

Programming Assignment 1

70 points

Global Employee Location Server Using Selective Repeat in the Data Link Layer Due: Friday, April 8, 2011 at 9 a.m.

This assignment consists of two components: a design paper and an implementation that supports both clients and a server for an international employee location database. This program is to be completed in three-person teams. Each team must reserve a one-hour live demonstration of your project in an available slot on April 8 or 11.

The general assignment is to build a **concurrent server** that handles requests from two or more clients. [*Linux socket calls are used with TCP to emulate a physical layer for transmitting between clients and a concurrent server*]. Both the clients and the server will have a small application layer protocol that defines the interaction between clients and the server.

When a new client connects to the concurrent server, the server will *create* a child to handle all interactions with that client. Whether the server children are *forked as processes* or generated using *concurrent threads* is a team design decision. However, forked processes will then require implementation of shared memory mechanisms to handle concurrent accesses to the database structures required in the design of your application.

Employee Location Application Layer

Assume the *Global Computer Company* wants to maintain a database of its employees that includes a nine-digit ID number, last name, first name, position/title, a photo and the current employee location that includes street address and country. *Global* has advertised for contract bids to implement this network-based application. Your team must submit a design proposal for the employee location database system.

The prototype employee database handles a maximum of 100 employees and can be implemented using SQL or simply as an in-memory data structure that gets uploaded when the server starts. The location application supports two classes of clients: **authorized HR clients** that have the capability to input employee location information that is sent to the location database on the server and **query clients** that enable any *Global* employee to remotely query the database to retrieve employee information (that may include the latest photo on file). You should use the following specifics to define the sizes of the specific database fields:

id_number is a 9-digit identification number.
first_name is a non-blank ASCII string with maximum length of 15 characters.
last_name is a non-blank ASCII string with maximum length of 20 characters.
title is a non-blank ASCII string with maximum length of 15 characters.
address is a non-blank character string (40 character max) identifying the current location of the person.

country is a non-blank character string (15 character max) which uniquely identifies the current country where the employee resides. Note, this can be a country abbreviation to fit into the field.

Generic Application Layer Requirements

Given the time frame of this course, your design proposal needs to support only a few primitive operations in an application layer protocol that correspond to client/server interactions. However, the **minimum** requirement is at least **five** client request types that must include: inputting an employee photo to the database, downloading of an employee photo as part of query request from a client, and a command that requests information about all employees currently located in a single country. Your protocol also needs to provide a mechanism to guarantee only authorized HR personnel can put employee information into the database. The employee photos should be read in as jpg files and stored as jpg files on the server. Note, the abstraction is that the photo constitutes part of a single message at the application layer that is converted to packets at the network layer.

Your application layer design needs to provide a clean mechanism for the client to notify the server when it is done sending a message (and visa-versa) and a scheme to close the client's connection to the server cleanly.

The Network Layer

The *network layer* receives messages from the application layer and creates packets consisting of: 200 bytes of message payload, two bytes for packet sequence number, and any other control bytes specified in your team's design proposal. The *network layer* sends packets to the *data link layer* and waits for packets received from the *peer network layer* through the *data link layer*. Received packets can be data packets or ACK packets or piggybacked packets. The *network layer* is also responsible for taking packet payloads out of received data frames and constructing messages to be sent up to the application layer.

Initially, the *client network layer* makes a one-time call (note, this is an **emulation cheat**) to the physical layer to establish a connection with the *server network layer*. Once a connection has been established, the client network layer begins receiving 200 byte message 'chunks' from the *client application layer* and depositing each message chunk into a packet payload.

Data Link Layer

The data link layer receives packets from the network layer, creates frames, and sends frames to the physical layer for transmission. The data link layer also receives transmitted frames from the physical layer, extracts the payload, reassembles packets, and forwards packets to the network layer.

Frame Format

All frames need a frame-type byte to distinguish data and ACK frames. All data frames must have two bytes for the sequence number, two bytes for error-detection, and one end-of-packet byte. The data link layer sends data frames that contain from 1 to 128 bytes of payload (encapsulated data from the

network layer packet). ACK frames consist of **zero** bytes of payload, a two-byte sequence number, and a two-byte error detection field. Note - the setting of the end-of-packet byte indicates to the *receiving data link layer* that the current received frame is the last frame of a packet.

This assignment requires an implementation of the **Selective Repeat** sliding window protocol at the data link layer. Your design proposal may add other ‘overhead’ bytes to the frame structure deemed necessary to implement **Selective Repeat**. Your design needs to include two error-detection bytes. Your design will need to include sequence numbers in the frames and a mechanism for handling ACK’s. The minimum frame size is your choice, and it is your design decision whether to *piggyback* ACK’s or to send separate ACK frames. The data link layer continues to receive packets from its network layer until its sliding window is full of unACK’ed frames. This requires **multiple virtual timers** on both the client and server side.

Those teams seeking an A on the project and all BS/MS students must implement a sending and receiving window size of four frames. A sliding window of size one should be implemented and tested first. ALL teams struggling with this project can drop back and turn in a working version with a window size of one.

General Data Link Layer Issues

The **data link layer** has to check for transmission errors using the **error-detection** bytes. If the received data frame is in error, this event is recorded and the receiving process waits to receive another frame. While **CRC** at the bit level is possible, it is strongly recommended that you use a two-byte **XOR folding algorithm** of all the frame bytes to create your error-detection bytes and to validate frames upon reception. For the ACK frame, the error-detection bytes simply become a copy of the two-byte sequence number.

The Client Process Flow

Each Client will call the *physical layer* to establish a connection to the concurrent server. The *data link layer* then gets packets from the *network layer*, puts together frames and sends them to the *physical layer* to transmit. The *data link layer* flow depends on events coming from the *network layer*, the current availability within the sliding window, and events coming from the *physical layer*.

The Server Child Process Flow

This assignment requires implementation of a concurrent server that supports concurrent conversations with multiple clients.

Each *server data link layer process (or thread)* waits for frames from the *server physical layer* and passes packets up to the *server network layer*. Similar to the client side, the flow of the *server data link layer* depends on whether there is traffic from the Server to be sent back to the Client (either frames with packets from the *server network layer* or ACK frames that are part of the Selective Repeat protocol).

The Physical Layer

The *physical layer* uses Linux sockets to emulate the sending of constructed frames as actual TCP messages between the clients and the concurrent server. When the *client physical layer* first establishes a client connection to the concurrent server, it **must** send one TCP message to the server child process or thread to identify the client.

Frame Error Simulation

Since real TCP guarantees no errors at the emulated physical layer, your program must inject artificial transmission errors into your physical layer.

Force a **transmission error** in every 5th data frame sent by flipping any single bit in the error-detection bytes prior to transmission of the frame. Force a **transmission error** in every 7th ACK frame sent by using the same flipping mechanism. (i.e., data frames 5, 10, 15, ... sent will be perceived as in error by the receiver and ACK frames 7, 14, 21, ... sent will be perceived as error by the receiver.) Note: since both the clients and the server can send data frames and ACK frames each of these counts that trigger transmission errors must be kept relative to whichever node is sending data or ACK frames. When the *client or server data link layer* times out due to either type of transmission error, it retransmits the data frame with the correct error-detection byte such that the retransmission will not be received with an error.

Assignment Hints

- **[Design]** Plan your design in a modular fashion such that if everything is **not** totally working, you can turn in and demo some type of output that shows **exactly** what is working. Relaxing the sliding window scheme is one option when your project is in trouble.
- **[Documentation]** Your commented program **must** have a special section to explain the details of your specific design decisions. **Remember: This is a team project and all routines must specify the author as part of the documentation!! Team members may not receive the same grade on an assignment due to uneven workload.**
- **[DEBUG]** Include print statements in the various layers while debugging. You should use some type of verbose debugging flag that can be turned on and off.
- **[Performance Timing]** You **must** measure and print out the total execution time of the complete emulated transfer per client.
- **port numbers** - Your clients should have unique port numbers and the clients should treat the server port number like a “well-known” port number.

- **[Version retention]** Exercise extreme care at retaining old versions. Demo experience has shown that a project team should keep old versions (especially in the 11th hour at 3 a.m.) because **often** the newest version done in a rush yields significantly worse results and you need to be able to fall back to a previous version for the demo!

Deliverables

Each team will have to schedule a one hour demo. At the due date, each team should turn in a tarred/zipped file that includes the source code and a README file. At the demo, provide a short user guide that explains the client application commands and syntax.

The README file **must indicate how much of project is working at the time of the demo. It also needs to indicate the non-working parts and some indication of why specific modules are not working correctly. This README information is critical for teams to receive partial-credit for non-operational versions of program 1.**

See the Design Proposal Report for more detailed information about what to turn in at the project demonstration.