

CS 360

Application Layer

Chapter 7

SMTP / DNS / DNSSEC

BitTorrent

Regular Expressions

RFC 2396

The following line is the **regular** expression for breaking-down a URI reference into its components.

```
^((([^\/?#]+):)?(//([^/\?#]*)?([^?#]*) (\?([^\#]*)?)?(#(.*)?))?)?
```

12 3 4 5 6 7 8 9

The numbers in the second line above are only to assist readability; they indicate the reference points for each subexpression (i.e., each paired parenthesis). We refer to the value matched for subexpression <n> as \$<n>. For example, matching the above expression to

```
http://www.ics.uci.edu/pub/ietf/uri/#Related
```

results in the following subexpression matches:

```
$1 = http:
$2 = http
$3 = //www.ics.uci.edu
$4 = www.ics.uci.edu
$5 = /pub/ietf/uri/
$6 = <undefined>
$7 = <undefined>
$8 = #Related
$9 = Related
```

```
http://en.wikipedia.org/wiki/Regular_expression#POSIX_Basic_Regular_Expressions
```

For fun, look up a reg ex to validate email addresses.

Grammars

RFC 1034

`<domain> ::= <subdomain> | " "`

`<subdomain> ::= <label> | <subdomain> "." <label>`

`<label> ::= <letter> [[<ldh-str>] <let-dig>]`

`<ldh-str> ::= <let-dig-hyp> | <let-dig-hyp> <ldh-str>`

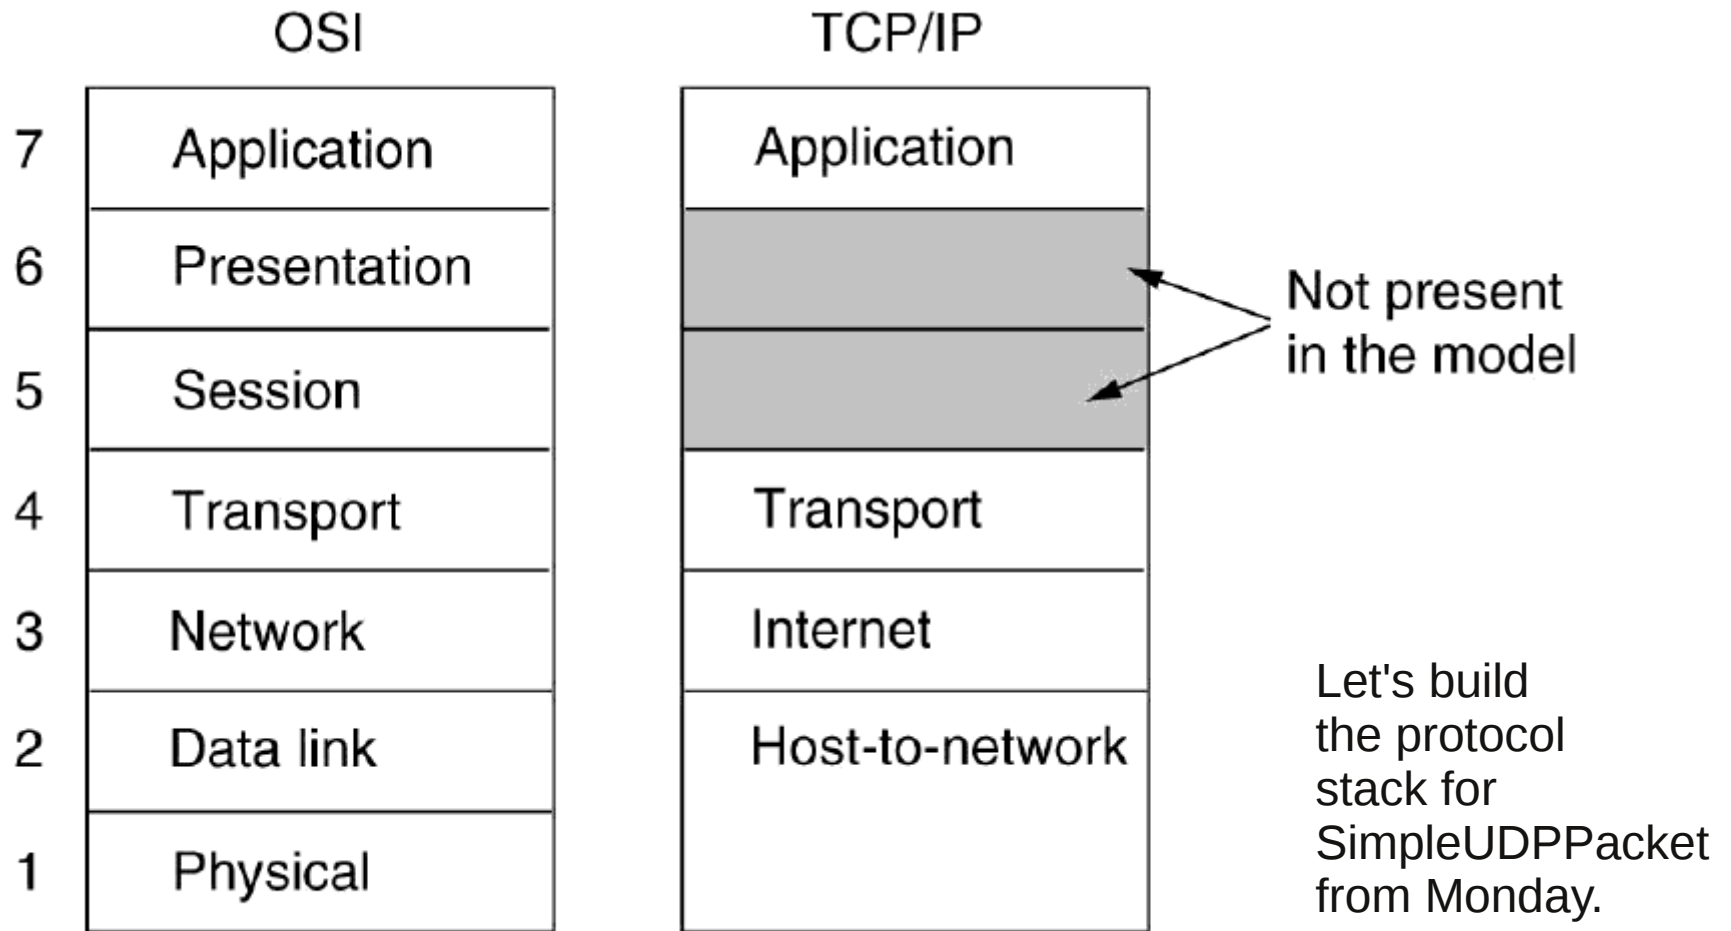
`<let-dig-hyp> ::= <let-dig> | "-"`

`<let-dig> ::= <letter> | <digit>`

`<letter> ::= any one of the 52 alphabetic characters A through Z in upper case and a through z in lower case`

`<digit> ::= any one of the ten digits 0 through 9`

Network Models (to remind us)



Computer Networks, 4th edition, Tanenbaum, page 43.
similar image on page 46 of the 5th edition.

Application Layer

- Where the meaningful work happens
 - SMTP/IMAP/POP3 (mail)
 - DNS
 - DNSSEC (section 8.9.2)
 - Streaming Media (section 7.4)
 - RTP / RTSP
 - digital encoding of media
 - Content Delivery (section 7.5)
 - P2P | BitTorrent | Chord

Client/Server Model

- A server at a **well-known** IP address listens on a **well-known** port
- A client connects, requests data, etc

Keyword	Decimal	Description	
-----	-----	-----	
echo	7/tcp	Echo	/etc/services
echo	7/udp	Echo	
daytime	13/tcp	Daytime (RFC 867)	
daytime	13/udp	Daytime (RFC 867)	
qotd	17/tcp	Quote of the Day	
qotd	17/udp	Quote of the Day	
ftp-data	20/tcp	File Transfer [Default Data]	
ftp-data	20/udp	File Transfer [Default Data]	
ftp-data	20/sctp	FTP	
ftp	21/tcp	File Transfer [Control]	
ftp	21/udp	File Transfer [Control]	
ftp	21/sctp	FTP	
ssh	22/tcp	The Secure Shell (SSH) Protocol	
ssh	22/udp	The Secure Shell (SSH) Protocol	
ssh	22/sctp	SSH	
telnet	23/tcp	Telnet	
telnet	23/udp	Telnet	
smtp	25/tcp	# Simple Mail Transfer	
smtp	25/udp	# Simple Mail Transfer	

netstat -a | less

Active Internet connections (servers and established)

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State
tcp	0	0	*:amanda	*:*	LISTEN
tcp	0	0	*:mysql	*:*	LISTEN
tcp	0	0	*:sunrpc	*:*	LISTEN
tcp	0	0	*:ssh	*:*	LISTEN
tcp	0	0	localhost:ipp	*:*	LISTEN
tcp	0	0	localhost:smtp	*:*	LISTEN
tcp	0	0	zeus.cs.pacificu.:54632	ada.cs.pacificu.e:ldaps	ESTABLISHED
tcp	0	0	zeus.cs.pacificu.ed:ssh	64.59.233.248:45566	ESTABLISHED
tcp	0	0	*:https	*:*	LISTEN
tcp	0	0	*:sunrpc	*:*	LISTEN
tcp	0	0	*:www-http	*:*	LISTEN
tcp	0	0	*:ssh	*:*	LISTEN

List of open "files"

lsuf

Which process is using which file/socket?

need to be root

fuser -v -n tcp 80 # verbose, tcp, port 80

Transport: Connected/Connectionless

- Connectionless:
 - series of unrelated packets
- Connected
 - stream of data

Telnet

- Very basic, TCP application
- Connect to an address and port and type away!
 - just echos to the screen the data it receives
- Insecure
 - sends data and passwords in *clear text*
- Zeus & Lab machines
 - no telnet servers are running!
 - now we use **ssh**!

Great for testing out your server (if your protocol is ASCII text)

Telnet to the Web server

```
Address Port
chadd@coffee:~> telnet zeus.cs.pacificu.edu 80
```

```
Trying 64.59.233.197...
```

```
Connected to zeus.cs.pacificu.edu.
```

```
Escape character is '^]'.
```

```
GET /chadd/index.html HTTP/1.1
```

```
Host: zeus.cs.pacificu.edu
```

```
<blank line, just [CRLF]>
```

```
HTTP/1.1 200 OK
```

```
Date: Fri, 03 Feb 2012 21:21:37 GMT
```

```
Server: Apache/2.2.21 (Linux/SUSE)
```

```
Last-Modified: Wed, 01 Feb 2012 19:43:57 GMT
```

```
ETag: "8c100b-3a16-4b7ec4cf20793"
```

```
Accept-Ranges: bytes
```

```
Content-Length: 14870
```

```
Content-Type: text/html
```

```
<?xml version="1.0" encoding="UTF-8"?>
```

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
```

```
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
```

```
<html xmlns="http://www.w3.org/1999/xhtml">
```

E-Mail

- RFC 822 – **ASCII** Email messages
 - <http://tools.ietf.org/html/rfc822>
 - RFC 2822
- Protocols
 - SMTP: Simple Mail Transport Protocol (RFC 821, RFC 1123)
 - POP3: Post Office Protocol (RFC 1939)
 - IMAP: Internet Mail/Message Access Protocol (RFC 1064)
 - why so many?
- User Agent (mail reader)
- Transfer Agent

SMTP

- Creating a new message

```
zeus$ telnet smtp.mailexample.net 25
220 smtp.mailexample.net ESMTP qpsmtpd 0.33-dev ready; send us your mail, but not your
spam.
HELO cs360.com
250 mailexample.net says hello to cs360.com
MAIL FROM: <professor@cs360.com>
250 sender ok
RCPT TO: student1138@cs360.com
250 recipient ok
DATA
354 Send mail; end with "." on a line by itself
From: professor@cs360.com
To: student1138@cs360.com
Subject: Cheap Stuff!
Hello! Would you like to buy something?
.
250 Message accepted
QUIT
221 mailexample.net closing connection
```

- Why does this promote spam?
 - what is spam?

what is an open relay?

POP3

- Retrieving messages

```
zeus$ telnet pop3.mailexample.net 110
+OK POP3 server ready
USER chadd
+OK
PASS iwantmail
+OK login successful
LIST
1 2505
2 14302
.
RETR 1
(send message 1)
DELE 1
QUIT
+OK POP3 server disconnecting
```

- POP3 may use *plaintext* passwords
- TLS or SSL *could* be used to encrypt the session

RFC 822 Email Message Syntax

B.1. SYNTAX

message = *field *(CRLF *text)

field = field-name ":" [field-body] CRLF

field-name = 1*<any CHAR, excluding CTLs, SPACE, and ":">

Notation:

field-body = *text [CRLF LWSP-char field-body]

1*mWORD

WORD must appear
in repetition between
1 and *m* times

1*WORD

WORD must appear
in at least once, and
may be repeated

*WORD

WORD may be repeated

MIME

base64?
octet?

- Email is ASCII
 - uuencode/uudecode (in the old days)
- Multipurpose Internet Mail Extensions (RFC 2045)
 - allows us to send non ASCII data via email
 - examples?
- No such thing as a free lunch, what does this cost us?

```
MIME-Version: 1.0
Content-type: image/jpeg
Content-Transfer-Encoding: base64
```

- where else is this used?

- where else do we send all data as ASCII?

<http://tools.ietf.org/html/rfc2045#section-6.8>

Received: from M.mailexample.net (M.mailexample.net [127.0.127.04])
by circular.mailexample.net (8.12.11.20060308/8.12.5) with ESMTP id IOUIKMEa011117;
Tue, 30 Jan 2007 13:20:22 -0500

Received: from dispatch.mailexample.net (dispatch.mailexample.net [127.0.128.60])
by M.mailexample.net (8.12.10/8.12.5) with ESMTP id IOUIKEAQ022812
for <list@M.mailexample.net>; Tue, 30 Jan 2007 13:20:14 -0500 (EST)

Received: from [127.0.130.105] (wedge.pc.mailexample.net [127.0.130.105])
by dispatch.mailexample.net (8.13.1/8.12.5) with ESMTP id IOUIKCIP006424
(**version=TLSv1/SSLv3 cipher=DHE-RSA-AES256-SHA bits=256 verify=NO**);
Tue, 30 Jan 2007 13:20:12 -0500

Message-ID: <45BF8C5C.8040108@mailexample.net>

Date: Tue, 30 Jan 2007 13:20:12 -0500

From: Da Boss <boss@mailexample.net>

User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.7.11) Gecko/20050728

MIME-Version: 1.0

To: list@mailexample.net

Subject: Comments on talk titles

Content-Type: text/plain; charset=us-ascii; format=flowed

Content-Transfer-Encoding: 7bit

X-CSD-MailScanner-Information: Please email staff@mailexample.net for more information

X-CSD-MailScanner: Found to be clean

X-CSD-MailScanner-SpamCheck: not spam, SpamAssassin (score=-1.44,required 5)

X-CSD-MailScanner-From: boss@mailexample.net

Hello! How are you? Do you like the titles of these talks?

Network Stack Review

- Message

Connectionless

- Stream

Connection-oriented

- Routing

Transport
Internet/Network

DNS

- Domain Name System (RFC 1034, 1035, 2181)
 - What is DNS?
 - When do we use it?
 - What is a domain?
 - what does this address mean: zeus.cs.pacificu.edu

DNS

- How does it work?
 - originally, just ONE file, `hosts.txt`, that was copied around to all the machines on the Internet (ARPANET) every night
 - `/etc/hosts` file still exists in UNIX
 - look here first, then queries the DNS server
 - on zeus: `cat /etc/hosts | more`
 - on Windows `system32\drivers\etc\hosts`
 - hmmm. what havoc could we wreak by writing to this file?
- Zones:
 - non-overlapping areas in the DNS
 - each zone as its own Name Server (plus a back up or two)
 - the Name Server contains the *authoritative* records for all hosts in the zone
 - not cached, always correct

DNS Root Servers

- 13 root servers spread across the globe
 - <http://d.root-servers.org/>
 - University of Maryland, College Park
 - In the basement of the Computer Science Department
 - each “root server” is really a cluster of servers



Map of the Root Servers



<http://www.icann.org/correspondence/root-map.gif>

Need an Address?

/etc/resolv.conf

- Need to find an address?
 - Use the *resolver* to look it up via a name
 - *resolver* – a network application distributed as part of an OS
 - UDP packet is sent to the local DNS nameserver
 - UDP packet is sent back with the *Resource Record*
 - why UDP?
- Resource Record
 - **Domain Name:** pacificu.edu (string)
 - **TimeToLive:** How stable is this record (int, seconds)
 - **Class:** In – Internet (string)
 - **Type:** **A** – IPv4 Address, **AAAA** - IPv6 Address, **SOA** – Authority Info, **NS** – Name Server, **MX** - mail exchange (string)
 - **Value:** Data (IP address)

DNS Protocol

<http://tools.ietf.org/html/rfc1035>

Header	
Question	the question for the name server
Answer	RRs answering the question
Authority	RRs pointing toward an authority
Additional	RRs holding additional information


The header contains the following fields:

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
ID															
QR	Opcode		AA	TC	RD	RA	Z	RCODE							
QDCOUNT															
ANCOUNT															
NSCOUNT															
ARCOUNT															

ID
QR
OpCode
AA
TC
RD
RA
Z (Updated: RFC 2065 section 6.1)
RCode
xxCount

DNS

Resource Record



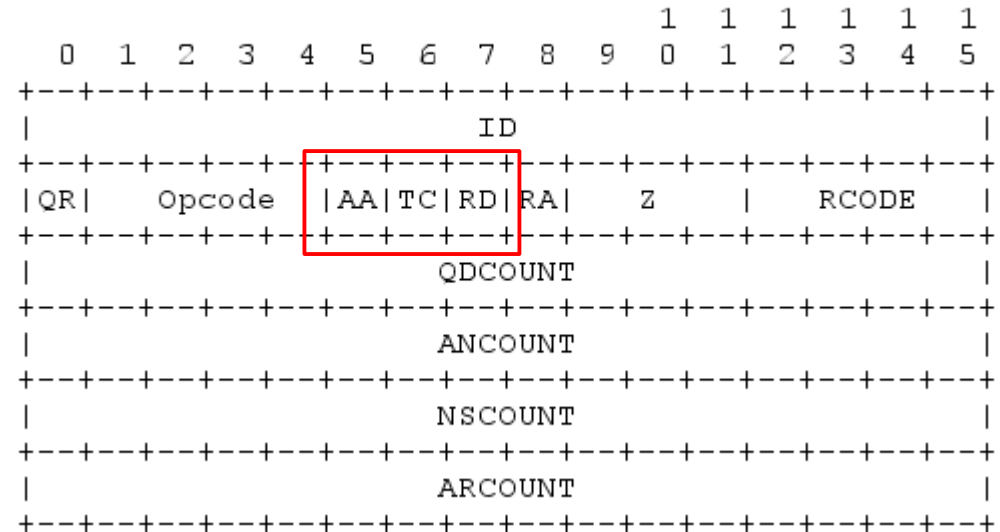
Name
Type
Class
TTL
RDLength
RData

Compression.

Bit Flipping in Java

| bit-wise OR
& bit-wise AND
>> shift right
<< shift left

The header contains the following fields:



```
byte tmpByte;  
byte RD=1, TC=0, AA=1;
```

```
tmpByte =(byte) ( (AA<<2) | (TC << 1) | RD);
```

Let's find an address

- `resolve(zeus.cs.pacificu.edu);`
- query local Name Server for the record
 - if that address is found, return it and stop.
- query (root server) Name Server for .edu TLD Servers
- query (edu server) Name Server for pacificu.edu
- query (pacificu.edu) Name Server for cs.pacificu.edu
- query (cs.pacificu.edu) Name Server for zeus.cs.pacificu.edu
 - *cache* the Resource Records retrieved since you might need them again soon
 - what problems are there with caching?

```
coffee$ nslookup
```

```
> cnn.com
```

```
Server:      64.59.233.200
```

```
Address:    64.59.233.200#53
```

```
Non-authoritative answer:
```

```
Name:  cnn.com
```

```
Address: 157.166.255.18
```

```
Name:  cnn.com
```

```
Address: 157.166.255.19
```

```
Name:  cnn.com
```

```
Address: 157.166.226.25
```

```
Name:  cnn.com
```

```
Address: 157.166.226.26
```

```
> set querytype=soa
```

```
> cnn.com
```

```
Server:      64.59.233.200
```

```
Address:    64.59.233.200#53
```

```
Non-authoritative answer:
```

```
cnn.com
```

```
origin = ns1.timewarner.net
```

```
mail addr = hostmaster.turner.com
```

```
serial = 2012020301
```

```
refresh = 28800
```

```
retry = 7200
```

```
expire = 604800
```

```
minimum = 3600
```

```
Authoritative answers can be found from:
```

```
ns1.timewarner.net    internet address = 204.74.108.238
```

DNS in action

- What is going on here?
- Why do we have so many answers?
- What is non-authoritative?
- Why is this in the Application Layer, IP does routing, right?

How do I register a domain name?

- Use a company called a *registrar*
 - for a fee, they maintain lists of available domain names
 - you provide an IP address, they provide a DNS Name
 - previously, one company did this: Network Solutions
 - now a huge number of companies do this
 - many of them provide other services (web/mail hosting, etc)
- What about those companies that let you register a domain name for your dialup/DSL connection?
 - www.dyndns.com, www.tzo.com
 - why is a dialup/DSL connection a problem?

DNS: What can go wrong?

- http://news.zdnet.com/2100-1009_22-6156944.html?tag=nl.e589
- Root DNS servers were flooded with traffic (servers: *F, I, M, G, L*)
 - early morning Tuesday (West coast time)
- How can this affect the Internet?
- What mechanisms are in place in DNS to mitigate this type of attack?
- Did you notice a problem?
 - In 2002 a similar attack shutdown 9 of the 13 root servers

Example Code

zeus.cs.pacificu.edu/home/cs360s12/SVNROOT/CS360Utils

```
/* BIG ENDIAN
 *
 * -----
 * |high| low| VALUE
 * -----
 * | 00 | 01 | = 0x1
 * -----
 *      x   x+1 ADDRESS
 *
 * The high value byte is in the LOW address and thus
 * is the first byte read/written when using a ByteBuffer.
 *
 */
```

Data is transferred on the network in Network Byte Order; Big Endian

In C, you must use `htonl()` and `ntohl()`

Your CPU [usually] defines Endianness.
x86 is Little Endian

However:

Java is ALWAYS big Endian.
Some CPUs support both.

```

/**
 * Reads two bytes from a ByteBuffer that represent a 16-bit
 * unsigned short in Network Byte Order (Big Endian) and
 * transforms that unsigned short into a 4 byte signed int.
 * @param bb The ByteBuffer to read from
 * @return the signed int representation.
 */
public static final int unsignedShortFromBB(ByteBuffer bb)
{
    int i = 0;

    // get the high value byte
    i = (bb.get() & 0xFF) << 8 ;

    // get the low value byte
    i = i | ((bb.get() & 0xFF) );

    //System.err.println("VALUE:" +i);

    return i;
}

```



```

/**
 * Writes the int value into the ByteBuffer in the format of a
 * 16-bit unsigned short in Network Byte Order (Big Endian).
 *
 * @param bb the ByteBuffer to write to
 * @param value the value to write to the ByteBuffer
 * @return the int value
 * @throws Exception if the int value to be put into the ByteArray
 * is out of range of 16 bit unsigned int
 */
public static final int unsignedShortToBB(ByteBuffer bb, int value)
    throws Exception
{
    if( value < 0 || value > 65535)
    {
        throw new Exception("Overflow");
    }

    // get the high value byte for writing
    byte b = (byte) ((value & 0xff00) >> 8);
    bb.put(b);

    // get the low value byte
    b = (byte) ((value & 0xff));
    bb.put(b);

    return value;
}

```

Peer to Peer

(7.5)

Read all of 7.4 for Monday!

- What is a peer?
- What is a client?
- High level idea?
- Challenges?

How is p2p different than client/server?

How is this similar/different to/from DNS?

Peer-to-Peer

- *Computer Networking: A Top-Down Approach Featuring the Internet, 3rd edition.* Kurose, Ross. [In my office if you want to read it.](#)
- Ethical/Legal issues:
 - <http://iptps03.cs.berkeley.edu/final-papers/copyright.pdf>
 - <http://freenetproject.org/papers/freenet-ieee.pdf>
 - Legit uses?

Napster



- Very early P2P system
 - 1999
 - shutdown by court order
- Centralized index
 - upload your list of shared data
 - receive IP address of peers sharing data you want

Pop up and share



Gnutella (0.4)

- Decentralized
- Protocol, many clients
 - LimeWire (uses Gnutella and BitTorrent)
 - morpheus, BearShare, Gtk-Gnutella....
- Bootstrap
 - Creates an *overlay network*
 - Query flooding: send a query to all your peer
 - each peer forwards on the query if they don't have the data
 - max number of hops
 - send response back through the path the query took
 - how is this good?
 - how is this bad?
 - how can we fix it?



Gnutella 0.6

- Add ultrapeers
 - each peer is connected to small number of ultrapeers
 - each ultrapeer is connected to many ultrapeers
 - high out degree
 - Lower max number of hops
 - send results directly to requester

Bit Torrent

- How is this fundamentally different?

- .torrent file?

- tracker

- swarm

- seeder

- leecher

- chunks

Reward good behavior

randomly select peers

trade chunks with peers with
best performance
unchoked

poor performing nodes will
get choked off


Chord

- Hashtable (key, data)
- DHT (Distributed Hash Table)
 - structured
 - decentralized
 - fault tolerance
 - scalability
- Hash the key to find the containing node
 - move data to correct node
 - store *data* at node *hash(key)*
 - the *data* will outlive you

Built for more persistent storage
Could build a distributed
file system
CFS

replicate data among nodes
move data as nodes join/leave

overlay network



OceanStore

- “OceanStore is a global **persistent** data store designed to scale to billions of users. ... atop an infrastructure comprised of **untrusted servers**.”
- Uses Chimera
 - an implementation of DHT
 - similar to Tapestry and Pastry