

## Programming Assignment 5

*Assigned: March 27**Due: April 14, 11:59:59 PM.**Weight: 6x*

## 1 Introduction

In this project, you will implement Chord [2]. To do so, you must read the paper describing the Chord protocol, algorithms, and implementation which is available at the following URL:

<http://www.cs.berkeley.edu/~istoica/papers/2003/chord-ton.pdf>

Chord is a peer-to-peer lookup protocol which enables the mapping of a key (or identifier) to a peer (or node) in the network. Chord exposes a single function “lookup” to higher-layer applications:

$$\text{lookup}(\langle \text{Key} \rangle) \rightarrow \langle \text{Host} \rangle$$

Chord uses local information and communicates with other peers to find the host (IP address and port) that a given key is mapped to.

Applications may then be built on top of this service that Chord provides. For example, a distributed hash table (DHT) application may distribute the key-value storage of a hash table across many peers and support the typical operations (i.e., insert, get, and remove). The DHT may be implemented on top of Chord by storing the key-value pairs at the node that Chord mapped to the given key.

As described in the paper, Chord nodes communicate via remote procedure calls (RPCs). You should rely on and extend your work on Assignment 4 to support the functions that each node exports to other nodes in the systems. You will continue to use Google Protocol Buffers (GPB) [1] to serialize and deserialize the structured data sent as part of the RPC protocol.

## 2 Protocol

The Chord protocol and algorithms are described in the paper [2]. You should read the paper to learn about the design of Chord, prior to starting your own implementation. The implementation in the paper, shown in Figure 5 and Figure 6, is presented as pseudocode. The pseudocode involves both local and remote function call invocations, determined by the node at which a function call is invoked. Note that this pseudocode is extended by Section IV.E.3 “Failure and Replication”, with changes to the ‘stabilize’, ‘closest\_preceding\_node’, and ‘find\_successor’ function calls. Additionally, there is a new function call ‘get\_successor\_list’ hinted at in the description, which returns a node’s list of successors and is used in the extended function calls. We are providing the GPB protocol message types for you, which may be found in the ‘assignment5’ directory in the ‘materials’ repository.

For this project, you may safely ignore the following portions of the paper:

1. Section IV.E.4 “Voluntary Node Departures” - All node departures will be treated as node failures.

2. Section II “Related Work”, Section V “Evaluation”, Section VI “Future Work”, Section VII “Conclusion”.

### 3 Chord Client

The Chord client will be a command line utility which takes the following arguments:

1. **-a** <**String**> = The IP address that the Chord client will bind to, as well as advertise to other nodes. Represented as an ASCII string (e.g., 128.8.126.63). Must be specified.
2. **-p** <**Number**> = The port that the Chord client will bind to and listen on. Represented as a base-10 integer. Must be specified.
3. **-ja** <**String**> = The IP address of the machine running a Chord node. The Chord client will join this node’s ring. Represented as a ASCII string (e.g., 128.8.126.63). Must be specified if (**-jp**) is specified.
4. **-jp** <**Number**> = The port that an existing Chord node is bound to and listening on. The Chord client will join this node’s ring. Represented as a base-10 integer. Must be specified if (**-ja**) is specified.
5. **-ts** <**Number**> = The time in milliseconds between invocations of ‘stabilize’. Represented as a base-10 integer. Must be specified, with a value in the range of [1,60000].
6. **-tff** <**Number**> = The time in milliseconds between invocations of ‘fix\_fingers’. Represented as a base-10 integer. Must be specified, with a value in the range of [1,60000].
7. **-tcp** <**Number**> = The time in milliseconds between invocations of ‘check\_predecessor’. Represented as a base-10 integer. Must be specified, with a value in the range of [1,60000].
8. **-r** <**Number**> = The number of successors maintained by the Chord client. Represented as a base-10 integer. Must be specified, with a value in the range of [1,32].

An example usage to start a new Chord ring is:

```
chord -a 128.8.126.63 -p 4170 -ts 30000 -tff 10000 -tcp 30000 -r 4
```

An example usage to join an existing Chord ring is:

```
chord -a 128.8.126.63 -p 4171 -ja 128.8.126.63 -jp 4170 -ts 30000 -tff 10000 -tcp 30000 -r 4
```

The Chord client will open a TCP socket and listen for incoming connections on port (**-p**). If neither (**-ja**) nor (**-jp**) are specified, then the Chord client starts a new ring by invoking ‘create’. The Chord client will initialize the successor list and finger table appropriately (i.e., all will point to the client itself).

Otherwise, the Chord client joins an existing ring by connecting to the Chord client specified by (**-ja**) and (**-jp**) and invoking ‘join’. The initial steps the Chord client takes when joining the network are described in detail in Section IV.E.1 “Node Joins and Stabilization” of the Chord paper.

Periodically, the Chord client will invoke various stabilization routines in order to handle nodes joining and leaving the network. The Chord client will invoke ‘stabilize’, ‘fix\_fingers’, and ‘check\_predecessor’ every (**-ts**), (**-tff**), and (**-tcp**) milliseconds respectively.

**Commands** The Chord client will handle commands by reading from `stdin` and writing to `stdout`. There are two command that the Chord client must support: ‘Lookup’ and ‘PrintState’.

‘Lookup’ takes as input an ASCII string (e.g., “Hello”). The Chord client takes this string, hashes it to a key in the identifier space, and performs a search for the host that is the successor to the key (i.e., the owner of the key). The Chord client then outputs the IP address and port information of the host.

‘PrintState’ requires no input. The Chord client outputs its local state information at the current time, which consists of:

1. The Chord client’s own node information
2. The node information for all nodes in the successor list
3. The node information for all nodes in the finger table

where “node information” corresponds to the identifier, IP address, and port for a given node.

An example sequence of command inputs and outputs at a Chord client is shown below. There are four participants in the ring (including the Chord client itself) and `-r` is set to 2. You may assume the same formatting for the input, and must match the same formatting for your output. “>” and “<” represent `stdin` and `stdout` respectively. The input commands will be terminated by newlines (`‘\n’`).

```
> Lookup Hello
< Hello f7ff9e8b7bb2e09b70935a5d785e0cc5d9d0abf0
< 31836aeaab22dc49555a97edb4c753881432e01d 128.8.126.63 2001

> Lookup World
< World 70c07ec18ef89c5309bbb0937f3a6342411e1fdd
< 7d157d7c000ae27db146575c08ce30df893d3a64 128.8.126.63 2002

> PrintState
< Self 31836aeaab22dc49555a97edb4c753881432e01d 128.8.126.63 2001
< Successor [1] 6fa8c57336628a7d733f684dc9404fbd09020543 128.8.126.63 2002
< Successor [2] 7d157d7c000ae27db146575c08ce30df893d3a64 128.8.126.63 2003
< Finger [1] 6fa8c57336628a7d733f684dc9404fbd09020543 128.8.126.63 2002
< Finger [2] 6fa8c57336628a7d733f684dc9404fbd09020543 128.8.126.63 2002
...
< Finger [159] 7d157d7c000ae27db146575c08ce30df893d3a64 128.8.126.63 2003
< Finger [160] f4d60480373006cb24147cd17765000f14aadca3 128.8.126.63 2004
```

## 4 Grading

Your project grade will depend on the parts of the project that you implement. Assuming each part has a “good” implementation, the grades are as follows:

Grade	Parts Completed
A	Resilient to $<$ (-r) node failures.
B	Uses finger tables to route efficiently.
C	Handles many nodes joining with updates to successor lists. Makes routing choices based on successor lists.

## 5 Additional Requirements

1. Your code must be submitted as a series of commits that are pushed to the origin/master branch of your Git repository. We consider your latest commit prior to the due date/time to represent your submission.
2. The directory for your project must be called ‘assignment5’ and be located at the root of your Git repository.
3. You must provide a Makefile that is included along with the code that you commit. We will run ‘make’ inside the ‘assignment5’ directory, which must produce a ‘chord’ executable also located in the ‘assignment5’ directory.
4. You must submit code that compiles in the provided VM, otherwise your assignment will not be graded.
5. Your code must be -Wall clean on gcc/g++ in the provided VM, otherwise your assignment will not be graded. Do not ask the TA for help on (or post to the forum) code that is not -Wall clean, unless getting rid of the warning is the actual problem.
6. You are not allowed to work in teams or to copy code from any source.

## References

- [1] Google Protocol Buffers. <https://developers.google.com/protocol-buffers>.
- [2] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: A Scalable Peer-to-peer Lookup Protocol for Internet Applications. *Networking, IEEE/ACM Transactions on*, 2003.