

Computer Communications - CSC1 551

Copyright © William C. Cheng

Peer-to-Peer File Sharing System

- Peer-to-peer: if one is using the system, one must be sharing his/her resources
- Application layer network named *SERVANT*
- the words "packet" and "message" are interchangeable
- messages are sent via flooding (no loops)
- reply to messages are sent along the path it was received
- Bi-directional connections between neighbors
- A→B and B→A must use the same connection
- Two types of nodes
 - beacon nodes*: well known addresses that every node knows, fully connected to form the core of the network
 - regular nodes
 - operationally, beacons are just like regular nodes (except a beacon does not need to *join* the network)

1

Computer Communications - CSC1 551

Copyright © William C. Cheng

Message Format

2

Computer Communications - CSC1 551

Copyright © William C. Cheng

Requests

For example, join request

- Requests are *flooded* to the entire *SERVANT* network
- anonymity of the message senders and message receivers
- must avoid loops
- nodes must *cache* a copy of any flooded messages
- (1) for *loop detection*, i.e., drop duplicate messages based on *UID*
- (2) to *route a response message* to the originator of the corresponding request message
- message cache expires after *MsgLifetime* or *GetMsgLifetime*

3

Computer Communications - CSC1 551

Copyright © William C. Cheng

CS51

Final Project Part (1)

Bill Cheng

<http://merlot.usc.edu/cs51-f12>

1

Computer Communications - CSC1 551

Copyright © William C. Cheng

Message Types

- Part (1): form and maintain the network (45% project grade)
 - Join
 - Hello
 - Keapalive
 - Notify
 - Status
 - Check (keep things connected)
- Part (2): think google and napsster (35% project grade)
 - Store
 - Search
 - Get
 - Delete

2

Computer Communications - CSC1 551

Copyright © William C. Cheng

UID

Probably/itically unique

```

char *getUID()
{
    char *obj_buf;
    int void_buf_sz;
}

static unsigned long seq_no=(unsigned long)1;
char sha1_buf[SHA_DIGEST_LENGTH], str_buf[104];

sprintf(str_buf, "%s%11d",
        node_inst_id, (long)seq_no++);
SHA1(str_buf, strlen(str_buf), sha1_buf);
memcpy(void_buf, sha1_buf,
        min(void_buf_sz, sizeof(sha1_buf)));
return void_buf;
    
```

3

Computer Communications - CSC 551

Network Formation

startup-14012.inl

```

[init]
MinNeighbors=1
[beacons]
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
  
```

nunki:14012
nunki:14013
nunki:14014
nunki:14015

beacon nodes are fully connected

- beacon node
- regular node
- neighbors
- temporary
- connection

Copyright © William C. Cheng

Computer Communications - CSC 551

Join

startup-14012.inl

```

[init]
MinNeighbors=2
[beacons]
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
  
```

nunki:14012
nunki:14013
nunki:14014
nunki:14015

join rply (distance=2781)

flooding join req

beacon nodes are fully connected

- beacon node
- regular node
- neighbors
- temporary
- connection

whole network, send join rply

send join request to a beacon

join request is flooded to the

Copyright © William C. Cheng

Computer Communications - CSC 551

Join

startup-14012.inl

```

[init]
MinNeighbors=2
[beacons]
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
  
```

nunki:14012
nunki:14013
nunki:14014
nunki:14015

init_neighbor_list

↑

nunki:14013 (557)
nunki:14014 (71)
nunki:14015 (1022)

beacon nodes are fully connected

- beacon node
- regular node
- neighbors
- temporary
- connection

send join request to a beacon

join request is flooded to the

whole network, send join rply

flooding stopped if packet

already seen

sort replies, write bottom ones to

init_neighbor_list in HomeDirectory

Copyright © William C. Cheng

Computer Communications - CSC 551

Responses

For example, join response

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Common Header															
Join Request															
OID															
Port															
Hostname															

- Do not know exactly who initiated the join (only know the OID of the join request)
- use the OID of the join request to do routing
- intermediate nodes must cache a copy of the join request message and which link it came from in order to send the join response
- join request initiator uses *JoinTimeout*

Copyright © William C. Cheng

Computer Communications - CSC 551

Join

startup-14012.inl

```

[init]
MinNeighbors=2
[beacons]
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
  
```

nunki:14012
nunki:14013
nunki:14014
nunki:14015

join req

beacon nodes are fully connected

- beacon node
- regular node
- neighbors
- temporary
- connection

Copyright © William C. Cheng

Computer Communications - CSC 551

Join

startup-14012.inl

```

[init]
MinNeighbors=2
[beacons]
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
  
```

nunki:14012
nunki:14013
nunki:14014
nunki:14015

join rply (distance=557)
join rply (distance=71)
join rply (distance=1022)

beacon nodes are fully connected

- beacon node
- regular node
- neighbors
- temporary
- connection

send join request to a beacon

whole network, send join rply

flooding stopped if packet

already seen

Copyright © William C. Cheng

Copyright © William C. Cheng

14

Legend:

- beacon node
- regular node
- neighbors
- temporary connection

Beacon nodes are fully connected

- send join request to a beacon
- join request is flooded to the whole network, send join reply
- flooding stopped if packet already seen
- sort replies, write bottom ones to `int_neighbor_list`
- restart, no need to `join`, just say hello
- say hello back

Startup logs:

```

nunki:14012 [init]
MinNeighbors=1
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14007 [init]
MinNeighbors=2
nunki:14012 [beacon]
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
int_neighbor_list
nunki:14013
nunki:14014
nunki:14015
  
```

Startup-14012.ini

Copyright © William C. Cheng

16

Node Goes Down

Startup logs:

```

nunki:14004 [init]
MinNeighbors=1
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14012 [beacon]
nunki:14009 [init]
MinNeighbors=3
nunki:14010 [init]
MinNeighbors=1
nunki:14007 [init]
MinNeighbors=2
nunki:14008 [init]
MinNeighbors=1
nunki:14010=
nunki:14011=
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
int_neighbor_list
nunki:14013
nunki:14014
nunki:14015
  
```

Startup-14003.ini

Copyright © William C. Cheng

18

Node Goes Down (Cont...)

Startup logs:

```

nunki:14004 [init]
MinNeighbors=1
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14012 [beacon]
nunki:14008 [init]
MinNeighbors=1
nunki:14010 [init]
MinNeighbors=2
nunki:14007 [init]
MinNeighbors=1
nunki:14003 [init]
MinNeighbors=3
nunki:14010=
nunki:14011=
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
int_neighbor_list
nunki:14013
nunki:14014
nunki:14015
  
```

Startup-14003.ini

Messages (even if there are nodes connected to it from below)

network is partitioned, nunki:14015 goes down

nunki:14010 will not get any check response messages

Copyright © William C. Cheng

13

Legend:

- beacon node
- regular node
- neighbors
- temporary connection

Beacon nodes are fully connected

- send join request to a beacon
- join request is flooded to the whole network, send join reply
- flooding stopped if packet already seen
- sort replies, write bottom ones to `int_neighbor_list`
- restart, no need to `join`, just say hello

Startup logs:

```

nunki:14012 [init]
MinNeighbors=2
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14007 [init]
MinNeighbors=1
nunki:14012 [beacon]
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
int_neighbor_list
nunki:14013
nunki:14014
nunki:14015
  
```

Startup-14007.ini

Copyright © William C. Cheng

15

SERVANT NETWORK

Startup logs:

```

nunki:14004 [init]
MinNeighbors=1
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14012 [beacon]
nunki:14009 [init]
MinNeighbors=3
nunki:14010 [init]
MinNeighbors=1
nunki:14007 [init]
MinNeighbors=2
nunki:14008 [init]
MinNeighbors=1
nunki:14010=
nunki:14011=
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
int_neighbor_list
nunki:14013
nunki:14014
nunki:14015
  
```

Startup-14003.ini

Copyright © William C. Cheng

17

Node Goes Down (Cont...)

Startup logs:

```

nunki:14004 [init]
MinNeighbors=1
nunki:14013 [beacon]
nunki:14014 [beacon]
nunki:14015 [beacon]
nunki:14012 [beacon]
nunki:14008 [init]
MinNeighbors=1
nunki:14010 [init]
MinNeighbors=2
nunki:14007 [init]
MinNeighbors=1
nunki:14003 [init]
MinNeighbors=3
nunki:14010=
nunki:14011=
nunki:14012=
nunki:14013=
nunki:14014=
nunki:14015=
int_neighbor_list
nunki:14013
nunki:14014
nunki:14015
  
```

Startup-14003.ini

How to Manage Timers

- Networking programming often requires you to manage many timers
 - e.g., MsgLifetime
 - you need to implement a message cache, keyed on UOID
 - drop duplicate messages
 - route response messages
 - every time you cache a message, conceptually, you should start a timer
 - when the timer expires, you can remove the message from your message cache data structure
 - need to cache a message for quite a long time
 - you can end up with thousands of timer
 - e.g., KeepAliveTimeout
 - not as many timers, but you need to keep track of them

Data Structures

- List
 - from warmup #2
 - Efficient data structure
 - Binary Search Tree (BST)
 - you can use libavl if you don't have something you already like
 - not required
 - Bloom Filter
 - required
 - for part (2), you don't have to worry about it for now

Node Startup Configuration File (Cont...)

Table driven

```

typedef struct tagKwInfo {
    int id;
    char *key;
    /* what else? */
    KwInfo;
} KwInfo;

#define KW_PORT 1001
#define KW_HOMEDIR 1002
#define KW_PERMISSIONS 1003

static KwInfo gKwInfo[] = {
    { KW_PORT, "port" },
    { KW_HOMEDIR, "homedir" },
    { KW_PERMISSIONS, "permissions" },
};

int main() {
    char *psz_key=NULL, *psz_value=NULL;
    if (getKeyValue(line, '&key', '&psz_key', &psz_value) == 0) {
        KwInfo *pwi=gKwInfo;
        while (pwi) {
            if (strcmp(pwi->key, psz_key) == 0) {
                break;
            }
            pwi=pwi->next;
        }
        if (pwi) {
            if (strcmp(pwi->value, psz_value) == 0) {
                return 0;
            }
            psz_key = pwi->key;
            psz_value = pwi->value;
        } else {
            return -1;
        }
    }
    return 0;
}
  
```

Node Startup Configuration File (Cont...)

- Don't complain it takes too much effort to parse the file!!
 - suggestion: utility file (for the rest of your grad school)
 - char *GetLine(FILE*): read an arbitrary long line
 - use malloc() and realloc()
 - void TrimBlanks(char*): get rid of leading and trailing space and tab characters
 - int GetKeyValue(char *buf, char separator, char **ppsz_key, char **ppsz_value): get key and value from an input buffer

```

if (psz_value == NULL) return ERR_CANNOT_FIND_SEPARATOR;
psz_value += strlen(buf, separator);
TrimBlanks(psz_value);
if (psz_key != NULL) *ppsz_key = buf;
if (psz_value != NULL) *ppsz_value = psz_value;
  
```

Node Startup Configuration File

```

[init]
Port=14014
Location=294967295
HomeDir=/VORHOME/servant/14014
LogFileName=servant.log
AutoShutdown=60
MT=25
MagiFecTime=60
MagiFecTime=60
getKwInfoRetTime=600
InitNeighbors=3
KeepAliveTime=5
MinNeighbors=2
NCheck=0
NeighborScore=1
NeighborScore=0.1
CacheSize=1000
[beacon]
Retry=15
foo.usc.edu:12311
foo.usc.edu:12312
foo.usc.edu:12313
foo.usc.edu:12314
  
```

- Port is the well-known port that this node listens to
- the black keys in the [init] section are optional
- check the spec for their default values

Node Goes Down (Cont...)

```

nunk:14004
nunk:14013
nunk:14012
nunk:14014
nunk:14015
nunk:14007
nunk:14010
nunk:14008
nunk:14009
nunk:14015
  
```

Neighbors lists:

```

nunk:14003: [init], MinNeighbors=1, Neighbors=2
nunk:14004: [init], Neighbors=3
  
```

Neighbors lists for nunk:14003 and nunk:14004 include nunk:14015, which is crossed out.

Copyright © William C. Cheng

25

How to Manage Timers (Cont...)

➤ Solution: have a timer that goes off every second

- for all the timeouts that are specified as multiple of seconds
- if a timer suppose to go off in 9 seconds, does it matter if it goes off every 9.7 seconds later?
- if a timeout is suppose to be for 15 seconds, initialize a count of 15
- every time the timer goes off, *scrub* all timer-related data structures
- if a count reaches zero, delete the object from the data structure
- you can use a timer thread for this
- if you have events that needs to be timed-out in resolution of multiple hundreds of milliseconds, use another timer that goes off every 100 milliseconds

➤

```

int main(int argc, char *argv[])
{
    gshutdown = FALSE;
    while (!gshutdown)
    {
        Process();
        Cleanup();
        return 0;
    }
}

```

Here is a simple way to implement soft restart:

➤

- only set `gshutdown` to TRUE if you want the program to exit (such as when the autoshtutdown timer goes off)
- otherwise, you are doing a *soft restart*
- in `Cleanup()`, you can clean up everything
- kill all threads, free up all memory, reset all variables (except `gshutdown`)
- keep the *state* of your program in your node's *HomeDir*

Computer Communications - CSC1 551

Copyright © William C. Cheng

28

A Design, Just A Design (Cont...)

➤ Event-driven style

- Identify all your threads
- draw them as circles
- on the previous slide, there are 2 threads to handle communication with a neighboring node
- this is not the only way, you need to decide what you are most comfortable with
- Identify all your shared data structures
- draw them as queues
- use shared data structures for thread-to-thread communication and thread synchronization
- protect each shared data structure with a mutex
- there may be other shared data structures that needs to be protected by mutexes, e.g., logfile

Computer Communications - CSC1 551

Copyright © William C. Cheng

30

Keep Track of Neighbors (Cont...)

➤ Solution: use a connection data structure/object

- Each "connection" has a unique numeric ID
- monotonically increase it when you need a new ID
- Store neighbor hostname and port numbers in it
- Store socket descriptors in it
- Write a bunch of utility functions/methods for it

➤ When you want to refer to something related to a neighbor, have it refer to a connection object

- For a network-read thread, have it reference a "connection" (or a connection ID)
- For a message in the message cache, have it reference a connection number

➤ Hopefully, this can remove some ambiguities

Computer Communications - CSC1 551

Copyright © William C. Cheng

27

A Design, Just A Design

Computer Communications - CSC1 551

Copyright © William C. Cheng

29

Keep Track of Neighbors

➤ How do you keep track of neighbors so you can look it up?

- For example, if you get a message from a neighbor and want to forward it to all other neighbors
- should you use socket descriptor number to distinguish different neighbors?
- probably not a good idea
- socket descriptors get reused as you lose and gain connections
- should you use hostname and port numbers?
- may not be a good idea
- neighbors go up and down

➤

Computer Communications - CSC1 551