15-440, Fall 2011
Project Assignment P1: Distributed Password Cracker
Assigned: Sept. 6
Part A Due: Thurs., Sept. 22, 10:00 PM
Part B Due: Thurs., Sept. 29, 10:00 PM

Randy Bryant (`Randy.Bryant@cs.cmu.edu`) is the lead person for this assignment.

1 Logistics

You must work on this project individually. You are free to discuss high-level design issues with other people in the class, but every aspect of your actual implementation must be entirely your own work.

You can get a .tar file from the class AFS directory:

```
/afs/cs.cmu.edu/academic/class/15440-f11
```

Retrieve the file `P1/P1-handout.tar`. Use the Linux tar command to extract the initial set of directories and files that you will use. See the file `README.txt` for a description of the provided files and the organization of the code directories. We have included executable versions of the echo client and server programs described for Part A. These will execute on any Linux machine supporting x86-64. We have also provided a compiled version of our implementation of the LSP protocol that you can use to evaluate your implementation.

We will notify you later on how to hand in your assignments.

2 Overview

In this assignment you will construct a simple distributed system that can harness the power of many processors to speed up a compute-intensive task. In your implementation, you will incorporate features required to create a *robust* system, handling lost or duplicated Internet packets, as well as failing clients and servers. You will also learn the value of creating a set of layered abstractions in bridging between low-level network protocols and high-level applications.

Your system will implement a distributed password cracker. Most systems minimize the cases where a password is stored or transmitted explicitly, since a minor breach of security would enable an adversary to get a copy of the password and then have full access to many system resources. Instead, a *secure hash function* h is applied to the password P to give a *hash signature* $H = h(P)$. Function h is designed so that it cannot easily be inverted. That is, knowing the value of H provides no real information about the value of P. Rather than storing and passing passwords around explicitly, systems strengthen the security of their password management by only keeping their hash signatures.

As an example, in early Unix systems, information about all user accounts, including the hash signatures of their passwords, was stored in an unencrypted and unprotected file `/etc/passwd`. When a user logged in and gave password P, the system would compute $h(P)$ and see if it matched the signature stored in the file for that user. The security of this approach was premised on the assumption that an attacker would be unable to determine the value of a password given its hash signature. Nowadays, Unix systems store the hash signatures in less accessible locations, because it became practical to crack passwords using brute-force approaches, similar to the one you will implement.

One commonly used secure hash function is SHA-1, developed by the National Security Agency. It generates a 20-byte hash signature of an arbitrary sequence of bytes. Here are some examples of applying SHA-1 to text strings:

P	$h(P)$
cat	`9d989e8d27dc9e0ec3389fc855f142c3d40f0c50`
bat	`9b64d4ad2ee2fb847c0d7c96a9480b183cafd31b`
bar	`62cdb7020ff920e5aa642c3d4066950dd1f01f4d`

As these examples illustrate, despite the similarity among text strings, their hash values are very different.

One simple attack on a scheme that relies on the difficulty of inverting a secure hash function is to run a brute-force search, in which we enumerate possible passwords and see if they hash to H. Our measurements show that a typical Andrew Linux machine can compute SHA-1 hashes at a rate of around 10,000 per second. Runnning sequentially, a brute force cracker would require around 9 hours to try all possible passwords consisting of at most 6 lower-case characters. But, if we could harness the power of 100 machines, then we could reduce this time to around 5 minutes.

Any time we want an application to run across tens or hundreds of machines, we must devise ways to make the system robust. Inevitably, some of these machines will stop working, or we will encounter cases where machines become disconnected. Thus, writing a distributed password cracker gives you a chance to hone your skills in designing reliable distributed systems.

The project is split into two parts:

A: Implement the *Live Sequence Protocol*, a homegrown protocol for providing reliable communication with simple client and server Application Program Interfaces (APIs) on top of the Internet UDP protocol.

B: Implement the password cracker application. You will find that the abstract interface provided by LSP will make this part go much smoother than it would if you were just hacking UDP code directly

3 Part A: Live Sequence Protocol

The low-level Internet Protocol (IP) provides what is referred to as an "unreliable datagram" service, allowing one machine to send a message as a packet to another, but with the possibility that the packet will be either dropped or duplicated. In addition, as an IP packet hops from one network node to another, its size is limited to a specified maximum number of bytes, known as the *Maximum Transmission Unit* (MTU). Typically, packets of up to 1500 bytes can safely be transmitted along any routing path, but going beyond this can become problematic.

Very few applications make use of the IP service directly. Instead, they are written in terms of one of the following protocols:

UDP: Also an unreliable datagram service, but it allows packets to be directed to different logical destinations on a single machine, known as *ports*. This makes it possible to run multiple clients or servers on a single machine.

TCP: A reliable streaming service, in which a series of arbitrary-length messages is transmitted by breaking each message into multiple packets at the source and then reassembling them at the destination. TCP handles such issues as dropped packets, duplicated packets, and preventing the sender from overwhelming both Internet bandwidth and the buffering capabilities at the destination.

The Live Sequence Protocol (LSP) provides features that lie somewhere between UDP and TCP, but it also has features not found in either protocol.

- Unlike UDP or TCP, it is specialized to support a client-server communication model
- The server maintains *connections* between a number of clients, each of which is identified by a numeric *connection identifier.*
- Communication between the server and a client consists of a sequence of discrete messages in each direction.
- Message sizes are limited to fit within single UDP packets (around 1000 bytes).
- Messages are sent *reliably*: exactly one copy of each message is received, and messages are received in the same order they are sent.
- The server and clients monitor the status of their connections and detect when the other side has become disconnected.

3.1 LSP: Network Perspective

We will first describe LSP in terms of the messages that flow between the server and one of its clients. Each LSP message contains three values:

Connection ID: A unique, 32-bit, nonzero number assigned by the server to identify this connection. Your implementation may chose any scheme for assigning IDs. Our implementation simply assigns IDs sequentially, starting with 1.

Sequence Number: A 32-bit number that is incremented with each data message sent, starting with number 1.

Payload: A sequence of bytes, with a format and interpretation determined by the application.

We will use the notation (id, sn, D) to describe a packet with connection ID id, sequence number sn, and payload D, where an empty payload is denoted nil.

There are three categories of message:

Connection Request: Having form $(0, 0, nil)$, sent by a client to the server to establish an initial connection.

Data: Having the form (id, i, D) where the payload is nonempty, and this is the i^{th} data message sent in this direction for this connection.

Acknowledgment: Having the form (id, i, nil), sent to acknowledge either a connection request ($i = 0$) or a data message ($i > 0$).

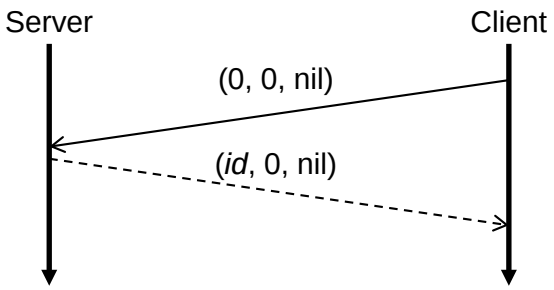

Figure 1: Establishing connection. The server assigns the connection ID and responds to the client with an acknowledgment message (shown as a dashed line.)

Figure 1 illustrates how a connection is established. In this figure, the vertical lines with downward arrows denote the passage of time on both the client and the server, while the lines crossing horizontally denote messages being sent between the two. The client initiates the connection by sending a connection request to the server. The server assigns a connection ID and responds to the client with an acknowledgment message containing this ID, sequence number 0, and an empty payload.

Figure 2 illustrates a normal communication sequence between the server and a client. Data messages can be sent in either direction, and different series of sequence numbers are maintained for each direction. The figure illustrates the transmission of data values D_i and D_{i+1}, having sequence numbers i and $i + 1$ from the client to the server. These are followed by a transmission of data value D_j, having sequence number j, from the server to the client.

Observe that every data message is followed by an acknowledgment message in the opposite direction. Each side must wait for the acknowledgment to be received before it can send another data message. Note, however, that it is entirely possible for one side to receive a data message from the other while waiting for

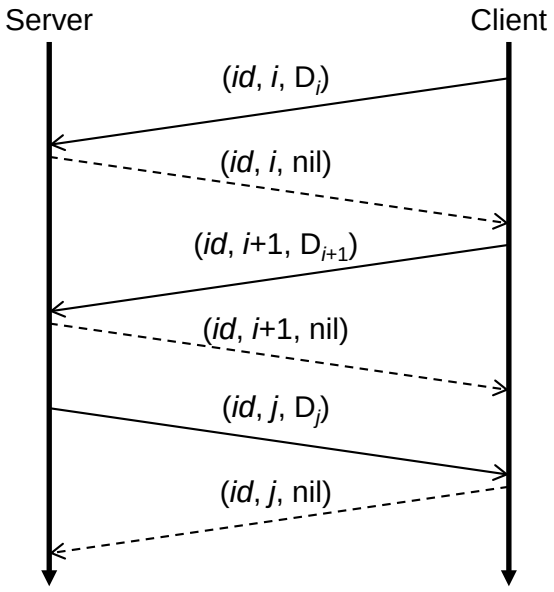

Figure 2: Normal communication. Message number i must be acknowledged before message $i + 1$ can be sent. Acknowledgment messages are shown as dashed lines.

the acknowledgment of a data message it has already sent—the two sides operate asynchronously, and IP does not guarantee that packets arrive in the same order they are sent.

We can see that the basic protocol illustrated in Figures 1 and 2 is not at all robust. On one hand, the presence of sequence numbers would make it possible to detect when a message has been dropped or duplicated. However, if any message—connection request, data message, or acknowledgment—gets dropped, the linkage in either one or both directions will stop working, with both sides waiting for messages from the other.

To make LSP robust, we incorporate a simple time trigger into the servers and clients. That is, timers fire periodically on the clients and server, dividing the flow of time for each process into a sequence of *epochs*. Let us denote this time interval δ. Our default value for δ is two seconds, although this can be varied.

Each time the epoch timer fires a client takes the following actions:

- Resend a connection request, if the original connection request has not yet been acknowledged.
- Send an acknowledgment message for the most recently received data message, or an acknowledgment with sequence number 0 if no data messages have been received.
- If a data message has been sent, but not yet acknowledged, then resend the data message.

The server performs a similar set of actions for each of its connections:

- Acknowledge the connection request, if no data messages have been received.
- Acknowledge the most recently received data message, if any.

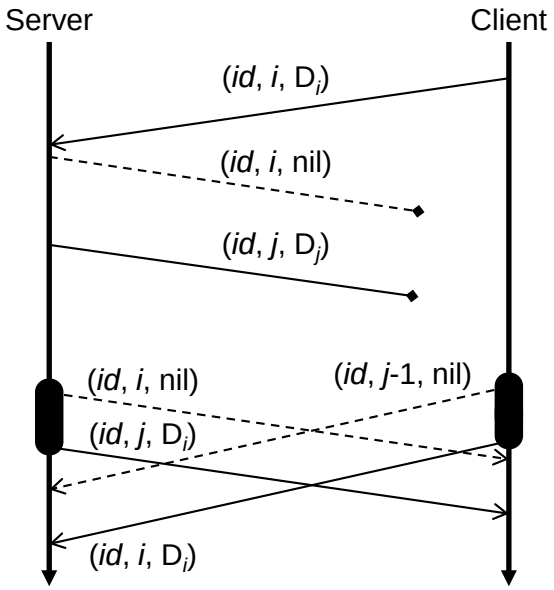

Figure 3: Epoch activities. Both sides send acknowledgments (shown as dashed lines) for the most recently received data and (possibly) resend any unacknowledged data.

- If a data message has been sent, but not yet acknowledged, then resend the data message.

Figure 3 illustrates how the epoch events make up for failures by the normal communication. We show the occurrence of the epoch timer as a large black oval on each time line. In this example, the client attempts to send data D_i, but the acknowledgment message gets dropped. In addition, the server attempts to send data D_j, but the data message gets dropped. When the epoch timer triggers on the client, it will send an acknowledgment of data message $j - 1$, the last data message received, and it will resend data D_i.

Assuming the server has an epoch event around the same time (there is no requirement that they occur simultaneously), we can see that the server will send an acknowledgment for data D_i, and it will resend data D_j.

We can see in this example that duplicate messages can occur: for example, the server received two copies of D_i. For most cases, we can use sequence numbers to detect any duplications. That is, each side maintains a counter indicating which sequence number it expects next and discards any message that does not match the expected number. One case of duplication requires special attention, however. It is possible for the client to send multiple connection requests, with one or more requests or acknowledgments being dropped. The server must track the host address and port number of each connection request and discard any for which that combination of host and port already has an established connection.

One feature of our handling of epochs is that there will be at least one message transmitted in each direction between client and server on every epoch. As a final feature, we will track at the endpoint of each connection the number of epochs that have passed since a message was received from the other end. Once this count reaches some number K, we will declare that the connection has been broken and notify the application that it must take some action. Our implementation sets this limit K as 5. Thus, if nothing is received from

the other end over a total period of $K \cdot \delta$, having a default value of 10 seconds, then the connection will be terminated.

3.2 LSP: Application Program Interface (API) Perspective

We will now provide a view of LSP from the perspective of a Go programmer. We require you to implement the exact API shown here to facilitate automated testing.

3.2.1 Formatting Packets

Each LSP message consists of three fields, declared as a Go structure:

```
type Packet struct {
        Connid int     // Connection ID
        Seqnum int     // Sequence number
        Payload []byte // Message payload
}
```

This structure is converted to a UDP packet by *marshaling* it into a byte sequence according to the JSON data interchange standard. This can be done in Go using the `Marshal` function in the `json` package.

3.2.2 Client API

An application calls the function `NewLspClient` to set up and initiate the activities of an LSP client. This function is declared as follows:

```
func NewLspClient(host string, port int) *LspClient
```

where `host` is the domain name of the server, and `port` is the port on which the server is operating. This function returns a pointer to a data structure of type `LspClient`, which you must define. If the client cannot be initiated, e.g., because the server is not available, the function returns `nil`.

The application client simply reads and writes data messages. It can also *close* a connection, causing the client to cease transmitting or receiving messages. All of the details of establishing a connection, acknowledging messages, and handling epochs is hidden from the application programmer. The following function declarations define this interface:

```
// Client Read.  Returns nil when connection lost
func (cli *LspClient) Read() []byte

// Client Write.  Should not send nil
func (cli *LspClient) Write(payload []byte)

// Close connection.
func (cli *LspClient) Close()
```

As the comment indicates, the `Read` function returns `nil` when the client has become disconnected from the server. Furthermore, the application should never call `Write` with a payload of `nil`

3.3 Server API

The API for the server is similar to that for a client, except that every message is tagged by its connection ID. The server is set up and initiated by calling the function `NewLspServer`, defined as follows:

```
func NewLspServer(port int) *LspServer
```

Its only parameter is the port number that the server should use. The function returns a pointer to a data structure of type `LspServer`, which you must define. If the server cannot be initiated, the function returns `nil`.

The server application can also read and write messages and close connections. In its case, however, every one of these operations includes a connection ID:

```
// Read from connection.  Return nil when connection lost
func (srv *LspServer) Read() (int, []byte)

// Server Write.  Should not send nil
func (srv *LspServer) Write(id int, payload []byte)

// Close connection.
func (srv *LspServer) Close(id int)
```

The `Read` operation returns two values: the connection ID indicating the client from which the message was received and the message payload. If the payload is `nil`, this indicates that the client has become disconnected. The `Write` and `Close` operations include a connection ID as an argument, indicating which connection should be written to or closed. Closing a connection indicates that the server should stop transmitting or receiving messages for this connection.

3.3.1 Setting LSP Parameters

In addition to the basic operations of the client and server, your LSP implementation must support the following operations to facilitate automated testing:

```
// Set length of epoch (in seconds)
func SetEpochLength(delta float32)

// Set number of epochs before timing out
func SetEpochCount(k int)
```

These functions set the parameters δ and K for the epoch duration and the maximum number of epochs that are allowed to pass without receiving any packets from the other side of a connection.

```
// Set fraction of packets that get dropped along each connection
func SetDropRate(rate float32)
```

Setting this parameter to a value r with $0.0 < r \leq 1.0$ should cause program to randomly "drop" packets with probability r. That means that your code that either sends or receives packets should generate a random

number x between 0.0 and 1.0 and, if $x < r$, then your program should either fail to send the packet or proceed as if the incoming packet had never been received.

3.4 LSP Application Example

As an illustration of how applications can be written to use LSP, we show the core functions for a simple echo client and server. Our client is a bit nonstandard in that it *tokenizes* the input, splitting the words of each line typed by the user into a sequence of messages that are sent to the server. This feature provides a convenient way to generate many messages in quick succession for the connection.

The following is the client code:

```
func runclient(cli *lsp.LspClient) {
        for {
                var s string
                // Get next token from input
                n, _ := fmt.Scan(&s)
                if (n <= 0) {
                        cli.Close()
                        return
                }
                // Send to server
                cli.Write([]byte(s))
                // Read from server
                payload := cli.Read()
                if payload == nil {
                        fmt.Printf("Lost contact with server\n")
                        return
                }
                fmt.Printf("[%s]\n", string(payload))
        }
}
```

The `runclient` function is given a pointer to an `LspClient` that has already been created. It runs a loop where it gets the next token from the input (via the `Scan` function), and then calls `Write` to send the text to the server and `Read` to get back the echoed text. It exits when it reads a `nil` packet, indicating that the connection to the server has been lost. It closes the client when an end-of-file is typed as input.

The following is the server code:

```
func runserver(srv *lsp.LspServer) {
        for {
                // Read from client
                id, payload := srv.Read()
                if payload == nil {
                        fmt.Printf("Connection %d has died\n", id)
                } else {
                        // Echo back to client
                        srv.Write(id, payload)
```

```
                }
        }
}
```

The `runserver` function is given a pointer to an `LspServer`. Within its loop, it proceeds by first receiving a connection ID and message via the `Read` function. If the payload is `nil`, this indicates that that client connection has been lost. Otherwise, it sends the payload back to the same connection from which it was received.

3.5 Requirements and Style Issues

Although Go has packages providing conventional synchronization primitives, such as locks, semaphores, and condition variables, you should avoid using these. Instead, use the synchronization constructs provided by Go, namely channels and the `select` statement.

You are free to formulate your own design, but one approach is to view both your client or your server as consisting of three asynchronous event handlers:

Network handler: Reads packets from the UDP connection.

Application handler: Receives messages from the application.

Epoch handler: Activates every time the epoch timer fires.

Each of these handlers executes as a "goroutine" and coordinates and communicates with one another via channels and other data structures.

You will be graded for this part of the assignment both on how well your program can pass a set of automated and manual tests, as well as the quality of your design and implementation. You should think carefully about an overall *architecture* for your system: how you will structure the different components in your client and server, how they will communicate and coordinate with one another, and what data structures you will require. Include in your code documentation on this architecture as a set of comments. The quality of your documentation will be considered in grading.

The `lsp` subdirectory in the code distribution includes code to automatically test your implementation. You will also find it useful to test your code by compiling the echoserver and echoclient applications using your implementation of LSP. Try running with packet drop rates as high as 0.25.

4 Part B: Password Cracker

Your password cracker will involve implementing three different components:

Request: An LSP client that sends a user-specified password cracking request to the server, receives and prints the response, and then exits.

Worker: An LSP client that continually accepts cracking requests from the server, exhaustively checks for a password over a specified range of strings, and then responds to the server with the search result.

Server: An LSP server that manages the entire cracking enterprise. At any time, it can have any number of workers available, and it can receive any number of requests. For each request, it splits the request into smaller jobs, and farms these out to workers. It monitors the returning results and ultimately replies back to the request client.

We will consider only passwords consisting of lower case letters. We can limit the amount of effort by any cracking job by specifying a range of possible passwords, giving the alphabetically smallest and largest passwords to attempt. So for example the range `aaaa` to `zzzz` designates all 4-character passwords. (Our jobs will always involve passwords of some fixed length.)

Format	From–To	Use
`j`	W–S	Join request
`c` *`hash lower upper`*	R–S, S–W	Crack request
`f` *`pass`*	W–S, S–R	Password found
`x`	W–S, S–R	Password not found

Figure 4: Message types for password cracker. In the "From–To" column, 'W' denotes a worker, 'S' denotes the server, and 'R' denotes a request client.

Figure 4 shows the types of messages that are sent among the different system components. These are sent as strings (formatted with spaces between the fields) as the payloads of LSP packets. The overall operation of the system proceeds as follows:

- The server is started using the following command, specifying the port number for the server to use:

  `./server` *`port`*

- One or more workers are started using the following command, specifying the address and port number of the server:

  `./worker` *`host:port`*

- When a worker starts, it sends a join request message to the server, letting the server know that it is available.
- The user generates a cracking request by giving the following command, specifying the address and port of the server, the hash signature to be inverted, and the length of the password

  `./request` *`host:port hash len`*

- The request client should generate a crack request message giving lower and upper values `aa`$\cdots$`a` and `zz`$\cdots$`z`, where the number of `a`'s and `z`'s is based on the length of the desired password.
- The server breaks this request into more manageable-sized jobs (You should choose a suitable maximum job size.) It then starts farming the jobs out to workers by sending them crack requests with a smaller range of possible passwords.

- A worker responds to the server that it has either found a password or that it has not.
- Once a worker finds a password, the server can relay this result to the request client and cease any further effort for this request.
- If the workers exhaust the full range of the original request without finding a password, then the server responds back to the request client that it failed to find a password.

The request client should print its results on standard output as follows. You must match this format precisely in order for our automated testers to work.

- If it finds password *pass*, it should print

  ```
  Found:  pass
  ```

- If it does not find a password, it should print

  ```
  Not Found
  ```

- If the client loses the connection to the server, it should print

  ```
  Disconnected
  ```

We will assume that the server operates all the time, but it is quite possible that a request client or some of the workers can drop out. You should take the following actions when different system components fail:

- When a worker loses contact with the server it should shut itself down.
- When the server loses contact with a worker, it should reassign any job that the worker was handling to a different worker.
- When the server loses contact with a request client, it should cease working on any cracking request being done on behalf of the client. (You need not forcibly terminate a job on a worker; just wait for it to complete and ignore the results.)
- When a request client loses contact with the server, it should exit with an error message.

Your server will need to implement a *scheduler* to assign workers to incoming client requests. You should design a scheduler that balances loads across all requests, so that the number of workers assigned to each outstanding request is roughly equal. Your code should contain documentation on how your scheduler achieves this requirement.

4.1 Testing and Evaluation

You should test your code on actual hash signatures, varying the number of workers, and handling multiple requests simultaneously. You should also experiment by terminating active workers and starting new ones to demonstrate the robustness of your system. With multicore processors, you should be able to see performance gains even when adding more workers on a single machine. You should be able to set the drop rates for the different components to nonzero values (say 0.10 or 0.20) and still have the system operate successfully.

We will provide automated tests for your system components closer to the due date.

5 Useful Resources

You may find the following resources helpful:

- Wikipedia entries for SHA-1 and JSON.
- Online Go documentation, especially the language specification, the package documentation, and the web page titled "Effective Go."
- Chapter 3 of the online book *Network programming with Go* by Jan Newmarch. This document provides helpful examples of networking code, including using UDP.

6 Miscellaneous Advice

Here are some lessons we learned in implementing our own solutions:

- Be careful of the distinction between '=' and '`:=`' in Go.
- Go has different "zero" values for different data types. For integers and floating-point numbers, it's 0; for pointers, it's `nil`, for strings it's `""`.
- The random number generators in the `rand` package start with the same seed everytime you run a program. If you want things to be random from one run to the next, try calling:

```
rand.Seed(time.Nanoseconds())
```

- The use of the SHA-1 hashing function is a bit obscure. Calling `sha1.New()` returns an object of type `Hash`, from the `hash` package. The basic idea is to write a sequence of bytes using `Write`, and then compute the hash with a call to `Sum`.
- The Go library JSON marshaler only marshals structure fields having upper-case names.
- You might want to implement your own library to provide functions to the different components to aid in password cracking. We suggest you follow the Makefile structure seen with the `lsp` subdirectory.

- Goroutines do not use preemptive scheduling. That means that if your worker runs in single-threaded mode (the default), the goroutine doing the cracking will prevent the goroutines implementing LSP activities from communicating with the server, possibly causing a timeout. The solution is to call the function `GOMAXPROCS` in the `runtime` package to specify running at least two parallel threads. You can do this near the start of your `main` function.