content
 - Explicit session setup
- Internet broadcast
 - Sender announces multicast stream
 - No explicit session setup
- Internet telephony and conferences:
 - Bi-directional data flow, live sources
 - (mostly) explicit session setup, mostly persons at both ends



Real-Time Streaming Protocol (RTSP)

Rough synchronization

- Media description in DESCRIBE response
- Timing description in SETUP response
- Fine-grained through RTP sender reports

Aggregate and separate control of streams possible

Combine several data (RTP) servers

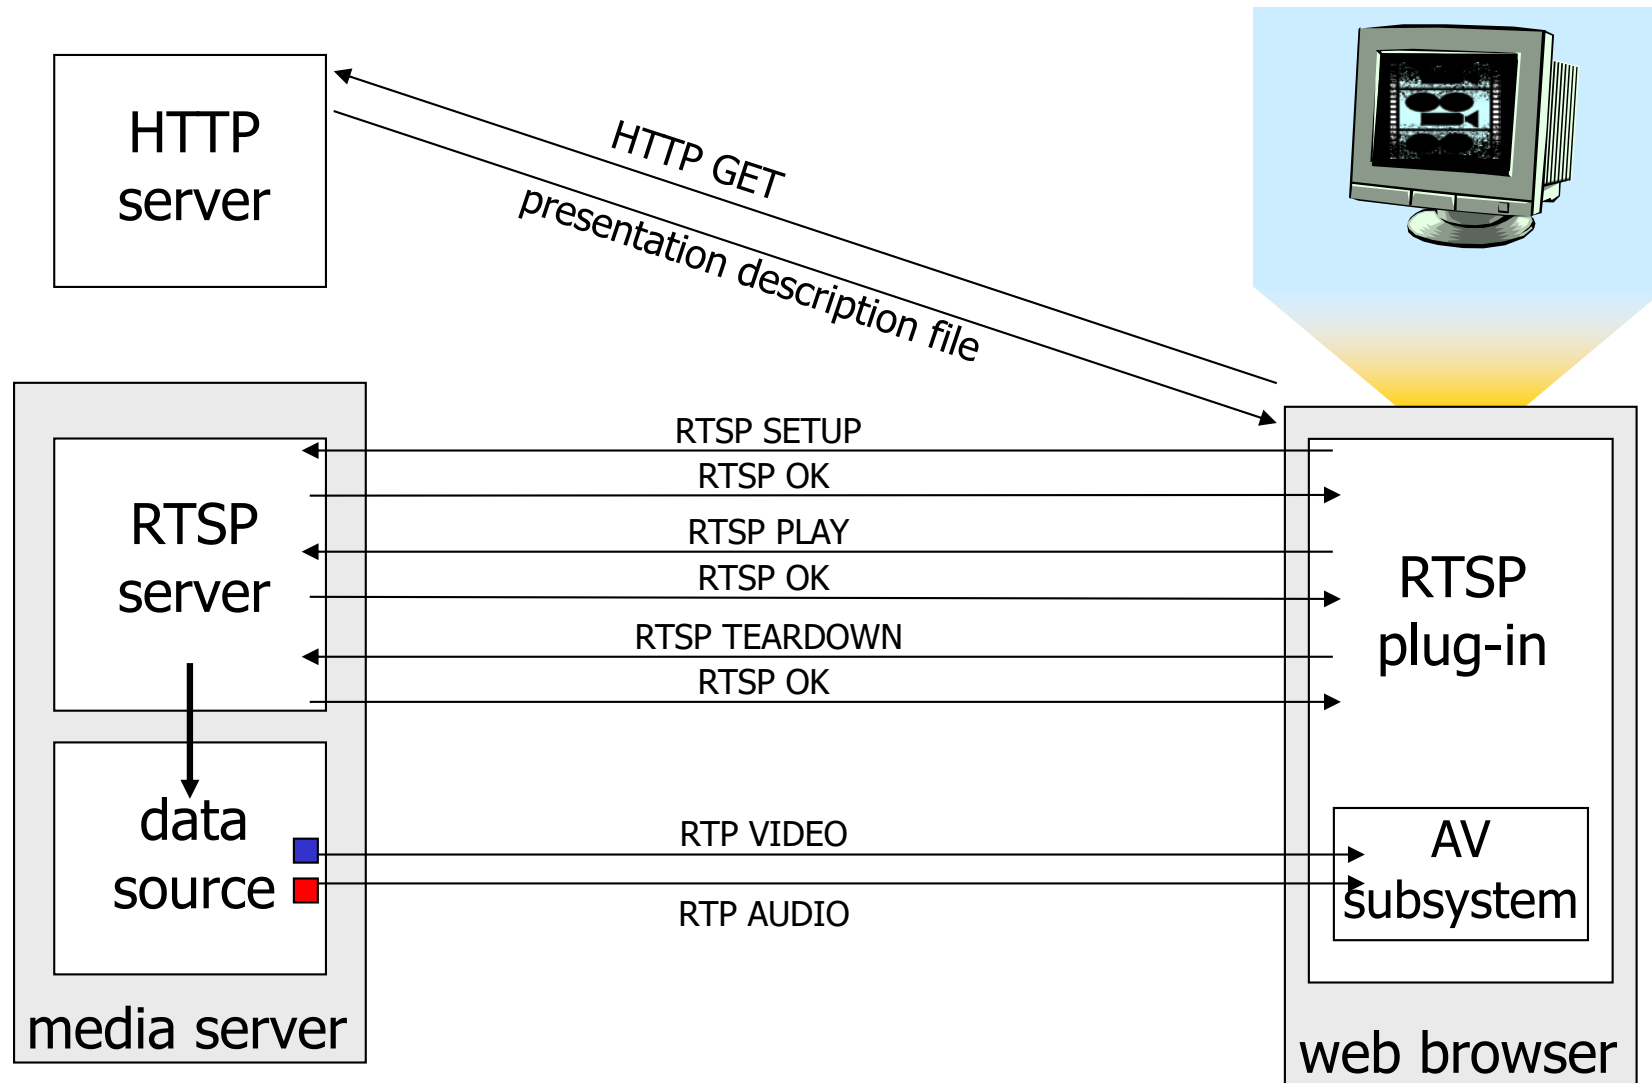
Load balancing by REDIRECT at connect time

Caching

- Much more difficult than web caching
 - interpret RTSP
 - but cache several RTP flows
- Cache must act as an RTP translator
 - otherwise it cannot guarantee to receive packets



RTSP Integration



Session Initiation Protocol (SIP)

Lightweight generic signaling protocol

Internet telephony and conferencing

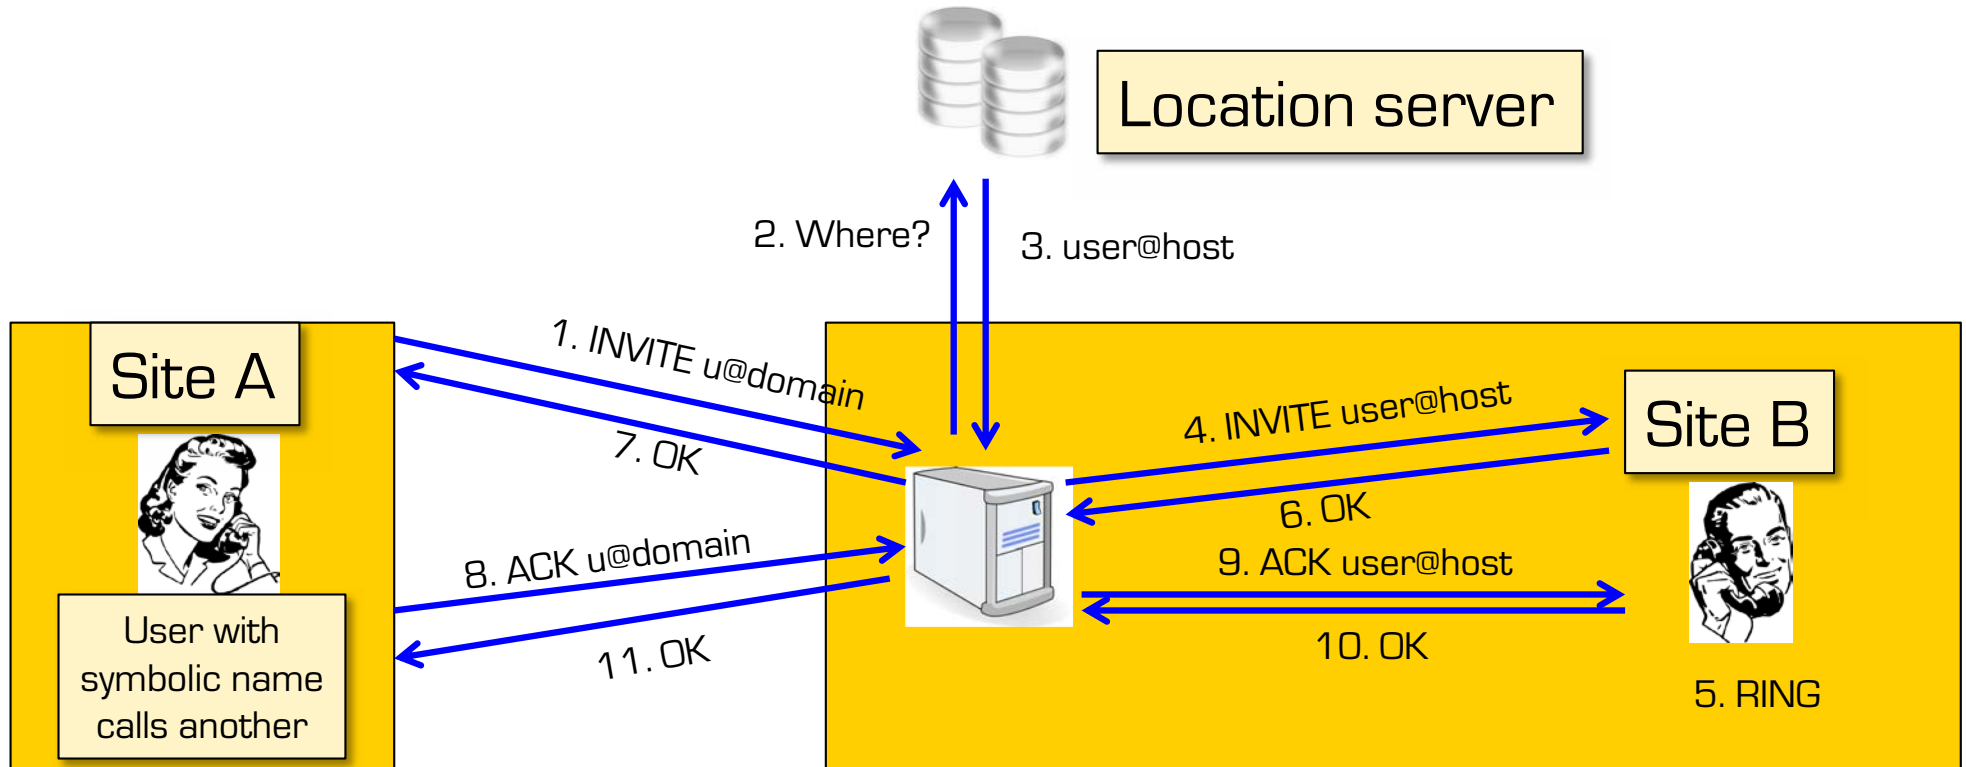
- call: association between number of participants
- signaling association as signaling state at endpoints (no network resources)

Several “services” needed

- Name translation
- User location
- Feature negotiation
- Call control
- Changing features



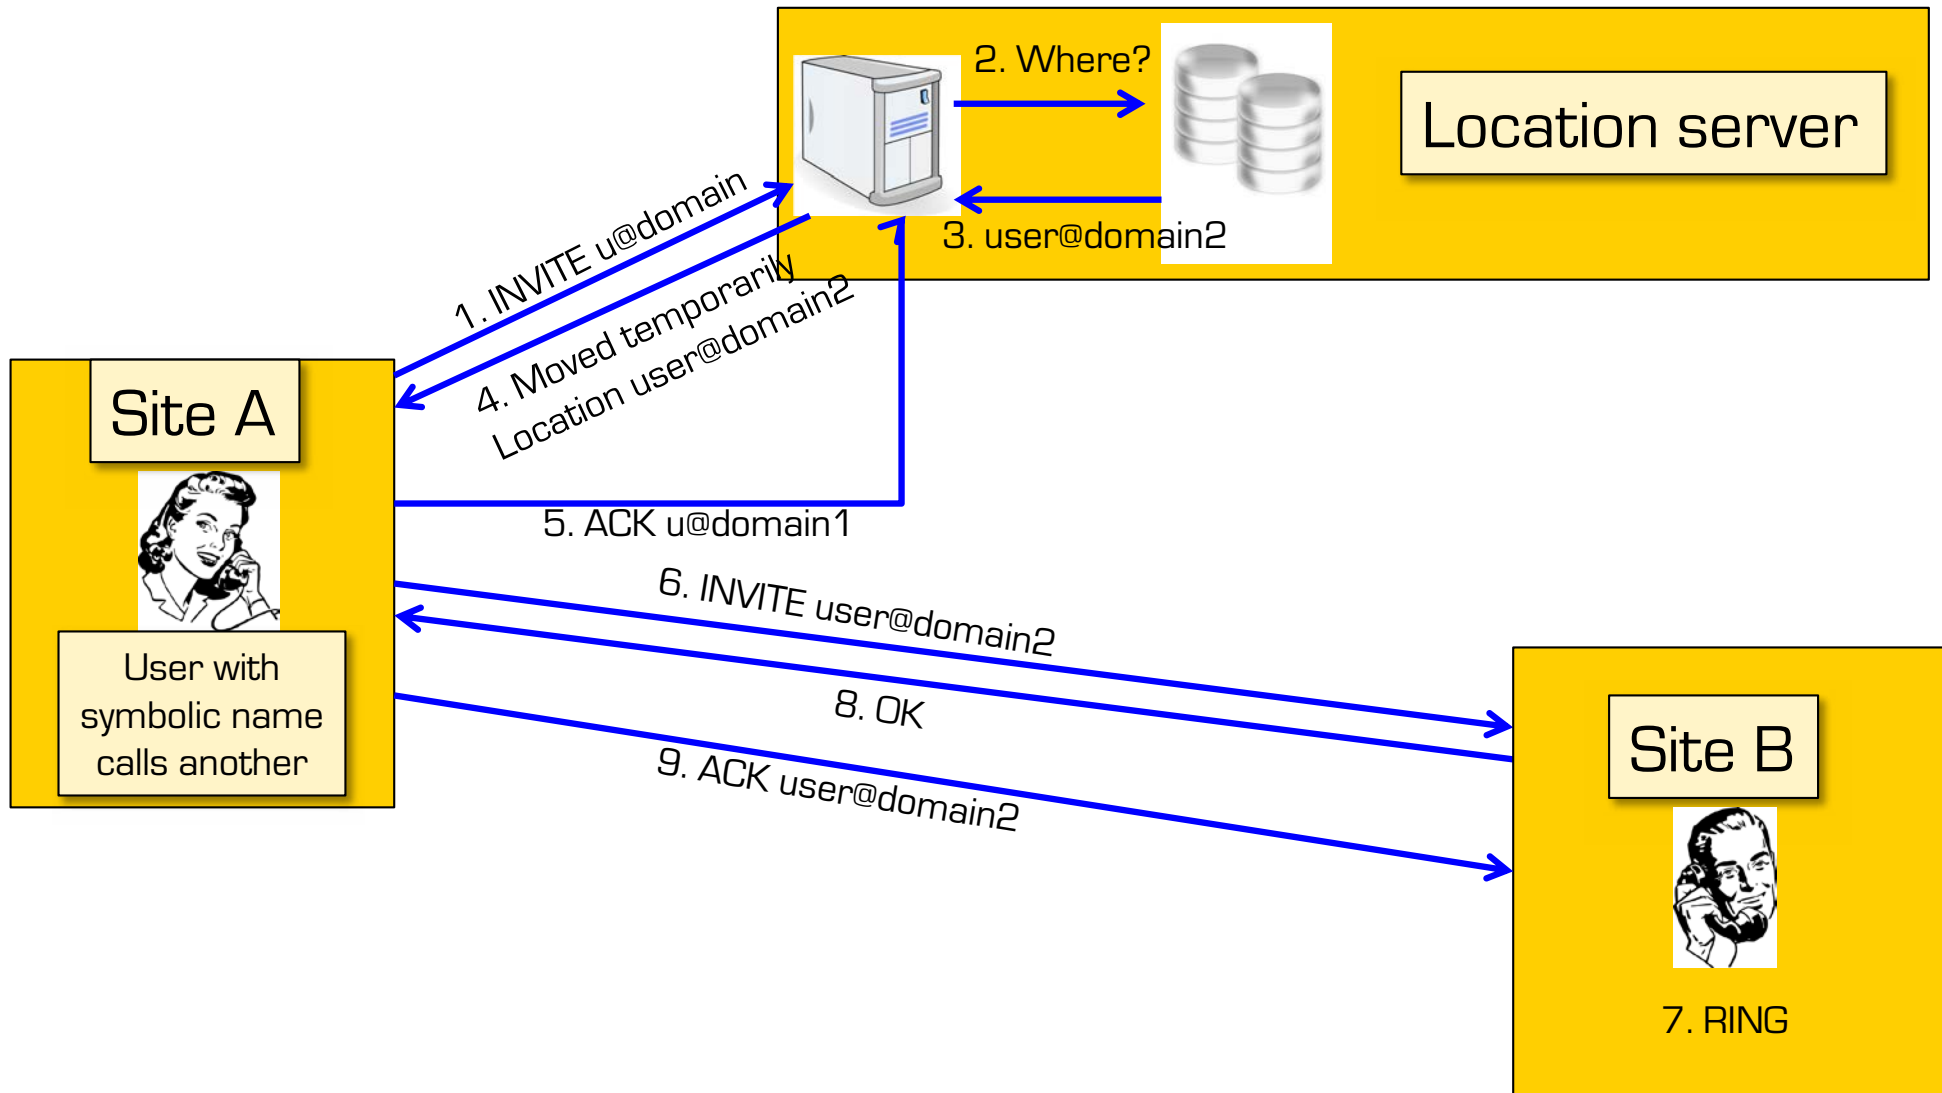
SIP Operation – Proxy Mode



Proxy forwards requests

- possibly in parallel to several hosts
- cannot accept or reject call
- useful to hide location of callee

SIP Operation - Redirect Mode



Summary

1. Signaling protocols: RTSP and SIP
2. Adaptation with RTP: Loss Delay Adjustment Algorithm (LDA)
3. How to adapt video quality?
4. Dynamic Adaptive Streaming over HTTP (DASH) and related techniques
5. Distribution Architectures
 - using the Zipf distribution
6. P2P
 - DHTs
 - Chord

