

Practice Midterm Exam Solutions

1. Short Answer

Part A: Scattering Inodes

One example:

- **Benefit:** fewer seeks between accessing an inode and accessing its payload data
- **Drawback:** increased complexity in finding inodes

Part B: Log Issues

The problem is that the append operation isn't idempotent: if the operation is done multiple times, it will produce a different result than if it is done only once. This is a problem because it's possible that the operation has already been reflected on disk at the time of a crash, so replaying the log may perform the operation a second time. In general we don't know whether operations have already been performed or not when a crash occurs, so it's important that all operations in the log are idempotent.

Part C: Doubly-Indirect Indexing

- Advantage: supports even larger files than original design.
- Disadvantage: accessing payload for very small files requires many disk accesses

Part D: Rename

Find the inode number of the file being moved by iteratively drilling toward it as **pathname_lookup** would. Remove file's **dirent** from parent directory payload, and then drill toward new parent directory of second argument, creating new **dirent**'s along the way as needed. Finally, append new **dirent** on behalf of new name, with new name (e.g. **index-w19.html**) and existing inode number.

Part E: Crash Recovery Tradeoffs

Generally, if we want to improve durability and consistency, we must sacrifice performance in order to perform additional operations / write more aggressively to avoid data loss. As an example, the block cache may have delayed writes of about 30 seconds to improve performance, but this means we may lose the last 30sec of data. If we wrote immediately, this would improve crash recovery, but at the expense of performance. (Another example is data logging – logging doesn't usually support payload data due to the amount of data that would log, thus impacting performance, but it means that we don't log payload changes so that data may be lost. If we did log payload data, we would have better crash recovery but the logging operations would be much more intensive and affect system performance.)

Part F: Multithreading

Three possibilities (there are more, these are just 3):

50 (if thread 1 runs to completion, then thread 2 runs to completion)

40 (if thread 2 runs to completion, then thread 1 runs to completion)

60 (if thread 1 increments, then thread 2 increments, then thread 1 multiplies, then thread 2 multiplies)

2. Duet

Sample Solution

```
static void duet(int incoming, char *one[], char *two[], int outgoing) {
    pid_t pids[2];
    int fds[2];
    pipe(fds);
    pids[0] = fork();
    if (pids[0] == 0) {
        close(fds[0]);
        close(outgoing);
        dup2(incoming, STDIN_FILENO);
        close(incoming);
        dup2(fds[1], STDOUT_FILENO);
        close(fds[1]);
        execvp(one[0], one);
    }

    close(incoming);
    close(fds[1]);
    pids[1] = fork();
    if (pids[1] == 0) {
        dup2(fds[0], STDIN_FILENO);
        close(fds[0]);
        dup2(outgoing, STDOUT_FILENO);
        close(outgoing);
        execvp(two[0], two);
    }

    close(outgoing);
    close(fds[0]);
    waitpid(pids[0], NULL, 0);
    waitpid(pids[1], NULL, 0);
}
```

3. Expression Evaluation

Sample Solution

```
typedef struct ThreadInfo {
    vector<int> v;
    mutex m;
} ThreadInfo;

static void evaluate(Expression& exp, ThreadInfo& info) {
    int result = exp.evaluate();
    info.m.lock();
    info.v.push_back(result);
    info.m.unlock();
}

static bool concurrentAnd(const vector<Expression>& expressions) {
    ThreadInfo info;

    vector<thread> threads;
    for (size_t i = 0; i < expressions.size(); i++) {
        threads.push_back(thread(evaluate, ref(expressions[i]),
                                ref(info)));
    }

    for (thread& t : threads) t.join();

    printResults(info.v);
}
```

4. Multiprocessing

Part A: Close

The `close(sp_supplyfd)` within the test program can only indicate the end of input that `sort` must process if all other references to the write end are closed. That doesn't happen unless we close it everywhere it's needed. If we don't, `sort` will continue waiting for more input forever thinking that more could come.

Part B: Thyme

Sample Solution

```
int main(int argc, char *argv[]) {
    struct timespec start;
    clock_gettime(CLOCK_REALTIME, &start);
    pid_t pid = fork();
    if (pid == 0) execvp(argv[1], argv + 1);
    waitpid(pid, NULL, 0);
    struct timespec finish;
    clock_gettime(CLOCK_REALTIME, &finish);
    print_elapsed_time(&start, &finish);
    return 0;
}
```