

Comparative Analysis of Methods to Handle Race Conditions in File-Based Concurrency Using PHP and a Custom Solution Without Using Existing Methods

Dr. Vimal K Chaudhari

Assistant Professor

Department Of Computer Science, VNSGU, Surat, India

vimal@vnsgu.ac.in

Abstract

Race conditions in file-based concurrency introduce significant challenges in web applications, mainly when PHP is used for server-side scripting. This paper try to discovers traditional ideas to mitigate race conditions, which includes file locking, temporary files, atomic operations, database systems, in-memory caching, message queues, and semaphores. Moreover, this paper recommends a unique custom solution that avoids PHP's built-in concurrency tools to handle race condition. The study examines the issue of race conditions in file-based systems, confirms the conventional ideas, and gives an unique custom solution that does not depend on PHP's in built locking mechanisms/functions or external libraries. Through a comparative analysis, the performance, reliability, and scalability of the custom solution are measured against traditional methods. The results tells us that the custom approach offers comparable reliability with enhanced flexibility, although at the cost of increased complexity and longer execution times. This research paper contributes to a in depth understanding of concurrency control in resource-constrained environments and provides a foundation for further optimization.

Key words: Race condition, PHP, concurrency control, inbuilt function PHP.

1. INTRODUCTION

Concurrency control is a very significant aspect of modern web development, mainly in systems where multiple processes or threads access shared resources at the same time. PHP, a widely-used server-side scripting language, most of the time implements file-based operations for tasks such as logging, caching, and data persistence. But, these operations can lead to race conditions when more than one requests attempts to read or write to the same file at a time, resulting in data inconsistency or corrupted or incorrect data or states. In built solutions to race conditions in PHP is to utilize functions like flock() for file locking or rely on external tools such as database transactions mechanism. Most of the time these methods often introduce dependencies and limitations, particularly in situations where simplicity or we don't have the database support. This paper explains and implement race conditions in file-based concurrency within PHP, analyses existing mitigation strategies, and gives us a custom solution that avoids native locking mechanisms and third-party tools. The projected custom method implements a file-based locking ideas with improvements to address common challenges such as deadlocks and scalability problems. By combining theoretical analysis, practical code implementations, and experimental outcome, this work tries to provide developers with a all-inclusive understanding of concurrency control options and their respective adjustments in PHP applications.

1.1 UNDERSTANDING RACE CONDITIONS

A race condition occurs when two or more processes or threads access shared resources at the same time, and the final result depends on the timing of their execution. In PHP's file-based concurrency, this demonstrates as more then processes reading and writing to the same file without synchronization, as implements in the following code example:

```
<?php
```

```

$file_ori = 'data_ori.txt';
for ($i = 0; $i < 10; $i++) {
    $pid = pcntl_fork();
    if ($pid == -1) die('Could not fork');
    elseif ($pid) continue;
    else {
        file_put_contents($file_ori, "Process $i writing\n", FILE_APPEND);
        exit;
    }
}
while (pcntl_waitpid(0, $status) != -1);
echo file_get_contents($file_ori);
?>

```

Without synchronization, the output may be incomplete or incorrect as it will over write previous result.

RELATED WORK

The problem of handling race-conditions in file-based concurrent accesses are well addressed in theoretical and practical researches especially for PHP and web applications. In this section we highlight some important works that develop fundamental principles and practical aspects whether they are related to PHP in particular or parallelism in general.

The official PHP Manual[1] serves as a basic reference, where built-in file system functions like flock() and file_put_contents are documented. It is a repository of all possible functions available, however does not look at alternative ways to have concurrency control or any custom methods.

Silberschatz et al. [2] provide fundamental principles of concurrency control such as mutual exclusion, synchronization and deadlock-avoidance in Operating System Concepts. Their fine discussion not being PHP-centric, it provides an interesting theoretical basis to understand race conditions in file systems. Also, Tanenbaum and Bos [3] describe in their Modern Operating Systems book operations on file systems and concurrency support with focus on atomic operations, operating system-level locking devices.

Studies on system scalability show that coarse-grained locking considerably limits performance under high concurrency. Paznikov *et al.* [4] verified that dynamic race detection mechanisms introduce significant runtime overhead, even if we applied advanced compilation techniques. These observations propose that heavyweight runtime methods are ill-suited for PHP environments, where execution time and response latency are critical limitations.

Subsequent research has tried to decrease this overhead using novel detection approaches. Zhang *et al.* [5] anticipated efficient timestamping techniques for sampling-based race detection, presenting measurable enhancements in performance. Though, such solutions rely on deep runtime or compiler integration and that's why cannot be directly applied to PHP's process-based execution model, which offers limited control over memory access and thread scheduling.

Concurrency testing frameworks additionally reveal that race conditions persist even when traditional synchronization is engaged. Li *et al.* [6] presented **Fray**, a concurrency testing platform that methodically explores execution interleavings. Their results tell us that many race conditions remain latent despite the presence of locking mechanisms, giving us, the limitations of lock-centric designs in real-world systems.

Dossche *et al.* [7] presented a context-sensitive static analysis approach capable of identifying race patterns in complex systems. Which is attractive for PHP-based applications due to very less runtime overhead.

Hardware-assisted methods have also been suggested to mitigate synchronization costs. Shastri *et al.* [8] presented HMTRace, a memory-tagging-based race detection mechanism that considerably reduces detection overhead. Even if it is not applicable to PHP directly, this work strengthens a broader research trend toward avoiding coarse-grained locking and minimizing shared-state contention.

The complexity of synchronization errors is further highlighted by Shi *et al.* [9], who inspected synchronization-reversal data races and confirmed that subtle concurrency violations may arise even in well-structured programs. Their findings propose that reliance on file-level atomicity alone is insufficient for confirming correctness in concurrent environments.

McKenney and Slingwine [10] presented Read-Copy Update (RCU), explaining how decreasing lock conflict can improve scalability without losing correctness. While RCU works at the kernel level, its fundamental principles gives conceptual guidance for designing lightweight concurrency mechanisms at the application level.

If we try to understand and analyse all of above studies , these studies show both the strengths and weaknesses of current concurrency control methods in PHP and related systems. Introductory works [2], [3] lay the theoretical and basic groundwork, while others [4]–[10] look into practical, real world and experimental methods. However, there is a significant gap in research on standalone and web based application, lock-free, and PHP-native concurrency solutions, which this study aims to fulfil the gap.

3. EXISTING METHODS FOR HANDLING RACE CONDITIONS

3.1 FILE LOCKS (FLOCK())

Mechanism: Uses inbuilt given locking to restrict file access, supporting exclusive (LOCK_EX) and shared (LOCK_SH) modes.

```
<?php
function SWrite($file1, $data1) {
    $fp1 = fopen($file1, 'a');
    if (flock($fp1, LOCK_EX)) {
        fwrite($fp1, $data1);
        flock($fp1, LOCK_UN);
        fclose($fp1);
        return true;
    }
    fclose($fp1);
    return false;
}
?>
```

3.2 TEMPORARY FILES

Mechanism: initially writes in temporary files, merging them into the target file atomically whenever gets the access of original file.

```
<?php
function writeWithTemp($file1, $data1) {
    $temp1 = tempnam(sys_get_temp_dir(), 'php_');
    file_put_contents($temp1, $data1);
    $fp1 = fopen($file1, 'a');
    if (flock($fp1, LOCK_EX)) {
        fwrite($fp1, file_get_contents($temp));
        flock($fp1, LOCK_UN);
    }
    fclose($fp1);
    unlink($temp1);
}
?>
```

3.3 ATOMIC OPERATIONS (RENAME())

Mechanism: take advantages of the atomicity of rename() for file replacement.

```
<?php
function atomicWrite($file1, $data1) {
    $temp1 = tempnam(sys_get_temp_dir(), 'php_');
    file_put_contents($temp1, $data1);
    return rename($temp1, $file1);
}
?>
```

3.4 DATABASE SYSTEMS

Mechanism: Uses database transactions mechanism for concurrency control.

```
<?php
$pdo1 = new PDO('mysql:host=localhost;dbname=test', 'user', 'pass');
$pdo1->beginTransaction();
$stmt1 = $pdo->prepare('INSERT INTO logs (msg) VALUES (?)');
$stmt1->execute(["Data"]);
$pdo1->commit();
?>
```

3.5 IN-MEMORY CACHING (REDIS)

Mechanism: Uses atomic operations like INCR and APPEND.

```
<?php
$redis1 = new Redis();
$redis1->connect('127.0.0.1', 6379);
$redis1->append('logs', "Data\n");
?>
```

3.6 SEMAPHORES

Mechanism: Uses semaphores for synchronization.

```
<?php
$semId1 = sem_get(ftok(__FILE__, 't'), 1);
sem_acquire($semId1);
file_put_contents('data1.txt', "Data\n", FILE_APPEND);
sem_release($semId1);
?>
```

4. PROPOSED CUSTOM SOLUTION

4.1 CONCEPT

As we are avoiding using built-in methods, we propose a custom file-based locking mechanism using a lock file with timeout and stale lock detection. This approach simulates mutual exclusion without flock() or external tools.

4.2 ALGORITHM

Step 1: Create a distinct temporary file: here each process creates a temporary file with a name based on timestamp and random number (e.g temp_[timestamp]_[random]. txt).

Step 2: Write to Tempfile : It writes data to above created tempfile.

Step 3: Test Target File Availability: This step tests the availability of target file by trying to rename its temp file into the target file with a unique suffix (e.g log_[sequence]. txt).

Step 4: Retry Logic: If the rename failed (another process is working), we simply sleep for a given fixed time (10ms e.g.) and retry with a higher sequential number.

Step 5: Merging Results: After any and all processes are completed, a cleanup script merges the sequenced files into a single log.

```
<?php
// Function to mimic a process writing data
function processData($data, $targetBaseName = "log") {
    // Step 1: Unique temp file
    $tempFile1 = "temp_" . uniqid(mt_rand(), true) . ".txt"; // Using uuid for exclusivity
    $maxRetries = 10;
    $initialDelay = 10; // milliseconds
    $sequence = 0;
```

```

// Step 2: Write data to file
if (file_put_contents($tempFile1, $data) === false) {
    echo "Error: Failed to write to temp file $tempFile1\n";
    return false;
}

// Step 3 & 4: Checking availability of target file and retry logic
$delay = $initialDelay;
while ($sequence < $maxRetries) {
    $targetFile1 = "$targetBaseName_{$sequence}.txt";

    if (@rename($tempFile1, $targetFile)) {
        echo "Success: Data written to $targetFile\n";
        return true;
    }
    // If rename not happen, another process may get active
    echo "Retry: Rename to $targetFile failed, retrying ($sequence/$maxRetries)... \n";
    $sequence++;
    usleep($delay * 1000);
    $delay *= 2;
}

// If max tries reached, clean up temp file and send failure
unlink($tempFile1);
echo "Error: Failed to rename $tempFile1 after $maxRetries attempts\n";
return false;
}

// Merge all sequenced files into a final log
function cleanup($targetBaseName = "log", $finalLog = "final_log.txt") {

    $files = glob("$targetBaseName_*.txt");
    if (empty($files)) {
        echo "No files to merge.\n";
        return;
    }
    sort($files); // sorting

    // final log file for writing
    $finalHandle = fopen($finalLog, 'w');
    if ($finalHandle === false) {
        echo "Error: Could not open $finalLog for writing\n";
        return;
    }

    // Merge
    foreach ($files as $file) {
        $content = file_get_contents($file);
        if ($content !== false) {
            fwrite($finalHandle, $content . PHP_EOL);
            echo "Merged: $file into $finalLog\n";
            unlink($file);
        } else {
            echo "Error: Could not read $file\n";
        }
    }
}

fclose($finalHandle);
echo "Cleanup complete: All files merged into $finalLog\n";
}

// function for running multiple processes
function simulateProcess($processId) {
    $data = "Data from process $processId at " . date('Y-m-d H:i:s');
    processData($data);
}

for ($i = 1; $i <= 5; $i++) {
    // In a real scenario, this could be forked or run in parallel

```

```

    simulateProcess($i);
}
sleep(1);
cleanup();
?>

```

5. METHODOLOGY

We tested each method with 100 simultaneous processes that attached 100-byte strings to a shared resource at the end. We repeated this 10 times on a system that ran PHP 8.2+, Apache 2.4+, and Ubuntu 22.04+, which had a 4-core CPU and 8+GB of RAM.

5.1 RESULTS

Method	Avg. Execution Time (s)	Data Integrity (%)	Scalability	Complexity	Dependencies
flock()	0.85	100	High	Low	None
Temporary Files	1.02	100	Moderate	Moderate	None
rename()	0.62	100	Moderate	Low	None
Database (MySQL)	1.45	100	High	High	Yes
Redis	0.38	100	High	Moderate	Yes
Beanstalkd	0.95	100	High	High	Yes
Semaphores	0.90	100	Moderate	High	None
Custom Solution	1.20	98	Low	Moderate	None

The above results encapsulate the performance, data integrity, scalability, complexity, and dependency requirements of numerous concurrency control mechanisms, such as PHP's built-in methods, external systems, and a custom solution. Redis is giving us fastest and most scalable option, but introduces dependencies. flock() and rename() are dependency-free, simple, and reasonably fast, making them ideal for small to medium concurrency workloads. Custom solution is a feasible lightweight, dependency-free alternative, trading slightly higher latency and minor data integrity risk for full control and zero external dependencies.

6. CONCLUSION

In this paper, we examined some ways to address race conditions resulting from PHP file-based concurrency and presented a custom solution for file-locking that does not rely on built-in mechanisms. Traditional approaches such as flock() and atomic operations provide a solid and efficient solution but the custom method can be a valid choice on setups with constraints. Further advances to the custom solution can be made by including adaptive polling, and future work can explore hybrid file-based and in-memory techniques.

8. FUTURE SCOPE

The above mention method is robust because the new method replaces PHP's internal file locking tools. Yet, a number of aspects are left to be examined in the future.

We could then write adaptive retry heuristics and tune them under different loads. Secondly, a nice thing to do could be combining this method with widely-used PHP frameworks like Laravel or CodeIgniter giving some practical scalability and real-world performance understanding. Third, cross-platform testing would provide a measure of portability and robustness on Windows, macOS, and distributed file systems etc.

REFERENCES

- [1] PHP Manual, "Filesystem Functions," *PHP Documentation*, 2025.
- [2] A. Silberschatz, P. B. Galvin, and G. Gagne, *Operating System Concepts*, 10th ed., Hoboken, NJ: Wiley, 2018.
- [3] A. S. Tanenbaum and H. Bos, *Modern Operating Systems*, 4th ed., Upper Saddle River, NJ: Prentice Hall, 2015.
- [4] A. Paznikov, A. Kogutenko, Y. Osipov, M. Schwarz, and U. Mathur, "Compiling Away the Overhead of Race Detection," *arXiv preprint arXiv:2512.05555*, Dec. 2025.
- [5] M. Zhang, et al., "Efficient Timestamping for Sampling-Based Race Detection," in *Proc. ACM SIGPLAN Conf. Programming Language Design and Implementation (PLDI)*, Montreal, Canada, Jun. 2025, pp. 123–135.
- [6] A. Li, et al., "Fray: An Efficient General-Purpose Concurrency Testing Platform for the JVM," *arXiv preprint arXiv:2501.12618*, Jan. 2025.

- [7] N. Dossche, B. Abrath, and B. Coppens, "A Context-Sensitive, Outlier-Based Static Analysis to Find Kernel Race Conditions," *arXiv preprint arXiv:2404.00350*, Apr. 2024.
- [8] J. Shastri, X. Wang, B. A. Shivakumar, F. Verbeek, and B. Ravindran, "HMTRace: Hardware-Assisted Memory-Tagging Based Dynamic Data Race Detection," *arXiv preprint arXiv:2404.19139*, Apr. 2024.
- [9] Z. Shi, U. Mathur, and A. Pavlogiannis, "Optimistic Prediction of Synchronization-Reversal Data Races," in *Proc. 46th Int. Conf. Software Engineering (ICSE)*, Madrid, Spain, May 2024, pp. 789–800.
- [10] P. E. McKenney and J. D. Slingwine, "Read-Copy Update: Using Execution History to Solve Concurrency Problems," in *Proc. IEEE Int. Parallel and Distributed Processing Symp. (IPDPS)*, Cancun, Mexico, Apr. 1998, pp. 123–130.
- [11] PHP Group, "Filesystem Functions," *PHP Manual*. [Online]. Available: <https://www.php.net/manual/en/ref.filesystem.php>

