UMass Boston cs444 proj2:
Hamming Code and RAID 2 for Error Correction
Posted Thursday, 15 Oct 2026
Due Saturday, 31 Oct 2026 at 11:59pm

J.H.DeBlois

1 Overview

For proj2, the second C code project, there are four parts: proj2-1 (dir), proj2-2 (required code), proj2-3 (readMe) and proj2-4 (code with incremental deliveries). We call them C2-1, C2-2, C2-3 and C2-4. The C2-2 includes handwritten and on-server, that gives C2-2a and C2-2b. The chart on page 1 of this description was discussed in class, along with endian.c, on Tues 13 Oct.

2 Introduction

This project develops **RAID 2** encoding and decoding based on the **Hamming(7, 4)** code. See class website for lecture **Hamming.pdf**, endianness, Error Correction Code (ECC) and **RAID 2** covered in lecture **ch5.pdf** on Input/Output.

Briefly, **RAID 2** splits every 4-bit nibble of a file into four data bits and adds three parity bits for error detection and correction. So ASCII character 'A' which is one byte (hex 41) is first split into two nibbles (hex 4 = binary 0100 and hex 1 = binary 0001).

Then, using **Hamming(7, 4)**, each 4-bit nibble will be encoded as 7 bits named: P1, P2, D1, P4, D2, D3, and D4. Thus, 'A' which is two nibbles will require 14 bits. However, for **RAID 2** the bits from each nibble go to 7 separate disks. We will use 7 files instead.

As a reminder, for **Hamming(7, 4)**, these are the parity rules: For P1, take even parity of D1, D2, D4. For P2, take even parity of D1, D3, D4. For P4, take even parity of D2, D3, D4.

This means that the P1 bits of each nibble in the encoded file are concatenated and stored on one disk drive; similarly for the other six types of bits. If one drive (for us, one file) fails, **Hamming(7, 4)** code can detect and correct the error.

Example: Consider an ASCII input file "ABCD", size 4 bytes. In hex bytes, it reads: 0x41, 0x42, 0x43, 0x44. The eight nibbles of the file are individual hex digits: 4,1,4,2,4,3,4,4. The following table lists the corresponding **Hamming(7, 4)** bits by nibble.

In this table, the original four bytes (left column 1) are encoded into 7 bit-type bytes (columns 2,3,4,5,6,7,8). Read them top to bottom. For example, column 2 is bit-type P1 with data 11101111_2, which is 0xEF in hex. Check that you can decipher the other columns. BE SURE YOU UNDERSTAND THE TABLE!

nibble	P1	P2	D1	P4	D2	D3	D4
4	1	0	0	1	1	0	0
1	1	1	0	1	0	0	1
4	1	0	0	1	1	0	0
2	0	1	0	1	0	1	0
4	1	0	0	1	1	0	0
3	1	0	0	0	0	1	1
4	1	0	0	1	1	0	0
4	1	0	0	1	1	0	0
	0xEF	0x50	0x00	0xFB	0xAB	0x14	0x44

3 C Code Programs to Write

The tasks for this project are to write two C programs. The first program shall be named `abcraid.c` where abc stands for your initials. It encodes one file xyz into seven files. It is run like this:

```
pe15> ./raid -f xyz
```

It creates seven files:

```
xyz.part0
xyz.part1
xyz.part2
xyz.part3
xyz.part4
xyz.part5
xyz.part6
```

These files store the bits of P1, P2, D1, P4, D2, D3, and D4, in that order. The command-line option `-f filename` is a required feature.

The second program shall be named `abcdiar.c` where abc is your initials and diar is raid spelled backwards. It decodes seven files back to the original file. It is run like this:

```
pe15> ./diar -f xyz -s 8
```

It looks for the seven files `xyz.part[0-6]`, decodes them, and creates a new file called `xyz.2`. The command-line options `-f filename` and `-s numberOfBytesInfilename` are required features.
You are also required to add a debug switch -d # to output intermediate results for each program.

4 Project Tasks

First, create a project folder under your `cs444` directory. This is C2-1, likely done as an InClass (IC).

```
pe15> mkdir ~/proj2
```

Next, copy five files from the instructor's directory.

```
pe15> cp /home/hdeblois/cs444/proj2/* ~/proj2
```

The files are the following.

```
-rw-r--r-- 1 hdeblois 5694072 Oct 26 21:33 completeShakespeare.txt
-rwxr-xr-x 1 hdeblois   13400 Oct 26 21:49 diarjhd
-rw-r--r-- 1 hdeblois     106 Oct 26 21:34 Makefile
-rwxr-xr-x 1 hdeblois   13344 Oct 26 21:49 raidjhd
-rw-r--r-- 1 hdeblois       8 Oct 26 21:42 test.txt
```

The two text files `completeShakespeare.txt` and `test.txt` will be used to test your programs. The `Makefile` can be used to compile your project. The two executables `raidjhd` and `diarjhd` were compiled from the instructor's code.

Prepare a `readMe.txt` to list how you developed your code in increments. This is C2-2a (handwritten, likely done InClass) and C2-2b (the on-server version that you keep adding to as you work). Put you name and ID on the top line. Enter your AI Use-Case for proj2. (See Syllabus: UC-1: none, UC-2: ahead, UC-3: during and UC-4: entire.) You can change your use case if you need to.

For each day you work, make an entry in your readMe.txt file with date, what you worked on and difficulties encountered. Keep a list of sources you use, including URLs for on-line documents and prompts used for any Large Language Models (LLMs). BE SURE YOU NAME YOUR FILE: readMe.txt!

Put the commands you use to test both programs in your readMe.txt. See Grading Rubric section for the list.

Keep incremental copies of your code on the server. For example, each day copy your code to a name like `abcraid_oct15.c` before you make changes. Then make your changes in `abcraid.c`. Code will be collected from the server often. Put a header comment in your code with the name of your code and your name, since some user names bear no resemblance to real names.

If you consult and/or use (small features of) code on the internet, list each source in your readMe.txt and cite the source in your code according to the MIT Writing Code formats. This applies to code you modify as well. For LLM generated code, mention it in your readMe and include the prompt and date in the code. For all code you bring in, mark the START of the section and the END of the section, so we know that the code you wrote yourself starts after each END marker. Do not cite fellow student's code because it is not a published source and you are not supposed to share your code. If applicable, please cite the code given you by the instructor.

Regarding compiling with a Makefile, the name of your C code file must match the name(s) you include in the `Makefile`. On the server, type: `man make` to see information about GNU Make Utility. Modify the given Makefile so it compiles both your programs (add the abc part, your initials, to each place the name raid or diar appears).

4.1 Task 1: to RAID

You can run `raidjhd` and see its output.

```
pe15> ./raidjhd -f test.txt
pe15> ls -l
pe15> ./raidjhd -f completeShakespeare.txt
pe15> ls -l
```

Your task is to write your own `abcraid.c` that encodes **RAID 2**. Use getopt for command line options. It accepts the command-line option `-f filename`. It accepts a -d # switch to request debug, and # could be 1 for your basic intermediate result or there could be several #s possible, showing debug1, debug2, etc printf output (if (debug1) printf, if (debug2, printf, etc.). This is the -d flag with optional argument so you need the colon!

4.2 Task 2: from RAID

You can run `diarjhd` and see its output.

```
pe15> ./diarjhd -f test.txt -s 16
pe15> ls -l
pe15> ./diarjhd -f completeShakespeare.txt -s 5694072
pe15> ls -l
```

Note that the sizes of these two test files are multiples of four. Thus, the last bytes of the RAID files are filled with real data bits. If the size is not a multiple of four, there will be trailing non-data bits, which must be ignored. However, you are not expected to handle this scenario — assume the sizes of all test files are multiples of four.

Your task is to write your own `abcdiar.c` that decodes **RAID 2** and corrects errors. Use getopt for command line options. It accepts the command-line option `-f filename` and `-s numberOfBytes`. It also accepts a -d # switch, as explained above. This is the -d flag with an optional argument so you need the colon!

You can verify that the restored files are identical to the originals:

```
pe15> diff test.txt test.txt.2
pe15> diff completeShakespeare.txt completeShakespeare.txt.2
```

There should be no output from the `diff` command. When there are no failed drives, we can use just parts 2, 4, 5, and 6 to reconstruct the original file.
Most importantly, **RAID 2** can perform error correction when one drive fails. The techniques for error detection and correction are described in the Hamming slides. You can test error correction by deliberately corrupting a RAID file. For example, you could type the first line given below:

```
pe15> cp completeShakespeare.txt completeShakespeare.txt.part2
pe15> ./diarjhd -f completeShakespeare.txt -s 5694072
pe15> diff completeShakespeare.txt completeShakespeare.txt.2
```

Since diar will correct it, there should be no output from the `diff` command. If you want to corrupt a different file, remember to run `raidjhd` again to restore the corrupted file back to the correct status before you corrupt another — **RAID 2** can handle only one failed drive.

It is best to write your C code on the server using a text editor, e.g., `nano` or `emacs`. Whether you write your code there or on a different machine, you need to compile your code on the server and run it there each time you work on it because the libraries may differ between machines. Incremental deliveries are required starting no less than 24 hours before the code is due. They protect you, in case your latest version doesn't run. Use the getopt.h library to set up your command-line options.

For command-line options, you may want to include the `getopt.h` library and use it to implement the options. As for variables, include int debug and the debug switch `-d` to set `int debug = 1` and use it to turn on `printf` statements.

5 Grading Rubric

Be sure to "make your code your own," by organizing it carefully and by knowing exactly how it works. Also, the instructor and graders may require you to come in and explain how particular sections of your code work. Know the algorithm before you start. Understand the table on page 1.

The first step in grading is to run your executable in your course directory and read your `readMe.txt`. After we check that your executable runs, we copy your code to another location, recompile it and test it more. In grading, we run Stanford's Measure of Software Similarity (MOSS) to check that no one duplicates someone else's code. All combinations are tested. If MOSS score greater than 45 percent similar, both students could be penalized.

Be sure to notice how the uppercase and lowercase letters are used in naming the files and directories that you are required to create (`Makefile`, `proj1`, `readMe.txt`). Names must match exactly or your submission will not be collected.

Any file changes made after 11:59 pm on the due date will will incur a late penalty, 1 percent per hour late down to a maximum loss of 20 percent.

- (15 points) You have the right files in the right places and your initials are added to the basic files names.
- (15 points) Both your programs compile without errors or warnings, the executables match the ones you submitted, and the input switches work
- (15 points) `readMe.txt` contains your full name, you list the six commands required (two to run your abcraid with each test file, two to run your abcdiar with each test file and two to show how to display debug for each program), you list of what was worked on in each increment and you list of sources consulted if any and LLM prompts and dates if any.
- (20 points) Your full name is in both programs, your code reads the CL correctly, your code is correctly indented, names are well-chosen, comments are useful, and citing is specific – including URL, whether copied or copied and modified, LLM prompt and date, whether copied or copied and modified and the beginning and end of each cited or LLM section is marked by START and by END markers so we know where your own code resumes
- (5 points) The sources you cite in your programs are a subset of the sources you cite in your readMe.txt and you did a code walkthru if requested
- (15 points) Tests with test.txt passed, output for completeShakespeare.txt is correct, and your debug switch shows intermediate results
- (15 points) After doing the suggested command that corrupts the output file your raid created, your diar can correct the error and diff shows no difference between original input file and corrected file.