

 Chair for Network Architectures and Services – Prof. Carle  
Department for Computer Science  
TU München

**Master Course  
Computer Networks  
IN2097**

Prof. Dr.-Ing. Georg Carle  
Christian Grothoff, Ph.D.

Chair for Network Architectures and Services  
Institut für Informatik  
Technische Universität München  
<http://www.net.in.tum.de>


 Technische Universität München

 **Chapter 2: Application Layer**

Our goals:


- conceptual, implementation aspects of network application protocols
  - transport-layer service models
  - client-server paradigm
  - peer-to-peer paradigm
- learn about protocols by examining popular application-level protocols
  - HTTP
  - DNS
- programming network applications
  - socket API

IN2097 - Master Course Computer Networks, WS 2009/2010 3

 **Chapter 2: Application layer**

- Principles of network applications
- Web and HTTP
- DNS
- P2P applications
- Socket programming with TCP
- Socket programming with UDP

IN2097 - Master Course Computer Networks, WS 2009/2010 2

 **Some network applications**

- e-mail
- web
- instant messaging
- remote login
- P2P file sharing
- multi-user network games
- streaming stored video clips
- voice over IP
- real-time video conferencing
- grid computing

IN2097 - Master Course Computer Networks, WS 2009/2010 4

## Creating a network application

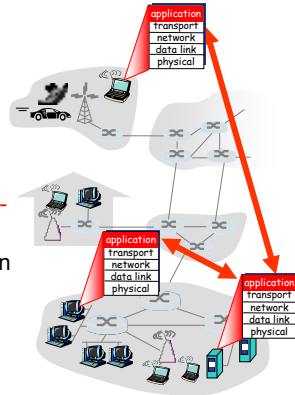
### write programs that

- run on (different) *end systems*
- communicate over network
- e.g., web server software communicates with browser software

### No need to write software for network-core devices

- Network-core devices do not run user applications
- applications on end systems allows for rapid application development, propagation

⇒ think of different viewpoint:  
what would be the benefits if you could program your router?



## Application architectures

- Client-server
- Peer-to-peer (P2P)
- Hybrid of client-server and P2P

## Chapter 2: Application layer

### Principles of network applications

- Web and HTTP
- DNS
- P2P applications
- Socket programming with TCP
- Socket programming with UDP

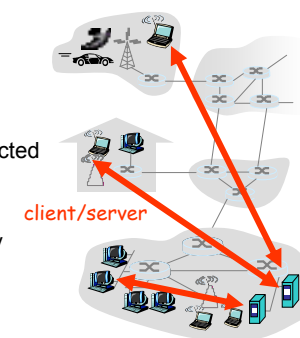
## Client-server architecture

### server:

- always-on host
- permanent IP address
- server farms for scaling

### clients:

- communicate with server
- may be intermittently connected
- may have dynamic IP addresses
- do not communicate directly with each other



### Pure P2P architecture

- no always-on server
- arbitrary end systems directly communicate
- peers are intermittently connected and change IP addresses

Highly scalable but difficult to manage

IN2097 - Master Course Computer Networks, WS 2009/2010 9

### Processes communicating

**Process:** program running within a host.

- within same host, two processes communicate using **inter-process communication** (defined by OS).
- processes in different hosts communicate by exchanging **messages**

**Client process:** process that initiates communication

**Server process:** process that waits to be contacted

- Note: applications with P2P architectures have client processes & server processes

IN2097 - Master Course Computer Networks, WS 2009/2010 11

### Hybrid of client-server and P2P

**Skype**

- voice-over-IP P2P application
- centralized server: authenticates user, finds address of remote party
- client-client connection: direct (not through server)

**Instant messaging**

- chatting between two users is P2P
- centralized service: client presence detection/location
  - user registers its IP address with central server when it comes online
  - user contacts central server to find IP addresses of buddies

IN2097 - Master Course Computer Networks, WS 2009/2010 10

### Sockets

- process sends/receives messages to/from its **socket**
- socket analogous to door
  - sending process shoves message out door
  - sending process relies on transport infrastructure on other side of door which brings message to socket at receiving process

- API: (1) choice of transport protocol; (2) ability to fix a few parameters (lots more on this later)

IN2097 - Master Course Computer Networks, WS 2009/2010 12



## Addressing processes

- ❑ to receive messages, process must have *identifier*
- ❑ host device has unique 32-bit IP address
- ❑ **Q:** does IP address of host suffice for identifying the process?



## Application-layer protocol defines

- ❑ Types of messages exchanged,
  - e.g., request, response
- ❑ Message syntax:
  - what fields in messages & how fields are delineated
- ❑ Message semantics
  - meaning of information in fields
- ❑ Rules for when and how processes send & respond to messages

### Public-domain protocols:

- ❑ defined in RFCs
- ❑ allows for interoperability
- ❑ e.g., HTTP, SMTP

### Proprietary protocols:

- ❑ e.g., Skype



## Addressing processes

- ❑ to receive messages, process must have *identifier*
- ❑ host device has unique 32-bit IP address
- ❑ **Q:** does IP address of host on which process runs suffice for identifying the process?
  - **A:** No, *many* processes can be running on same host
- ❑ *identifier* includes both **IP address** and **port numbers** associated with process on host.
- ❑ Example port numbers:
  - HTTP server: 80
  - Mail server: 25
- ❑ to send HTTP message to gaia.cs.umass.edu web server:
  - **IP address:** 128.119.245.12
  - **Port number:** 80
- ❑ more shortly...



## What transport service does an application need?

### Data loss

- ❑ some apps (e.g., audio) can tolerate some loss
- ❑ other apps (e.g., file transfer, telnet) require 100% reliable data transfer

### Timing

- ❑ some apps (e.g., Internet telephony, interactive games) require low delay to be "effective"
- ❑ frequently the applications also need timestamps (e.g. specifying playout time)

### Throughput

- ❑ some apps (e.g., multimedia) require minimum amount of throughput to be "effective"
- ❑ other apps ("elastic apps") make use of whatever throughput they get

### Security

- ❑ Some apps (e.g. Internet banking) require security services such as encryption, data integrity, ...

### Transport service requirements of common apps

Application	Data loss	Throughput	Time Sensitive
file transfer	no loss	elastic	no
e-mail	no loss	elastic	no
Web documents	no loss	elastic	no
real-time audio/video	loss-tolerant	audio: 5kbps-1Mbps video: 10kbps-5Mbps	yes, 100's msec
stored audio/video	loss-tolerant	same as above	yes, few secs
interactive games	loss-tolerant	few kbps up	yes, 100's msec
instant messaging	no loss	elastic	yes and no

IN2097 - Master Course Computer Networks, WS 2009/2010 17

### Internet apps: application, transport protocols

Application	Application layer protocol	Underlying transport protocol
e-mail	SMTP [RFC 2821]	TCP
remote terminal access	Telnet [RFC 854]	TCP
Web	HTTP [RFC 2616]	TCP
file transfer	FTP [RFC 959]	TCP
streaming multimedia	HTTP (e.g., Youtube), RTP [RFC 1889]	TCP or UDP
Internet telephony	SIP, RTP, proprietary (e.g., Skype)	typically UDP

IN2097 - Master Course Computer Networks, WS 2009/2010 19

- ### Internet transport protocols services
- TCP service:**
- *connection-oriented*: setup required between client and server processes
  - *reliable transport* between sending and receiving process
  - *flow control*: sender won't overwhelm receiver
  - *congestion control*: sender throttled when network overloaded
  - *does not provide*: timing, minimum throughput guarantees, security
- UDP service:**
- unreliable data transfer between sending and receiving process
  - does not provide: connection setup, reliability, flow control, congestion control, timing, throughput guarantee, or security
- Q:** why bother? Why is there a UDP?
- IN2097 - Master Course Computer Networks, WS 2009/2010 18

- ### Chapter 2: Application layer
- Principles of network applications
  - **Web and HTTP**
  - DNS
  - P2P applications
  - Socket programming with TCP
  - Socket programming with UDP
- IN2097 - Master Course Computer Networks, WS 2009/2010 20

## HTTP overview

**HTTP: hypertext transfer protocol**

- Web's application layer protocol
- client/server model
  - client:** browser that requests, receives, "displays" Web objects
  - server:** Web server sends objects in response to requests

PC running Explorer

Mac running Navigator

Server running Apache Web server

HTTP request

HTTP response

HTTP request

HTTP response

IN2097 - Master Course Computer Networks, WS 2009/2010 21

## HTTP connections

**Nonpersistent HTTP (v1.0)**

- At most one object is sent over a TCP connection.

**Persistent HTTP (v1.1)**

- Multiple objects can be sent over single TCP connection between client and server.

IN2097 - Master Course Computer Networks, WS 2009/2010 23

## HTTP overview (continued)

**HTTP uses TCP:**

- client initiates TCP connection (creates socket) to server, port 80
- server accepts TCP connection from client
- HTTP messages (application-layer protocol messages) exchanged between browser (HTTP client) and Web server (HTTP server)
- http1.0: TCP connection closed

**HTTP is "stateless"**

- server maintains no information about past client requests

aside

**Protocols that maintain "state" are complex!**

- past history (state) must be maintained
- if server/client crashes, their views of "state" may be inconsistent, must be reconciled

⇒ research by PhD candidate Andreas Klenk: stateless negotiation protocol suitable for Web services

IN2097 - Master Course Computer Networks, WS 2009/2010 22

## Nonpersistent HTTP

Suppose user enters URL  
`www.someSchool.edu/someDepartment/home.index`  
 (contains text, references to 10 jpeg images)

1a. HTTP client initiates TCP connection to HTTP server (process) at `www.someSchool.edu` on port 80

1b. HTTP server at host `www.someSchool.edu` waiting for TCP connection at port 80. "accepts" connection, notifying client

2. HTTP client sends HTTP *request message* (containing URL) into TCP connection socket. Message indicates that client wants object `someDepartment/home.index`

3. HTTP server receives request message, forms *response message* containing requested object, and sends message into its socket

time ↓

IN2097 - Master Course Computer Networks, WS 2009/2010 24

### Nonpersistent HTTP (cont.)

5. HTTP client receives response message containing html file, displays html. Parsing html file, finds 10 referenced jpeg objects

4. HTTP server closes TCP connection.

6. Steps 1-5 repeated for each of 10 jpeg objects

time

IN2097 - Master Course Computer Networks, WS 2009/2010 25

### Persistent HTTP

Nonpersistent HTTP issues:

- requires 2 RTTs per object
- Operating System overhead for *each* TCP connection
- browsers often open parallel TCP connections to fetch referenced objects

Persistent HTTP

- server leaves connection open after sending response
- subsequent HTTP messages between same client/server sent over open connection
- client sends requests as soon as it encounters a referenced object
- as little as one RTT for all the referenced objects

IN2097 - Master Course Computer Networks, WS 2009/2010 27

### Non-Persistent HTTP: Response time

**Definition of RTT:** time for a small packet to travel from client to server and back.

**Response time:**

- one RTT to initiate TCP connection
- one RTT for HTTP request and first few bytes of HTTP response to return
- file transmission time

**total = 2RTT + transmit time**

initiate TCP connection

RTT

request file

RTT

file received

time

time

time to transmit file

IN2097 - Master Course Computer Networks, WS 2009/2010 26

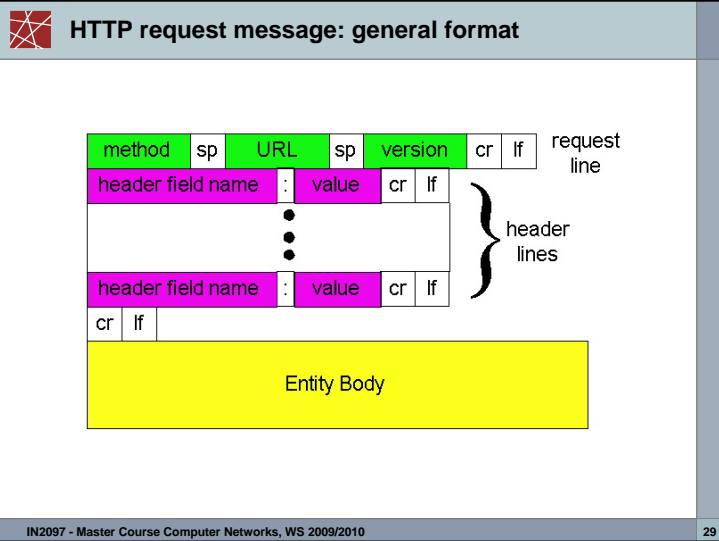
### HTTP request message

- two types of HTTP messages: *request, response*
- HTTP request message:**
  - ASCII (human-readable format)

```

request line (GET, POST, HEAD commands) → GET /somedir/page.html HTTP/1.1
header lines → Host: www.someschool.edu
User-agent: Mozilla/4.0
Connection: close
Accept-language: fr
Carriage return, line feed (extra carriage return, line feed) →
indicates end of message
  
```

IN2097 - Master Course Computer Networks, WS 2009/2010 28



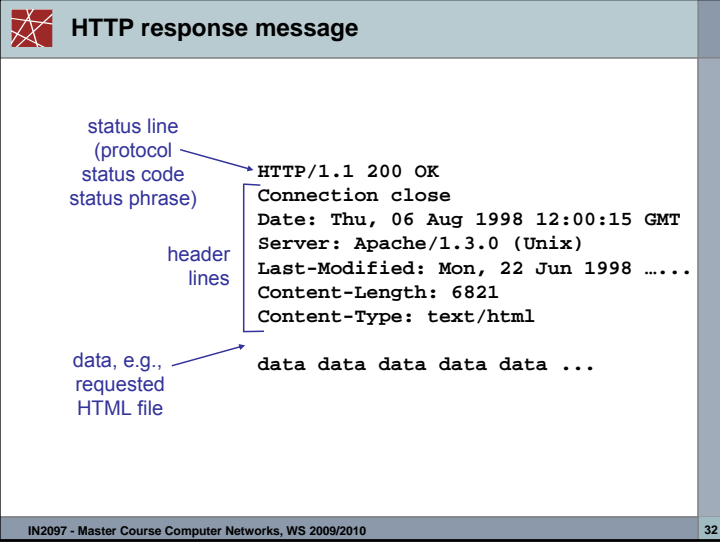
- ### Method types
- HTTP/1.0

  - GET
  - POST
    - asks server to leave requested object out of response
  - HEAD

HTTP/1.1

  - GET, POST, HEAD
  - PUT
    - uploads file in entity body to path specified in URL field
  - DELETE
    - deletes file specified in the URL field

- ### Uploading form input
- Post method:
- Web page often includes form input
  - Input is uploaded to server in entity body
- URL method:
- Uses GET method
  - Input is uploaded in URL field of request line:
- `www.somesite.com/animalsearch?monkeys&banana`





## HTTP response status codes

- In first line in server->client response message.
- A few sample codes:

### 200 OK

- request succeeded, requested object later in this message

### 301 Moved Permanently

- requested object moved, new location specified later in this message (Location:)

### 400 Bad Request

- request message not understood by server

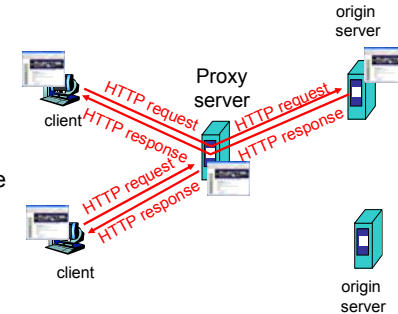
### 404 Not Found

- requested document not found on this server

### 505 HTTP Version Not Supported

## Web caches (proxy server)

- **Goal:** satisfy client request without involving origin server
- non-transparent web cache:  
user sets browser: Web accesses via cache
- browser sends all HTTP requests to cache
  - object in cache: cache returns object
  - else cache requests object from origin server, then returns object to client



## Trying out HTTP (client side) for yourself

1. Telnet to your favorite Web server:

```
telnet cis.poly.edu 80
```

Opens TCP connection to port 80 (default HTTP server port) at cis.poly.edu. Anything typed in sent to port 80 at cis.poly.edu

2. Type in a GET HTTP request:

```
GET /~ross/ HTTP/1.1
Host: cis.poly.edu
```

By typing this in (hit carriage return twice), you send this minimal (but complete) GET request to HTTP server

3. Look at response message sent by HTTP server!

## More about Web caching

- cache acts as both client and server
  - typically cache is installed by ISP (university, company, residential ISP)
- Why Web caching?**
- reduce response time for client request (Q.: under which condition is this statement true?)
  - reduce traffic on an institution's access link.
  - Internet dense with caches: enables "poor" content providers to effectively deliver content (but so does P2P file sharing)





## Chapter 2: Application layer

- Principles of network applications
- Web and HTTP
- **DNS**
- P2P applications
- Socket programming with TCP
- Socket programming with UDP



## Jon Postel

- Jon Postel (1943 – 1998)
  - Editor of RFC series
  - co-developer many Internet standards such as TCP/IP, SMTP, and DNS
  - Internet Assigned Numbers Authority (IANA)
  - "Be liberal in what you accept, and conservative in what you send."
  - obituary published in RFC 2468
- Postel Center at Information Sciences Institute, <http://www.postel.org/>
- Joe Touch  
Postel Center Director, USC/ISI  
Research Associate Professor, USC Dept. of Computer Science



## Paul Mockapetris

- „Father“ of DNS
- Did design DNS in 1983, while working at Information Sciences Institute (ISI) of University of Southern California (USC)
- DNS Architecture: RFCs 882, 883
- Obsoleted by RFCs 1034, 1035
- Company Nominum



## DNS: Domain Name System

**People:** many identifiers:

- Social Security Number, name, passport #

**Internet hosts, routers:**

- IP address (32 bit) - used for addressing datagrams
- "name", e.g., [www.yahoo.com](http://www.yahoo.com) - used by humans

**Q:** map between IP addresses and name ?

**Domain Name System:**

- *distributed database* implemented in hierarchy of many *name servers*
- *application-layer protocol* host, routers, name servers to communicate to *resolve* names (address/name translation)
  - note: core Internet function, implemented as application-layer protocol
  - complexity at network's "edge"

## DNS

Why not centralize DNS?

- ❑ single point of failure
- ❑ traffic volume
- ❑ distant centralized database
- ❑ maintenance

doesn't *scale!*

DNS services

- ❑ hostname to IP address translation
- ❑ host aliasing
  - Canonical, alias names
- ❑ mail server aliasing
- ❑ load distribution
  - replicated Web servers: set of IP addresses for one canonical name

IN2097 - Master Course Computer Networks, WS 2009/2010 45

## DNS: Root name servers

- ❑ contacted by local name server that can not resolve name
- ❑ root name server:
  - contacts authoritative name server if name mapping not known
  - gets mapping
  - returns mapping to local name server

13 root name servers worldwide

IN2097 - Master Course Computer Networks, WS 2009/2010 47

## Distributed, Hierarchical Database

```

graph TD
    Root[Root DNS Servers] --> com[com DNS servers]
    Root --> org[org DNS servers]
    Root --> edu[edu DNS servers]
    com --> yahoo[yahoo.com DNS servers]
    com --> amazon[amazon.com DNS servers]
    org --> pbs[pbs.org DNS servers]
    edu --> poly[poly.edu DNS servers]
    edu --> umass[umass.edu DNS servers]
    
```

Client wants IP for [www.amazon.com](http://www.amazon.com); 1<sup>st</sup> approx:

- ❑ client queries a root server to find com DNS server
- ❑ client queries com DNS server to get amazon.com DNS server
- ❑ client queries amazon.com DNS server to get IP address for [www.amazon.com](http://www.amazon.com)

IN2097 - Master Course Computer Networks, WS 2009/2010 46

## DNS root servers

- ❑ Only 13 physical servers?
  - No, there are 13 operators of a redundant set of DNS root servers
  - nine of the servers operate in multiple geographical locations using anycast routing (→ later), for better performance and more fault toleranc

IN2097 - Master Course Computer Networks, WS 2009/2010 48

## TLD and Authoritative Servers

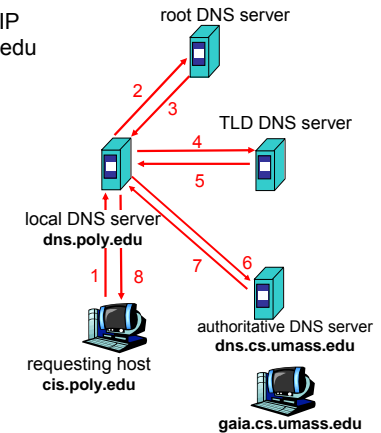
- **Top-level domain (TLD) servers:**
  - responsible for com, org, net, edu, etc, and all top-level country domains de, uk, fr, ca, jp...
  - the company Network Solutions maintains servers for com TLD
  - the company Educause for edu TLD
- **Authoritative DNS servers:**
  - organization's DNS servers, providing authoritative hostname to IP mappings for organization's servers (e.g., Web, mail).
  - can be maintained by organization or service provider

## DNS name resolution example

- Host at cis.poly.edu wants IP address for gaia.cs.umass.edu

### iterated query:

- contacted server replies with name of server to contact
- "I don't know this name, but ask this server"



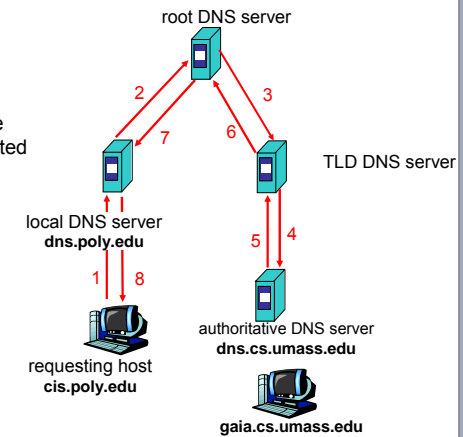
## Local Name Server

- does not strictly belong to hierarchy
- each ISP (residential ISP, company, university) has one.
  - also called "default name server"
- when host makes DNS query, query is sent to its local DNS server
  - acts as proxy, forwards query into hierarchy

## DNS name resolution example

### recursive query:

- puts burden of name resolution on contacted name server
- heavy load?





## DNS: caching and updating records

- once (any) name server learns mapping, it *caches* mapping
  - cache entries timeout (disappear) after some time
  - TLD servers typically cached in local name servers
    - Thus root name servers not often visited
- update/notify mechanisms designed by IETF
  - RFC 2136

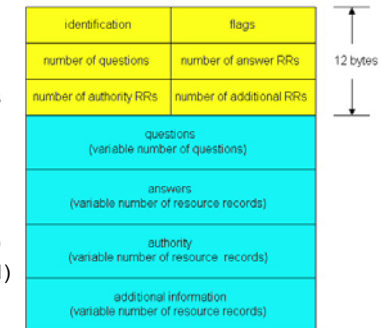


## DNS protocol, messages

DNS protocol: *query* and *reply* messages, both with same *message format*

### msg header

- **identification**: 16 bit # for query, reply to query uses same #
- **flags**:
  - query (0) or reply (1)
  - recursion desired (1)
  - recursion available (1)
  - reply is authoritative (1)



## DNS records

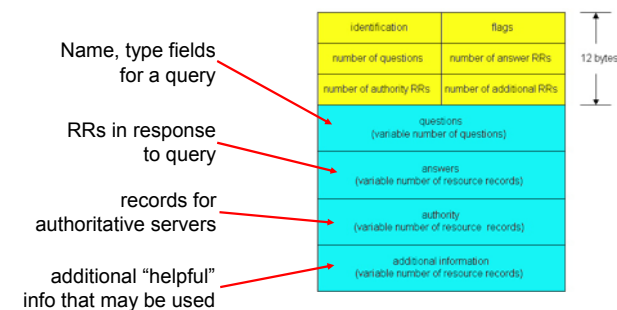
DNS: distributed database storing resource records (**RR**)

RR format: (name, value, type, ttl)

- Type=A
  - name is hostname
  - value is IP address
- Type=NS
  - **name** is domain (e.g. foo.com)
  - **value** is hostname of authoritative name server for this domain
- Type=CNAME
  - name is alias name for some "canonical" (the real) name
  - www.ibm.com is really servereast.backup2.ibm.com
  - value is canonical name
- Type=MX
  - value is name of mailserver associated with name



## DNS protocol, messages



## Inserting records into DNS

- example: new startup "Network Utopia"
- register name networkutopia.com at *DNS registrar* (e.g., Network Solutions)
  - provide names, IP addresses of authoritative name server (primary and secondary)
  - registrar inserts two RRs into com TLD server:  

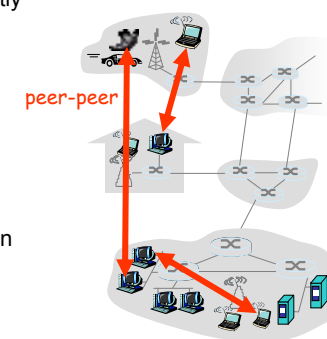
```
(networkutopia.com, dns1.networkutopia.com, NS)
(dns1.networkutopia.com, 212.212.212.1, A)
```
- create authoritative server Type A records for www.networkutopia.com;  
Type MX record for networkutopia.com
- **How do people get IP address of your Web site?**

IN2097 - Master Course Computer Networks, WS 2009/2010

57

## Pure P2P architecture

- *no* always-on server
- arbitrary end systems directly communicate
- peers are intermittently connected and change IP addresses
- **Three topics:**
  - File distribution
  - Searching for information
  - Case Study: Skype



IN2097 - Master Course Computer Networks, WS 2009/2010

59

## Chapter 2: Application layer

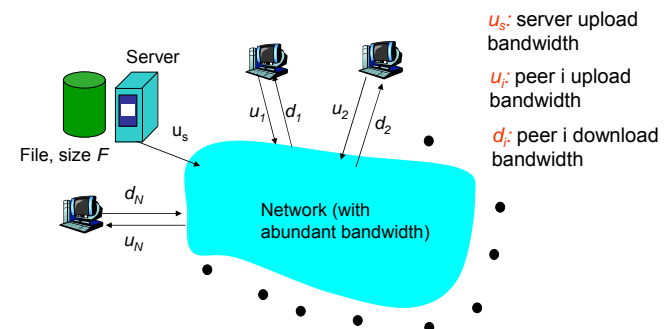
- Principles of network applications
- Web and HTTP
- DNS
- **P2P applications**
- Socket programming with TCP
- Socket programming with UDP

IN2097 - Master Course Computer Networks, WS 2009/2010

58

## File Distribution: Server-Client vs P2P

**Question:** How much time to distribute file from one server to  $N$  peers?



IN2097 - Master Course Computer Networks, WS 2009/2010

60

### File distribution time: server-client

- server sequentially sends N copies:
  - $NF/u_s$  time
- client i takes  $F/d_i$  time to download

Time to distribute  $F$  to  $N$  clients using =  $d_{cs} = \max \{ NF/u_s, F/\min(d_i) \}$   
client/server approach

increases linearly in  $N$  (for large  $N$ )

IN2097 - Master Course Computer Networks, WS 2009/2010 61

### Server-client vs. P2P: example

Client upload rate =  $u$ ,  $F/u = 1$  hour,  $u_s = 10u$ ,  $d_{min} \geq u_s$

N	Client-Server Time	P2P Time
0	0.0	0.0
5	0.5	0.4
10	1.0	0.5
15	1.5	0.6
20	2.0	0.65
25	2.5	0.68
30	3.0	0.7
35	3.5	0.7

IN2097 - Master Course Computer Networks, WS 2009/2010 63

### File distribution time: P2P

- server must send one copy:  $F/u_s$  time
- client i takes  $F/d_i$  time to download
- $NF$  bits must be downloaded (aggregate)  
fastest possible upload rate:  $u_s + \sum u_i$

$d_{P2P} = \max \{ F/u_s, F/\min(d_i), NF/(u_s + \sum u_i) \}$

IN2097 - Master Course Computer Networks, WS 2009/2010 62

### File distribution: BitTorrent

- P2P file distribution

**tracker:** tracks peers participating in torrent

**torrent:** group of peers exchanging chunks of a file

obtain list of peers

trading chunks

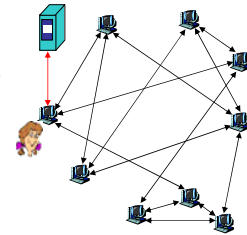
peer

IN2097 - Master Course Computer Networks, WS 2009/2010 64



## BitTorrent (1)

- file divided into 256KB *chunks*.
- peer joining torrent:
  - has no chunks, but will accumulate them over time
  - registers with tracker to get list of peers, connects to subset of peers ("neighbors")
- while downloading, peer uploads chunks to other peers.
- peers may come and go
- once peer has entire file, it may (selfishly) leave or (altruistically) remain

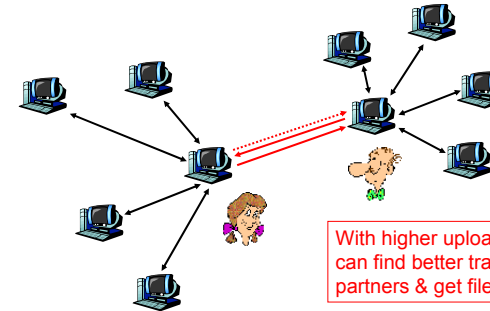


IN2097 - Master Course Computer Networks, WS 2009/2010

65

## BitTorrent: Tit-for-tat

- (1) Alice "optimistically unchokes" Bob
- (2) Alice becomes one of Bob's top-four providers; Bob reciprocates
- (3) Bob becomes one of Alice's top-four providers



IN2097 - Master Course Computer Networks, WS 2009/2010

67

## BitTorrent (2)

### Pulling Chunks

- at any given time, different peers have different subsets of file chunks
- periodically, a peer (Alice) asks each neighbor for list of chunks that they have.
- Alice sends requests for her missing chunks
  - rarest first

### Sending Chunks: tit-for-tat

- Alice sends chunks to four neighbors currently sending her chunks *at the highest rate*
  - re-evaluate top 4 every 10 secs
- every 30 secs: randomly select another peer, starts sending chunks
  - newly chosen peer may join top 4
  - "optimistically unchoke"

IN2097 - Master Course Computer Networks, WS 2009/2010

66

## Distributed Hash Table (DHT)

- DHT = distributed P2P database
- Database has (**key, value**) pairs;
  - key: ss number; value: human name
  - key: content type; value: IP address
- Peers **query** DB with key
  - DB returns values that match the key
- Peers can also **insert** (key, value) peers

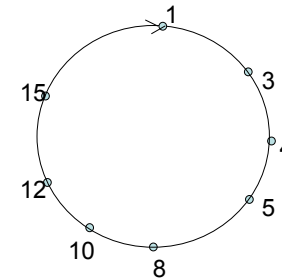
IN2097 - Master Course Computer Networks, WS 2009/2010

68

## DHT Identifiers

- Assign integer identifier to each peer in range  $[0, 2^n - 1]$ .
  - Each identifier can be represented by  $n$  bits.
- Require each key to be an integer in **same range**.
- To get integer keys, hash original key.
  - eg, key =  $h(\text{"Led Zeppelin IV"})$
  - This is why they call it a distributed "hash" table

## Circular DHT (1)



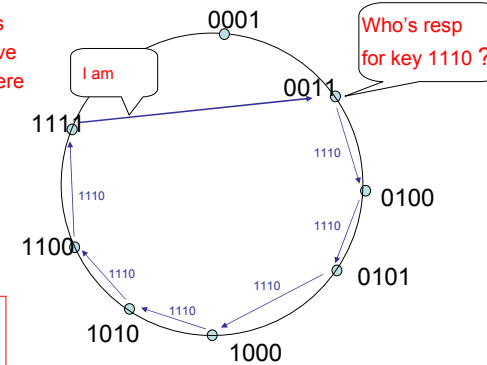
- Each peer *only* aware of immediate successor and predecessor.
- "Overlay network"

## How to assign keys to peers?

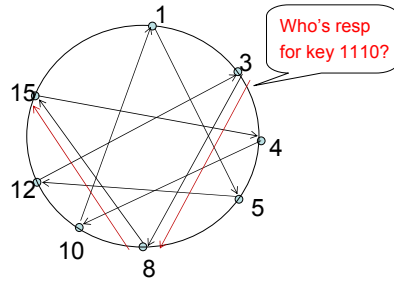
- Central issue:
  - Assigning (key, value) pairs to peers.
- Rule: assign key to the peer that has the **closest** ID.
- Convention in lecture: closest is the **immediate successor** of the key.
- Ex:  $n=4$ ; peers: 1,3,4,5,8,10,12,14;
  - key = 13, then successor peer = 14
  - key = 15, then successor peer = 1

## Circle DHT (2)

$O(N)$  messages on avg to resolve query, when there are  $N$  peers



### Circular DHT with Shortcuts



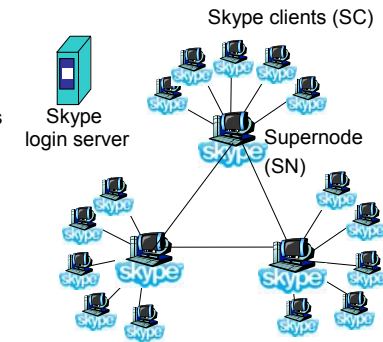
- Each peer keeps track of IP addresses of predecessor, successor, short cuts.
- Reduced from 6 to 2 messages.
- Possible to design shortcuts so  $O(\log N)$  neighbors,  $O(\log N)$  messages in query

IN2097 - Master Course Computer Networks, WS 2009/2010

73

### P2P Case study: Skype

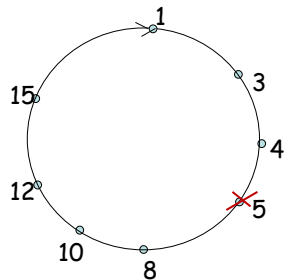
- inherently P2P: pairs of users communicate.
- proprietary application-layer protocol (inferred via reverse engineering)
- hierarchical overlay with SNs
- Index maps usernames to IP addresses; distributed over SNs



IN2097 - Master Course Computer Networks, WS 2009/2010

75

### Peer Churn



- To handle peer churn, require each peer to know the IP address of its two successors.
- Each peer periodically pings its two successors to see if they are still alive.

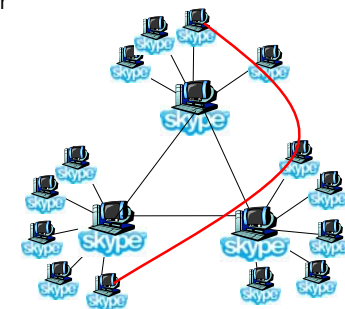
- Peer 5 abruptly leaves
- Peer 4 detects; makes 8 its immediate successor; asks 8 who its immediate successor is; makes 8's immediate successor its second successor.
- What if peer 13 wants to join?

IN2097 - Master Course Computer Networks, WS 2009/2010

74

### Peers as relays

- Problem when both Alice and Bob are behind "NATs".
  - NAT prevents an outside peer from initiating a call to insider peer
- Solution:
  - Using Alice's and Bob's SNs, Relay is chosen
  - Each peer initiates session with relay.
  - Peers can now communicate through NATs via relay



IN2097 - Master Course Computer Networks, WS 2009/2010

76



## Chapter 2: Application layer

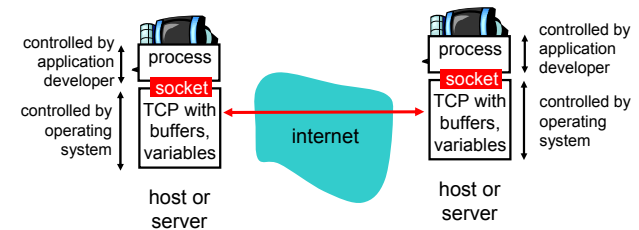
- ❑ Principles of network applications
- ❑ Web and HTTP
- ❑ DNS
- ❑ P2P applications
- ❑ **Socket programming with TCP**
- ❑ Socket programming with UDP



## Socket-programming using TCP

**Socket:** a door between application process and end-end-transport protocol (UCP or TCP)

**TCP service:** reliable transfer of **bytes** from one process to another



## Socket programming

**Goal:** learn how to build client/server application that communicate using sockets

### Socket API

- ❑ introduced in BSD4.1 UNIX, 1981
- ❑ explicitly created, used, released by apps
- ❑ client/server paradigm
- ❑ two types of transport service via socket API:
  - unreliable datagram
  - reliable, byte stream-oriented

**socket** —  
 a *host-local, application-created, OS-controlled* interface (a "door") into which application process can **both send and receive** messages to/from another application process



## Socket programming *with* TCP

### Client must contact server

- ❑ server process must first be running
- ❑ server must have created socket (door) that welcomes client's contact

### Client contacts server by:

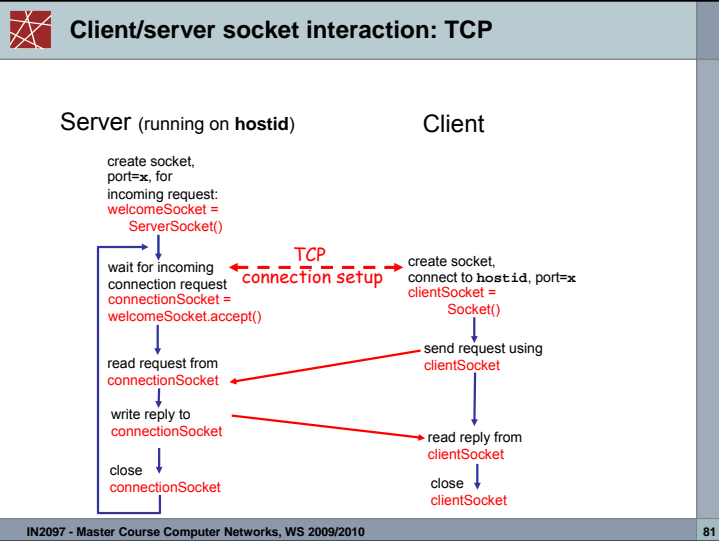
- ❑ creating client-local TCP socket
- ❑ specifying IP address, port number of server process
- ❑ When **client creates socket:** client TCP establishes connection to server TCP

- ❑ When contacted by client, **server TCP creates new socket** for server process to communicate with client

- allows server to talk with multiple clients
- source port numbers used to distinguish clients ([more in Chap 3](#))

### application viewpoint

*TCP provides reliable, in-order transfer of bytes ("pipe") between client and server*

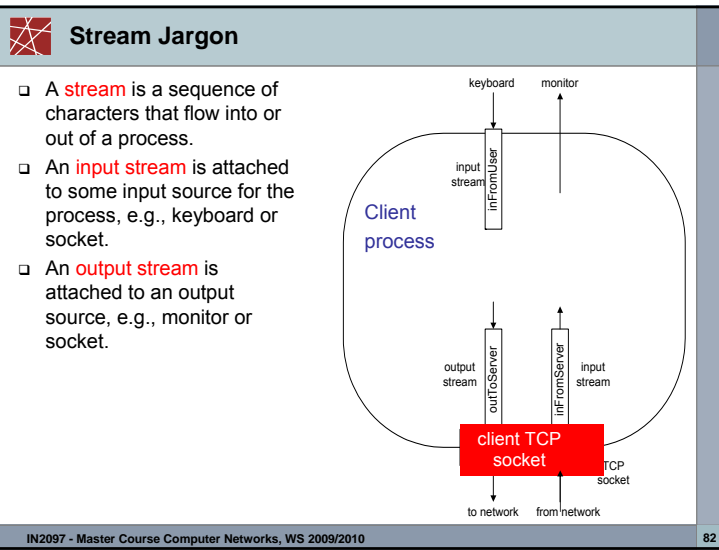


### Socket programming with TCP

**Example client-server app:**

- 1) client reads line from standard input (`inFromUser` stream) , sends to server via socket (`outToServer` stream)
- 2) server reads line from socket
- 3) server converts line to uppercase, sends back to client
- 4) client reads, prints modified line from socket (`inFromServer` stream)

IN2097 - Master Course Computer Networks, WS 2009/2010 83



### Example: Java client (TCP)

```

import java.io.*;
import java.net.*;
class TCPClient {

    public static void main(String argv[]) throws Exception
    {
        String sentence;
        String modifiedSentence;

        Create input stream → BufferedReader inFromUser =
                             new BufferedReader(new InputStreamReader(System.in));
        Create client socket, connect to server → Socket clientSocket = new Socket("hostname", 6789);
        Create output stream attached to socket → DataOutputStream outToServer =
                                                  new DataOutputStream(clientSocket.getOutputStream());
    }
}

```

IN2097 - Master Course Computer Networks, WS 2009/2010 84



## Socket programming *with* UDP

UDP: no "connection" between client and server

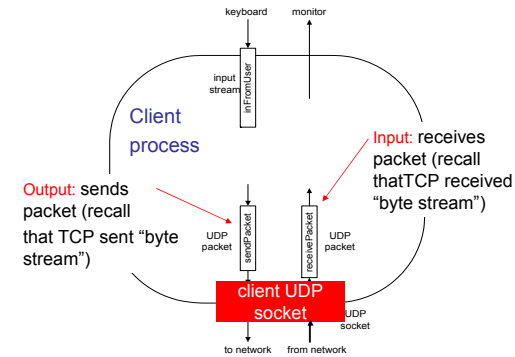
- no handshaking
- sender explicitly attaches IP address and port of destination to each packet
- server must extract IP address, port of sender from received packet

application viewpoint

UDP provides unreliable transfer of groups of bytes ("datagrams") between client and server

UDP: transmitted data may be received out of order, or lost

## Example: Java client (UDP)



## Client/server socket interaction: UDP

Server (running on `hostid`)

```
create socket,
port= x.
serverSocket =
DatagramSocket()
```

```
read datagram from
serverSocket
```

```
write reply to
serverSocket
specifying
client address,
port number
```

Client

```
create socket,
clientSocket =
DatagramSocket()
```

```
Create datagram with server IP and
port=x; send datagram via
clientSocket
```

```
read datagram from
clientSocket
close
clientSocket
```

## Example: Java client (UDP)

```
import java.io.*;
import java.net.*;

class UDPClient {
    public static void main(String args[]) throws Exception
    {
        Create input stream → BufferedReader inFromUser =
        new BufferedReader(new InputStreamReader(System.in));
        Create client socket → DatagramSocket clientSocket = new DatagramSocket();
        Translate hostname to IP address using DNS → InetAddress IPAddress = InetAddress.getByName("hostname");

        byte[] sendData = new byte[1024];
        byte[] receiveData = new byte[1024];

        String sentence = inFromUser.readLine();
        sendData = sentence.getBytes();
    }
}
```

### Example: Java client (UDP), cont.

```

Create datagram with data-to-send, length, IP addr, port } DatagramPacket sendPacket =
                                                                    new DatagramPacket(sendData, sendData.length, InetAddress, 9876);

Send datagram to server } clientSocket.send(sendPacket);

                                                                    DatagramPacket receivePacket =
                                                                    new DatagramPacket(receiveData, receiveData.length);

Read datagram from server } clientSocket.receive(receivePacket);

                                                                    String modifiedSentence =
                                                                    new String(receivePacket.getData());

                                                                    System.out.println("FROM SERVER:" + modifiedSentence);
                                                                    clientSocket.close();
                                                                    }
                                                                    }

```

IN2097 - Master Course Computer Networks, WS 2009/2010 93

### Example: Java server (UDP), cont

```

                                                                    String sentence = new String(receivePacket.getData());

Get IP addr, port #, of sender } InetAddress IPAddress = receivePacket.getAddress();
                                                                    int port = receivePacket.getPort();

                                                                    String capitalizedSentence = sentence.toUpperCase();

                                                                    sendData = capitalizedSentence.getBytes();

Create datagram to send to client } DatagramPacket sendPacket =
                                                                    new DatagramPacket(sendData, sendData.length, IPAddress,
                                                                    port);

Write out datagram to socket } serverSocket.send(sendPacket);
                                                                    }
                                                                    }

                                                                    End of while loop,
                                                                    loop back and wait for
                                                                    another datagram

```

IN2097 - Master Course Computer Networks, WS 2009/2010 95

### Example: Java server (UDP)

```

import java.io.*;
import java.net.*;

class UDPServer {
public static void main(String args[]) throws Exception
{
Create datagram socket at port 9876 } DatagramSocket serverSocket = new DatagramSocket(9876);

                                                                    byte[] receiveData = new byte[1024];
                                                                    byte[] sendData = new byte[1024];

                                                                    while(true)
                                                                    {
Create space for received datagram } DatagramPacket receivePacket =
                                                                    new DatagramPacket(receiveData, receiveData.length);

Receive datagram } serverSocket.receive(receivePacket);
                                                                    }
                                                                    }

```

IN2097 - Master Course Computer Networks, WS 2009/2010 94

### Chapter 2: Summary

our study of network apps now finished!

- application architectures
  - client-server
  - P2P
  - hybrid
- application service requirements:
  - reliability, bandwidth, delay
- Internet transport service model
  - connection-oriented, reliable: TCP
  - unreliable, datagrams: UDP
- specific protocols:
  - HTTP
  - DNS
  - P2P: BitTorrent, Skype
- socket programming

IN2097 - Master Course Computer Networks, WS 2009/2010 96





## Chapter 2: Summary

### Most importantly: learned about *protocols*

- typical request/reply message exchange:
  - client requests info or service
  - server responds with data, status code
- message formats:
  - headers: fields giving info about data
  - data: info being communicated
- *Important themes:*
- control vs. data msgs
  - in-band, out-of-band
- centralized vs. decentralized
- stateless vs. stateful
- reliable vs. unreliable msg transfer
- “complexity at network edge”