

Computer Communications - CSC 551

# CS551

## Final Project Part (2)

### Bill Cheng

*<http://merlot.usc.edu/cs551-f12>*

Copyright © William C. Cheng

## Part (2) Message Types

- Store
  - probabilistic storing of files
  - ◆ node that initiates STORE always store the file
  - ◆ use NeighborStoreProb to decide if it forwards to a particular neighbor
  - ◆ when a node gets a STORE request, use StoreProb to decide if it should cache a copy of the file
  - Search
  - Get
  - probabilistic/opportunistic caching of files
  - ◆ node that initiates GET always store the file
  - ◆ if forwarding GET response, use CacheProb to decide if it should cache a copy of the file
  - Delete

Part (2): think google and nasper (35% project grade)

Copyright © William C. Cheng

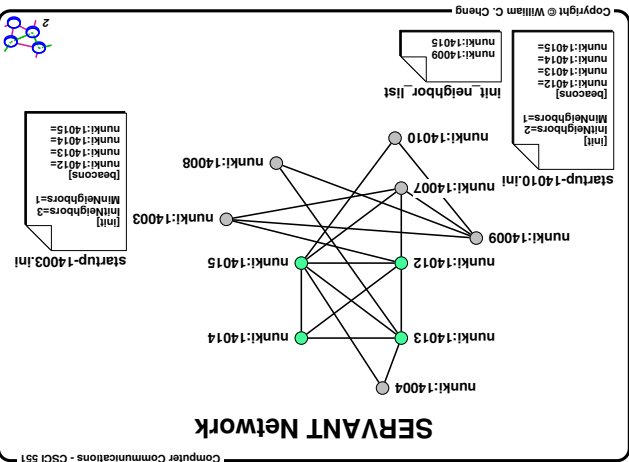
## Keywords

STORE command

```

store Blondiel.mp3 30
  category="audio mp3" \
  title="Blondie" \
  artist="Blondie" \
  title="Heart of Glass" \
  url="http://www.blondie.net/" \
  additional_keywords="debra haxy"

[metadata]
FileName=blondiel.mp3
FileSize=485526
SHA1=730764e28a5b66e3f95cead976e038d39bd89
Nonce=5cb2a2c6f224de5fcd45d601cd59f2db9d69
Keywords=categories audio mp3 artist blondie \
  title Heart of Glass \
  url http://www.blondie.net/ \
  additional_keywords debra haxy
BlkVector=
11000100000000420200000000000000000000000000000000
1000000000000000000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000000000000000000
0000210000000000000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000000000000000000
```



Copyright © William C. Cheng

## Part (2) Is Based On Part (1)

But,

- no JOIN
- every regular node will start with a good `init_neighbor_list` file
- make sure your code can parse it
- no CHECK
- do not initiate or forward CHECK messages
- the startup configuration file has `NoCheck=1`

Copyright © William C. Cheng

## Keywords (Cont...)

- Content-based addressing
- mini file system
- directory and files
- 2) directory contains description (metadata) of files no need for subdirectories
- Caching is a local behavior
- every node can have its own implementation

Copyright © William C. Cheng

## Searching (Cont...)

- GET (i.e., retrieving)
  - e.g., get 2 [<extfile>]
  - flood a GET request with a Field in the message
  - so that only one node will respond
  - you can create a Field when you create a SEARCH response message
  - keep Field in memory only
- Opportunistic caching
  - for nodes that did not initiate a GET request, cache the file
  - to increase performance (as the expense of extra storage)
  - with CacheProb
    - if CacheProb is 0.3, you should cache 30% of the time
    - call srand48 () during initialization
    - call drand48 (), if returned value < CacheProb, cache the file

Computer Communications - CSC1 551

Copyright © William C. Cheng

## Delete

- Delete a file
  - only the creator of a file can delete it
  - on file creation (i.e., STORE), generate a random password using getUID ()
  - this is a *one-time password*
  - calculate *nonce=SHA1(password)*
  - nonce* is part of *file metadata*
  - e.g., delete FileName=foo SHA1=6b6c... Nonce=fe18...
  - FileSpec is:
 

```
FileName=foo
SHA1=6b6c...
Nonce=fe18...
Password=27c3...
```
  - verifying one-time password
    - if SHA1(password) == nonce, delete the file

Computer Communications - CSC1 551

Copyright © William C. Cheng

## Bit-Vector (Cont...)

- 2 bit-vectors (n bits on the left and n bits on the right)
  - n = 512 for our project
  - concatenated into one 1024 bit string for storage in *File Metadata*, hexstring encoded
  - for a keyword k:
    - corresponding bit in left bit-vector: SHA1(k) mod n
    - corresponding bit in right bit-vector: MD5(k) mod n
  - Ex: single keyword, k = "categories"
    - echo -n "categories" | openssl sha1
    - 50b9e78177f37e3c747f67abcc8af36a44218f5
    - SHA1(k) mod n (same as taking the right-most 9 bits)
    - 0x0f5 ( = 245 in decimal)
    - echo -n "categories" | openssl md5
    - b0b5ccb4a195a07fd3eed14affb8695f
    - MD5(k) mod n = 0x15f ( = 351 in decimal)

Computer Communications - CSC1 551

Copyright © William C. Cheng

## Searching

- at commandline, think google.com but slightly different
- case-insensitive
- AND searches only
- e.g., search keywords="glass heart of" will only match a file with metadata containing *all* 3 words
- example of responses
 

```
[1] fileId=02adefc1dfc97a082fa18a5ef1e8c487259b7fb4
fileName=foo
FileSpec=I23
SHA1=b33a7b5fcbefcd3eaa547fb0f9a97b2a0ea94c
Nonce=1D7alb6fe169dae22518a865ab2e44c70fcab82
KeyWords=key1 key2 key3
FileID=45929c03a7c84687a73543cc348484edc3829496
FileName=bar
FileSpec=key1 key2
SHA1=6b6c5636c484d4759d20191c3023b8a29b2fe11
Nonce=fe1834fd8cd7356ca1e0c77ac38d387e228f94
KeyWords=key4 key5
...
```
- Searching
  - search keywords="glass heart of" will only match
  - AND searches only
  - case-insensitive
  - e.g., search keywords="glass heart of" will only match a file with metadata containing *all* 3 words
  - example of responses

Computer Communications - CSC1 551

Copyright © William C. Cheng

## Index Files

- You must implement 3 index structures to support 3 types of
  - one maps a filename to a list of file references
  - one maps a filename to a list of file references
  - one maps a SHA1 value to a list of file references
- Although the spec says that you need to use BSTs for filename and SHA1 indices, using a *sorted linear list* is fine
- When a node goes down, you need to *externalize* these index structures so that when you restart, it can recover the index structures quickly
- kwrd\_index maps a bit-vector to a list of file references
- name\_index maps a filename to a list of file references
- sha1\_index maps a SHA1 value to a list of file references

Computer Communications - CSC1 551

Copyright © William C. Cheng

## Bit-Vector

- Bit-vector as a simplest form of a *Bloom Filter*
- directory entry contains a bit-vector (long, e.g., 1024 bits)
  - map all possible words to the bit-vector
  - for example, use SHA1 mod 1024 to produce a bit index into the bit-vector
  - many words can map to the same bit index
  - take all keywords, compute bit index, set all these bits to 1
  - for a single-word query, compute bit index of this word
  - if the corresponding bit in a bit-vector is set, there is a *possible* match; in this case, do string compare
  - if the corresponding bit in a bit-vector is *not* set, there is *no possibility* of a match; try the next directory entry

Computer Communications - CSC1 551

