

Reliable Data Transfer

Please read all sections of this document before you begin coding.

Objectives

In this MP, you will implement reliable data transfer. You will need to use the knowledge of the reliable data transfer principles discussed in class to design and implement a sender and receiver that transfer a file across an unreliable link.

Guidelines

For this MP, you are encouraged to work with another student in the class, although you may work individually (however, you will be graded on the same scale and responsible for the same amount of work). Please find a partner as soon as possible in order to get a head start on the MP. If you cannot find a partner, post a message on the newsgroup. If this does not work, email the instructor.

You may discuss interpretations of the MP with other people from the class, but any form of group work with people other than your partner is forbidden. This includes code-sharing, which will be checked electronically.

You are required to make sure that your programs compile and run on the EWS Solaris machines (`g1snXX` and `dc1snXX`). Your programs will be tested on EWS machines only. Programs suffering compilation errors when tested by the course staff will earn no credit.

You may find it useful to refer to the Unix man pages as necessary. You may also want to refer to the links posted on the course webpage. If you use code from another source, credit it in your README. Stevens book on UNIX Programming may also prove helpful.

If you are borrowing heavily from a single source, check with the TAs or professors to confirm that you are still in line with the class policies. When in doubt, get permission to use a source well before the due date.

Please indent and document your code. Use meaningful names for variables and follow other style guidelines that enhance the clarity of your code.

Follow the guidelines in the “Hand In” section below when turning in this assignment.

Specification

This MP will require you to create a sender and receiver that implement reliable transfer over the UDP protocol. You will be required to deal with dropped, corrupted, duplicated, and reordered packets in your data transfer. You will also need to implement a sliding window protocol to support efficient data transfer over high-RTT links.

Sender

The sender should be run as follows:

```
sender filename listenport trollhost trollport
```

The sender will read the file specified by the filename and transfer it using a reliable transfer mechanism to the receiver sitting behind a troll. After completing the transfer, the sender should terminate and exit.

The sender will send data to a troll, sitting at `trollhost:trollport`, which will then forward data to the receiver. The sender should bind to `listenport` to receive acknowledgments and other signaling from the receiver. You should only use a single UDP socket for both sending and receiving data.

Receiver

The receiver should be run as follows:

```
receiver filename listenport trollhost2 trollport2
```

The receiver will need to bind to the UDP port specified on the command line and receive a file from the sender sent over the port. The file should be saved to `filename`. Note that you should make sure in your testing that the filename used by the receiver is not the same as the one used by the sender, since otherwise the receiver will overwrite the file and interfere with the sender.

The receiver will use the troll at `trollhost:trollport` to send acknowledgments and any other control signals to the sender. The receiver should exit once the transfer is complete.

Troll

To test your algorithm, use two trolls to connect the receiver and the sender. The troll can be found under `~ece438/MP4/troll`. The troll documentation may be found under `~ece438/MP4/troll.pdf`. You can set up your executable environment by running the commands below. Note that you will need to start the receiver and both trolls before you run the sender. You may also choose to run all the programs on the same machine, which is fine for testing, but your sender and receiver must operate properly when run on different hosts.

```
glsnXX> receiver outfile recvport glsnZZ trollport2
glsnYY> troll -h glsnXX -i recvport trollport
glsnZZ> troll -h glsnWW -i sendport trollport2
glsnWW> sender infile sendport glsnYY trollport
```

Note that although UDP will allow you to send large packets using IP fragmentation, we will configure the troll to discard any packets larger than 1472 bytes (in payload), so you should make sure that you restrict your packets to that size.

Design Document

To ensure that you are making sufficient progress on the project, you will be required to submit a design document as part of your project. The document must be handed in **in class on November 29th**; feedback will be provided to you by the following lecture about any potential problems in your project. You will also submit a revised version of this document with your project.

The design document should include the following components:

1. *Algorithm Overview*

Describe the details of your reliable transfer algorithm. You are welcome to use any of the algorithms described in class verbatim, or to combine ideas from the different algorithms we discussed. Draw the sender and the receiver state machines.

2. *Message Formats*

Describe the message formats you will be using for sending the file segments and any other signaling between the sender and the receiver. Describe any message types you will be using and the fields in them.

3. *Data Structures*

Describe the data structures you will use to buffer segments on the sending and receiving side. Describe some key variables that are keeping track of state (in particular, explain any variables you reference in the state machine). Also list any timers you will be using and what their purpose is.

4. *Work Breakdown*

Describe who will be responsible for which component of the project. Include items such as design, implementation, integration, debugging, testing, and documentation. Use the initial design document to plan the division of work and the final submitted one to document what was actually the case. (If you are working alone, you should omit this section.)

Evaluation and Strategies

Your project will be evaluated along three criteria. First, we will evaluate your design documentation and code for quality. We will then test your project against a number of “adversaries” who will reorder, delay, drop, duplicate, and corrupt packets in such a way as to make the reliable transfer task difficult. You will receive *correctness* points for transferring a file reliably against each adversary; we will use `diff` to compare the received file with the original. You will get partial credit based on how many adversaries you can “defeat” by transferring the file reliably.

If you correctly handle *all* adversaries, we will evaluate your *performance*, in terms of transfer time, against a benchmark and assign you a grade. There are three things to keep in mind regarding performance:

- If you do not get a perfect correctness grade, you will get *no* credit for performance.
- Performance is the smallest component to your grade.
- It is possible, however, to earn bonus points on performance and this can result in a grade of over 100 on the project.

In particular, you can get up to 80% of the grade by implementing a reliable but inefficient solution, and we recommend that you get that component working *first*, and then work on improving performance as there is time available.

To get full marks (before bonus) on the performance component, you will need to implement at least a sliding window algorithm and a timeout estimation algorithm.

Hand In

Place your source code in a directory called MP4 along with a Makefile that, on execution of the command `make`, causes the executable code for the programs `sender` and `receiver` to be generated by the compiler. Do not include any executable or test files with your submission; we will take off marks if you do.

Note that we have been lenient in the past for submissions with incorrect Makefiles, wrong executable names, or compilation errors. Due to the time constraints for grading MP4, we will be performing the tests using a script, therefore, if your project fails to generate the correct executables, you will get **no credit** for correctness or performance.

You should also include in your submission a file called DESIGN which includes the revised version of your design, along with notes of what has been changed since your original in-class submission and why, and a README file describing your implementation log and citing any sources that you used in solving the assignment.

Grading Guidelines

(20 pts) Initial Design

- This document will be evaluated for completeness, rather than correctness

(40 pts) Correctness

- Correct and complete files sent to the receiver under different adversaries.
- There will be several tests against different adversaries, and we will evaluate each one of them by using “`diff`” to check whether a file was transferred correctly.
- Each individual test will be graded on an all-or-nothing basis, but it *is* possible to get partial credit on this component (e.g. by passing half the tests.)

(20 pts+up to 20 pts bonus) Performance

- You will get **NO CREDIT** on this component if you do not pass all the correctness tests
- Grade will be assigned based on transfer time passing each test. We will perform a weighted average of the test times.
- You may receive up to 20 bonus points on this component if your solution achieves excellent performance.

(10 pts) Documentation

- README
- Revised design

(10 pts) Code quality

- Well commented
- Proper indentation
- Sensible variable names