

PROGRAMMING ASSIGNMENT 4

BUILDING A DISTRIBUTED, REPLICATED, AND FAULT TOLERANT FILE SYSTEM: CONTRASTING REPLICATION AND ERASURE CODING

Version 1.1

DUE DATE: Wednesday November 19th, 2025 @ 8:00 pm

OBJECTIVE

The objective of this assignment is to build a distributed, failure-resilient file system. The fault tolerance for files is achieved using two techniques: replication and erasure coding. As part of this assignment, you should identify the trade-off space involving these techniques. For example, your analysis could contrast storage efficiency, CPU overheads, and memory utilization. This assignment has several sub-items associated with it. There are 3 programs that you need to develop.

A Chunk Server responsible for managing file chunks. There will be one instance of the chunk server running on each machine.

1. A controller node for managing information about chunk servers and chunks within the system. There will be only 1 instance of the controller node.
2. A client which is responsible for storing, retrieving, and appending files in the system. The client is responsible for splitting a file into chunks and assembling the file back using chunks during retrieval.

All communications in this assignment are based on **TCP**. The assignment must be implemented in **Java** and the external jar files that you can use are listed towards the end of the assignment. You must develop all functionality yourself. This assignment may be modified to clarify any questions (and the version number incremented), but the crux of the assignment and the distribution of points will not change. This assignment will account for **10 points** towards your cumulative course grade. There are several components to this assignment, and the points-breakdown is listed in the remainder of the text. This assignment is to be done individually.

Generative AI Use and Consequences

Use of AI tools such as ChatGPT, Claude, Github Co-Pilot, or anything of their kind to write or "improve" your code or written work at **any** stage is prohibited; this includes the ideation phase. It is your responsibility to ensure that you don't have the GitHub Co-Pilot extension installed in your IDE; assignment solutions generated by Co-Pilot aren't written by you. Turning in code or an essay written by generative AI tools will be treated as turning in work created by someone else, namely an act of plagiarism and/or cheating. At a minimum, this will result in a 100% deduction (i.e., you will receive a -10/10). To ensure fairness and maintain integrity, grading will also include code reviews, interviews, and on-the-spot code modifications.

Ultimately, you will get out of the class what you put in. Simply copying and pasting code from generative AI tools is not only unethical, it robs you of the chance to learn. Here are four reasons why these generative AI tools undercuts your own education:

1. They take away the struggle that leads to understanding. They rob you of the ability to think and learn the concepts for yourself. Solving problems yourself is how concepts stick. If the AI does the work, what's left for you to learn?
2. You will struggle with the in-classroom quizzes and exams where you will not have access to these tools.
3. Yes, AI tools will become an important part of a software engineer's workflow. But to use them effectively later, you first need solid expertise in the subject matter; and, that only comes from practicing *without* them.
4. These tools are prone to generating imperfect or even incorrect solutions, so trusting them blindly can lead to bad consequences.

1 Fault Tolerant File System Design

In our discussions, we first start with fault tolerance using replication and then describe how the fault tolerance functionality achieved using replication can be achieved using erasure coding.

1.1 Fault Tolerance Using Replication

In this file system, portions (or **chunks**) of a file are dispersed over the set of available machines. There are multiple chunk servers in the system: on each machine there can be at most one chunk server that is responsible for managing chunks belonging to different files. A chunk server stores these chunks on its local disk (in most cases, this will be /tmp).

Every file that will be stored in this file system will be split into 64KB chunks. These chunks need to be distributed on a set of available chunk servers. Each 64KB chunk keeps track of its own integrity, by maintaining checksums for 8KB slices of the chunk. The message digest algorithm to be used for computing this checksum is SHA-1: this returns a 160-bit digest for a set of bytes. In Java, you can use `MessageDigest.getInstance("SHA-1")`; Individual chunks will be stored as regular files on the host file system.

File writes/reads will be done via the chunk servers that hold portions of the file. The chunk server adds integrity information to individual chunks before writing them to disk. Reads done by the chunk server will check for integrity of the chunk slices and will send only the content to the client (the integrity information is not sent).

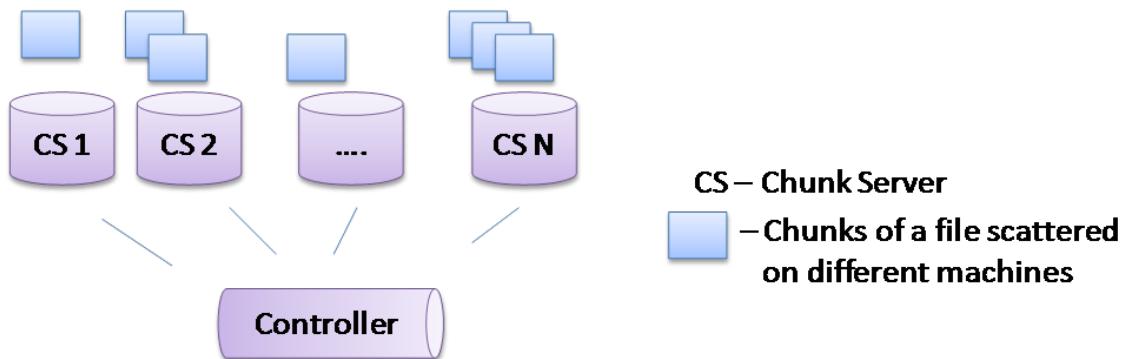


Figure 1: A file will be split into chunks and dispersed on multiple machines.

Each chunk being stored to a file needs to have metadata associated with it. If the file name is `/user/bob/experiment/SimFile.data`, chunk 2 of this file will be stored by a chunk server as `/tmp/<your netID>/chunk-server/user/bob/experiment/SimFile.data_chunk2`. This is an example of the metadata being encoded in the name of the file. There will be other metadata associated with the chunk: this additional information should not be encoded in the filename; this includes –

- Versioning Information: Multiple writes to the chunk will increment the version number associated with the chunk.
- Sequencing Information: There will be a sequence number associated with each chunk.
- Timestamp: The time that it was last updated.

Chunk Server and the Controller Node

Each chunk server will maintain a list of the files that it manages. For each file, the chunk server will maintain information about the chunks that it holds.

There will be one controller node in the system. This node is responsible for tracking information about the chunks held by various chunk servers in the system. It achieves this via heartbeats that are periodically exchanged between the controller and chunk servers. The controller is also responsible for tracking *live* chunk servers in the system. The controller does not store anything on disk, all information about the chunk servers and the chunks that they hold are maintained in memory.

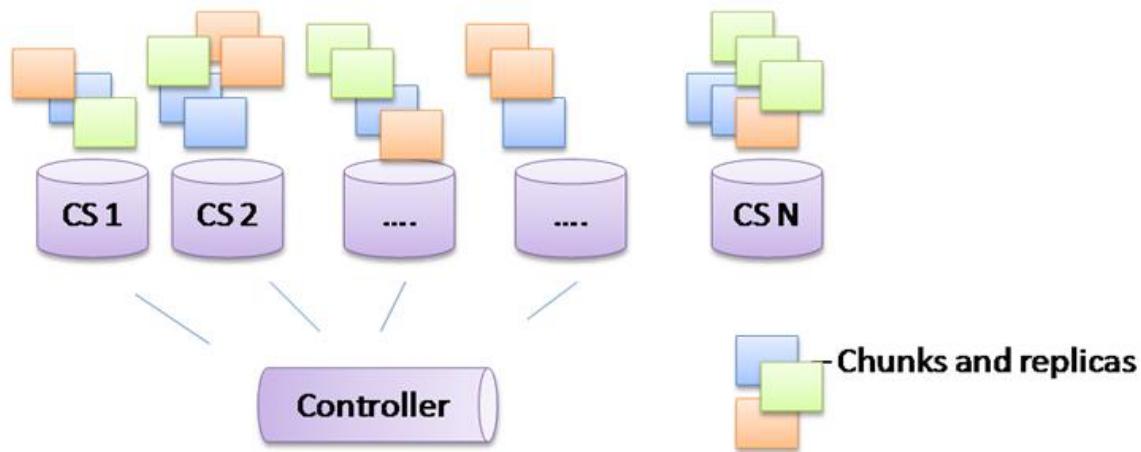


Figure 2: Distribution of chunks of a file and their corresponding replicas.

Heartbeats

The Controller Node will run on a preset host/port. A chunk server will regularly send heartbeats to the controller node. These heartbeats will be split into two

1. A major heartbeat every 60 seconds
2. A Minor heartbeat every 15 seconds

At the 2 minute mark ONLY the major heartbeat should be sent out.

The major heartbeat will include metadata information about ALL the chunks maintained at the chunk server. The minor heartbeat will include information about any newly added chunks. Additionally, when a chunk server detects file corruption, it will report this to the Controller Node.

All heartbeats will include information about the total number of chunks and free-space available at the chunk server. Free space information should be one of the metrics used for distribution of chunks on the set of available commodity machines. Each chunk server should assume to have 1GB of space at the beginning, so the free space is equal to 1GB - <space-used-so-far>.

The Controller is responsible for detecting chunk server failures when it does not receive heartbeats.

Replication of files

Each file should have a replication level of 3; this means that every chunk within the file should be replicated at least 3 times. When a client contacts the Controller node to write a file, the Controller will return a list of 3 chunk servers to which a chunk (64KB) can be written. The client then contacts these chunk servers to store the file. Rather than write to each chunk server directly, if there are 3 chunk servers **A**, **B** and **C** that were returned by the controller, the client will only write to chunk server **A**, which is responsible for forwarding the chunk to **B**, which in turn is responsible for forwarding it **C**. Propagating chunks in this fashion has the advantage of utilizing the bandwidths more efficiently. After the first 64KB chunk of a file has been written, the client (this should be managed transparently by your API) contacts the Controller to write the next chunk and repeat the process. *A given chunk server cannot hold more than one replica of a given chunk.*

Chunk data will be sent to the chunk servers and not the controller. The controller is only responsible for pointing the client to the chunk servers: chunk data *should not* flow through the controller.

Disperse a file on a set of available chunks servers (1 point)

You will take a file and ensure the storage of chunks of this file on different chunk servers. Each chunk of the file should be replicated 3 times. This chunk should be available on the local disk (/tmp) of the chunk server.

Deductions

1. If you use the controller to forward chunk data to the chunk servers **(-1 point)**
2. If more than 1 replica of a chunk is stored at the same chunk server **(-1 point)**

The command for the client is:

upload <source> <destination>

Uploads the file in the cluster. <source> is the local path of the file (e.g., ./project/data.txt) <destination> is the path in the cluster where the file will be stored (e.g., project/data.txt). <destination> can start with a / character or with the file/folder name, but it should not start with /. The chunks of this file should be stored in different chunk servers under the path **/tmp/<your netID>/chunk_server/project/data.txt_chunk<chunk_number>**

Uploading to the same destination of an existing file, should overwrite the existing file.

After issuing the upload command to the client, the client must print a list of ip addresses and ports where the chunks are stored in the following format.

<ip>:<port>
<ip>:<port>
<ip>:<port>
<ip>:<port>
...

In this list, the entries 1-3 indicate where chunk 1 is stored (and replicated). Entries 4-6 indicate where chunk 2 is stored and so on.

Reading a previously stored file (1 point)

During the testing process, you will have to read the file that was previously scattered over a set of chunk servers. For reading each 64 KB chunk, the client will contact the Controller and retrieve information about the chunk server that holds the chunk. Assuming there were no failures, the file read should match the file that was dispersed. The controller must return a **random** chunk server out of the set of chunk server that store the specified chunk.

Deductions

1. If you use the controller to forward chunk data from the chunk servers **(-1 point)**
2. If more than 1 replica of a chunk is accessed at the same time. A given read should result in only 1 copy of a chunk being accessed. **(-1 point)**

The command for the client is:

download <source> <destination>

Downloads the file from the cluster. <source> is the cluster path of the file (e.g., project/data.txt) <destination> is the local path where the file will be saved (e.g., ./project/data.txt). <source> can start with a / character or with the file/folder name, but it should not start with a ./

After issuing the download command to the client, the client must print a list of ip addresses and ports where the chunks are being retrieved from in the following format.

<ip>:<port>
<ip>:<port>
<ip>:<port>
<ip>:<port>

...

In this list, entry 1 indicates where chunk 1 was stored, entry 2 indicates where chunk 2 was stored and so on. Note that, unlike for the upload command, here we are not printing the locations of all 3 replicas of a chunk, but only the one location of the chunk that was used to download the file.

Tampering with chunks (1 point)

Next, we will go to an individual chunk file managed by your File System and tamper this by modifying the content of the file. This may be deleting/adding a line or a word to the file: this is done outside the purview of your chunk server. This should cause the file read to report a data corruption, and the specific chunk (and slice within it) that was corrupted. In this case the output of the download command should be:

<ip>:<port> <chunk-number> <slice-number> is corrupted

For each corrupted chunks

Note that chunk numbers and slice numbers both start from 1 (not 0).

Deductions

1. If you use the controller to detect corruptions of a chunk replica (-1 point)

The grader will execute the download command

Error Correction (2 points)

The contents of one of your chunks will be tampered with. A subsequent read of the file should detect this corruption and initiate a fix of this chunk slice.

If it is detected that a slice of a chunk is corrupted, contact other valid replicas of this chunk and perform error correction for the chunk slice. Error detections will be performed outside the heartbeat control message scheme. The control flow is through the Controller, but the data flow is between the chunk servers.

In this case the download command should first report all corrupted chunks following the format described in the previous section (**Tampering with chunks**), then it should print the location of all the chunks used to reconstruct the file as described in the **Reading a previously stored file** section.

Coping with failures of chunk servers (2 points)

We will terminate one/more of the chunk servers. In response to detection of failures of the chunk servers, the Controller should contact chunk servers that hold legitimate copies of the affected chunks and have them send these chunks to designated chunk servers. Note: The control flow is through the Controller, but the data flow is between the chunk servers.

The metadata maintained at the Controller is updated to reflect this. How are reads handled during this failure?

In this case the download command should first print the list of chunk servers that have failed then it should download the file and print the normal output. The list of failed chunk servers is in the following format:

<ip>:<port> has failed

<ip>:<port> has failed

...

1.2 Fault Tolerance Using Erasure Coding (3 points)

In the previous subsection, fault tolerance was achieved by replicating chunks. The storage requirements in a replication-based setting increase proportional to the number of replicas. Erasure coding offers an alternative to achieve the same degree of redundancy without the corresponding increase in storage costs.

In your scheme with erasure coding, you will take individual chunks, break it into k fragments, expand and encode with redundant pieces of information, and store across different sets of locations. Specifically, your chunks will be broken up into k fragments, erasure coded and expanded into n fragments. These n fragments are then dispersed over the available servers. Note that n must be greater than k ; furthermore, $m=n-k$ is the *degree of redundancy* since any of the k fragments can be used to reconstitute the chunk. For the purposes of this assignment, we will work with $k=6$ and $m=3$.

Similar to the GPS example that we looked at in class, one way to look at erasure coding is from the perspective of linear algebra. You have k variables and $k+m$ equations. We will be using **Reed-Solomon** as the erasure coding algorithm. The terminology typically used in erasure coding settings include the following: (1) The first k fragments are often referred to as the *primary or data shards*, and (2) the next m fragments are referred to as the *parity shards*.

The Reed-Solomon encoding/decoding library as well as the following code snippets are adopted from the open source code implementation available in <https://github.com/Backblaze/JavaReedSolomon>. The jar file has been made available on the course website at <http://www.cs.colostate.edu/~csx55/reed-solomon-erasure-coding.jar>. The following code snippet demonstrates how to use the provided library for encoding a given payload using Reed-Solomon scheme. Code used for some of the data manipulation using Java is omitted for brevity and to focus more on how to use the encoding and decoding APIs. Please follow the comments closely and implement the necessary sections. Also this code snippet assumes the number of data shards (k) is 4 and the number of parity shards (m) is 2.

```
public static final int DATA_SHARDS = 4;
public static final int PARITY_SHARDS = 2;
public static final int TOTAL_SHARDS = 6;

public static final int BYTES_IN_INT = 4;

// file size
int fileSize = (int) inputFile.length();

// total size of the stored data = length of the payload payload size
int storedSize = fileSize + BYTES_IN_INT;

// size of a shard. Make sure all the shards are of the same size.
// In order to do this, you can padd 0s at the end.
// This particular code works for 4 data shards.
// Based on the numer of shards, use a appropriate way to
// decide on shard size.
int shardSize = (storedSize + DATA_SHARDS - 1) / DATA_SHARDS;

// Create a buffer holding the file size, followed by the contents of
// the file
// (and padding if required)
int bufferSize = shardSize * DATA_SHARDS;
byte [] allBytes = new byte[bufferSize];

/* You should implement the code for copying the file size, payload and
padding into the byte array in here. */

// Make the buffers to hold the shards.
byte [] [] shards = new byte [TOTAL_SHARDS] [shardSize];

// Fill in the data shards
for (int i = 0; i < DATA_SHARDS; i++) {
    System.arraycopy(allBytes, i * shardSize, shards[i], 0, shardSize);
}

// Use Reed-Solomon to calculate the parity. Parity codes
// will be stored in the last two positions in 'shards' 2-D array.
ReedSolomon reedSolomon = new ReedSolomon(DATA_SHARDS, PARITY_SHARDS);
reedSolomon.encodeParity(shards, 0, shardSize);

// finally store the contents of the 'shards' 2-D array
```

The corresponding code snippet for decoding and recovering the original file is shown below.

```
public static final int DATA_SHARDS = 4;
public static final int PARITY_SHARDS = 2;
public static final int TOTAL_SHARDS = 6;

public static final int BYTES_IN_INT = 4;

// Read in any of the shards that are present.
// (There should be checking here to make sure the input
// shards are the same size, but there isn't.)
byte [] [] shards = new byte [TOTAL_SHARDS] [];
boolean [] shardPresent = new boolean [TOTAL_SHARDS];
int shardSize = 0;
int shardCount = 0;

// now read the shards from the persistance store
for (int i = 0; i < TOTAL_SHARDS; i++) {
    // Check if the shard is available.
    // If avaialbe, read its content into shards[i]
    // set shardPresent[i] = true and increase the shardCount by 1.
}

// We need at least DATA_SHARDS to be able to reconstruct the file.
if (shardCount < DATA_SHARDS) {
    return;
}

// Make empty buffers for the missing shards.
for (int i = 0; i < TOTAL_SHARDS; i++) {
    if (!shardPresent[i]) {
        shards[i] = new byte [shardSize];
    }
}

// Use Reed-Solomon to fill in the missing shards
ReedSolomon reedSolomon = new ReedSolomon(DATA_SHARDS, PARITY_SHARDS);
reedSolomon.decodeMissing(shards, shardPresent, 0, shardSize);
```

In your support for fault tolerance using erasure coding you are allowed to develop your own metadata schemes. The points distribution for the 3 points for this component are as follows:

1. Successful retrieval and assembly of erasure coded fragments into individual chunks and reconstruction of the entire file handling corruptions and failures. **(3 points)**

Commands and output formats are the same as for replication.

2 Third-party libraries and restrictions:

You are not allowed to use any 3rd party libraries other than for the Reed-Solomon Codes. The jar file for the Reed-Solomon codes will be posted on the course website. You are not allowed to download *any* other code from *anywhere* on the Internet. You are also not allowed to use RPC or distributed object frameworks to develop this functionality (there is a **10 point deduction** for this). You should not build GUIs for this application; in the context of this assignment, GUI-building is an auxiliary path (there is a **10 point deduction** for building a GUI). You can discuss the project with your peers at the architectural level, but the project implementation is an individual effort.

3 Testing Scenario

Commands to start the `controller node` (only one command at a time):

```
java csx55.dfs.replication.Controller portnum  
java csx55.dfs.erasure.Controller portnum
```

Commands to start the `chunk server` (only one command at a time):

```
java csx55.dfs.replication.ChunkServer controller-ip controller-port  
java csx55.dfs.erasure.ChunkServer controller-ip controller-port
```

Commands to start the `client` (only one command at a time):

```
java csx55.dfs.replication.Client controller-ip controller-port  
java csx55.dfs.erasure.Client controller-ip controller-port
```

We will test “replication” mode first and then “erasure” mode. We will never test “replication” and “erasure” mode together.

We will test your code with 1 controller node and between 10 and 20 chunk server nodes. We will stage large files, perform chunk corruptions and see if the system is able to detect, and crucially, recover from these data corruptions.

4 Rubric:

The auto-grader will assign points in this order

1 point	For correctly storing a file.
1 point	For correctly retrieving a file.
1 point	For reporting a corrupted chunk.
2 points	For fixing the corrupted chunk and correctly retrieving the file.
2 points	For detecting a failed chunk server, restoring the replication of chunks, and correctly retrieving the file.
3 points	For doing all the above with erasure coding.

5 What to Submit

Use **CANVAS** to submit a single .tar file that contains:

- The src folder containing all the Java files related to the assignment (please document your code)
- the build.gradle file you use to build your assignment
- A README.txt file containing a description of each file and any information you feel the GTA needs to grade your program.

Filename Convention: All classes should reside in a package called `csx55.dfs`. The archive file should be named as <FirstName>-<LastName>-HW<x>.tar. For example, if you are Cameron Doe then the tar file should be named Cameron-Doe-HW4.tar.

6 Version Change History

This section will reflect the change history for the assignment. It will list the version number, the date it was released, and the changes that were made to the preceding version. Changes to the first public release are made to clarify the assignment; the spirit or the crux of the assignment will not change.

Version	Date	Change
1.0	10/22/2025	First public release of the assignment
1.1	11/11/2025	Modified the path where ChunkServers should store their chunks. E.g., <code>/tmp/chunk_server</code> becomes <code>/tmp/<your_netID>/chunk_server</code> .