

Project Part 1: Data and Scalability

1. Team member names and student IDs:

Mythry Rudra - 989473114

Rani Supriya - 989471012

Naveena Voora - 989472599

2. Screenshots of the implementation tasks 1 through 6

The screenshot displays a web application interface with a sidebar on the left containing navigation items like 'Manifest', 'Service workers', 'Storage', 'Local storage', 'Session storage', 'IndexedDB', 'TodoDatabase', 'CompletedTasks', 'TodoList', 'status', 'Cookies', 'Private state tokens', 'Interest groups', and 'CacheStorage'. The main content area shows a table with 13 rows of data. The table has four columns: '#', 'Key (Key path: "status")', 'Primary key (Key path: "id")', and 'Value'. The 'Value' column contains JSON objects for each task, such as '{id: 1, task: 'Task 1', status: 'completed', dueDate: ...}'. Below the table, a console log shows the following output:

```
console.error(database error: , event.target.error),
}
Database already populated. Starting performance tests...
Time with dedicated object store for completed tasks: 3.300000044703484 ms
Time without optimization: 176.5 ms
Time with indexing: 184.7000000298023 ms
Time with read-only transaction: 349.5 ms
```

Project: Performance Optimization Techniques for NoSQL Databases Using Indexed DB

Objective:

To explore performance optimization techniques for large datasets within Indexed DB by simulating a 100,000-record 'Todo List' database. The objective is to evaluate the effectiveness of different optimization techniques and understand the impact of database design and configuration on retrieval speed.

Tasks and Results:

Baseline Test (Without Optimization)

Execution Time: 176.5 ms

Description: In this test, all records were retrieved without any specific optimization applied, providing a baseline performance measure for comparisons.

Lesson Learned: Running the test without optimization provided a useful benchmark against which other techniques were evaluated.

Challenge Faced: The retrieval time was notably high, demonstrating the inefficiency of fetching large datasets without any specific data retrieval structure.

Insight: The baseline test underlined the need for optimizations in large datasets, as the retrieval time was too high for practical applications in performance-critical scenarios.

Test with Indexing on Status Attribute

Execution Time: 184.7 ms

Description: This test used an index on the status attribute to optimize retrieval. The goal was to allow IndexedDB to locate "completed" tasks more efficiently by leveraging this index.

Lesson Learned: The indexed retrieval did not improve performance as anticipated, yielding a longer execution time compared to the baseline.

Challenge Faced: The index on "completed" tasks did not deliver a performance benefit, potentially due to the high ratio of records being retrieved or the mixed dataset structure, which made the index less effective.

Insight: This test emphasized that indexing is not always a straightforward solution to performance problems in NoSQL databases, especially with large datasets that involve mixed status fields.

Test with Read-Only Transaction

Execution Time: 349.5 ms

Description: A read-only transaction was used in this test, intending to reduce overhead by restricting the transaction to data reading operations only.

3. Lessons learned from the project

According to goals, utilizing a read-only transaction unexpectedly led to decreased performance.

It is clear that transaction settings do not always equal faster performance in IndexedDB because the read-only transaction optimization produced no performance gains.

The significance of verifying every possible optimization was brought to light by this assignment. Although read-only transactions are thought to reduce overhead, in fact they do not always improve efficiency, especially when working with large datasets in IndexedDB. One of the best ways to optimize big NoSQL databases is to separate frequently requested groups of data into separate object stores. This resulted in a demonstration of the significance of object store structure and data modeling for high-speed access in performance-critical applications.

4. Screenshot of the GitHub branch submission (e.g., project_1_team_1, project_1_team_n)

