

# Performance Analysis and Optimization of the Hurricane File System on the K42 Operating System

M.A.Sc. Thesis Oral Presentation

By David Tam  
`tamda@eecg.toronto.edu`

# Introduction

**Motivation:** K42 OS needs a scalable FS

- Many workloads, FS is point of contention
- FS critical component in achieving scalability

**Candidate:** Hurricane File System

- “scalable” & flexible
- shared-memory NUMA multiprocessors

**Goal:** Examine & (if necessary) improve scalability of HFS

**Approach:** Iterative experimentation



1. Microbenchmarks
2. Macrobenchmark

# Initial Scalability Results

## Microbenchmark:

Sequentially read separate, non-shared 4MB file

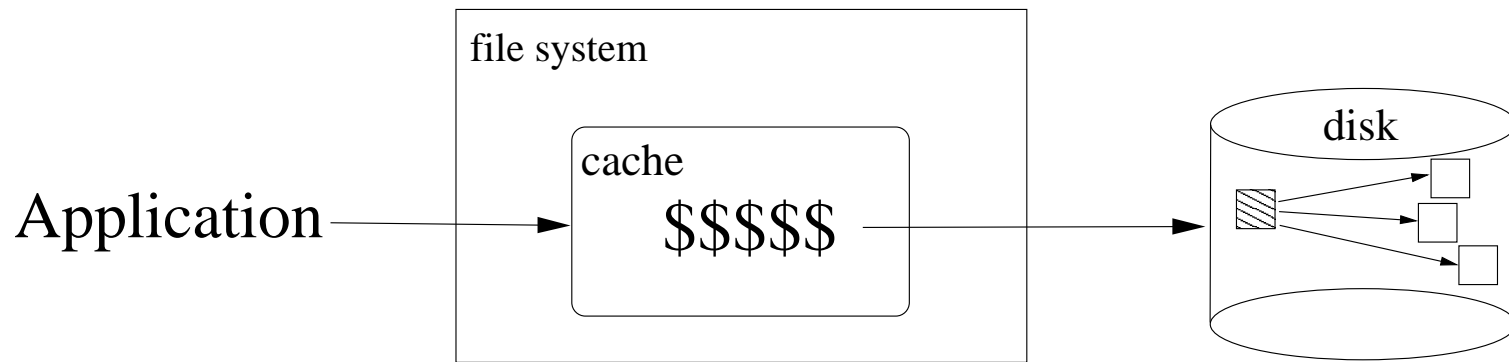
### SimOS



# Role of File System

**Abstraction:** files, directories, attributes, etc

**Implementation:** meta-data -- i.e. file layout on disk

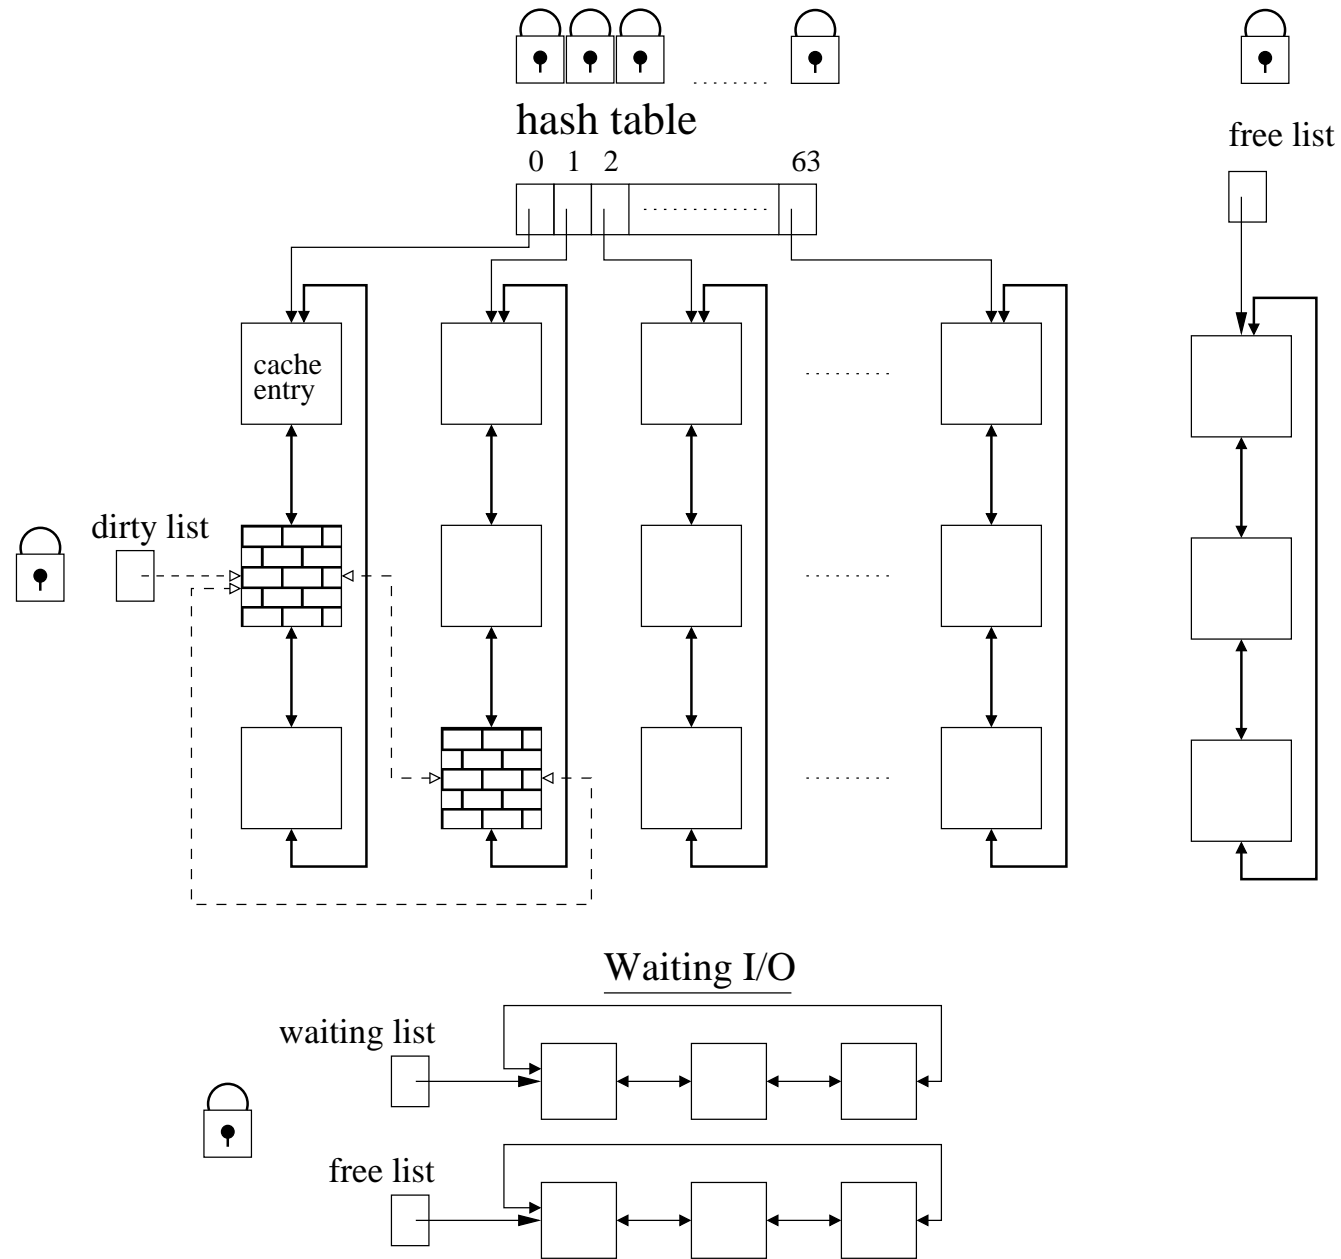


**Cache:** Implemented using hash tables & linked lists

## Hash Tables:

- Popular data structure
- Used throughout OSes & FSES
- $O(1)$  lookup
- Tunable parameters & policies

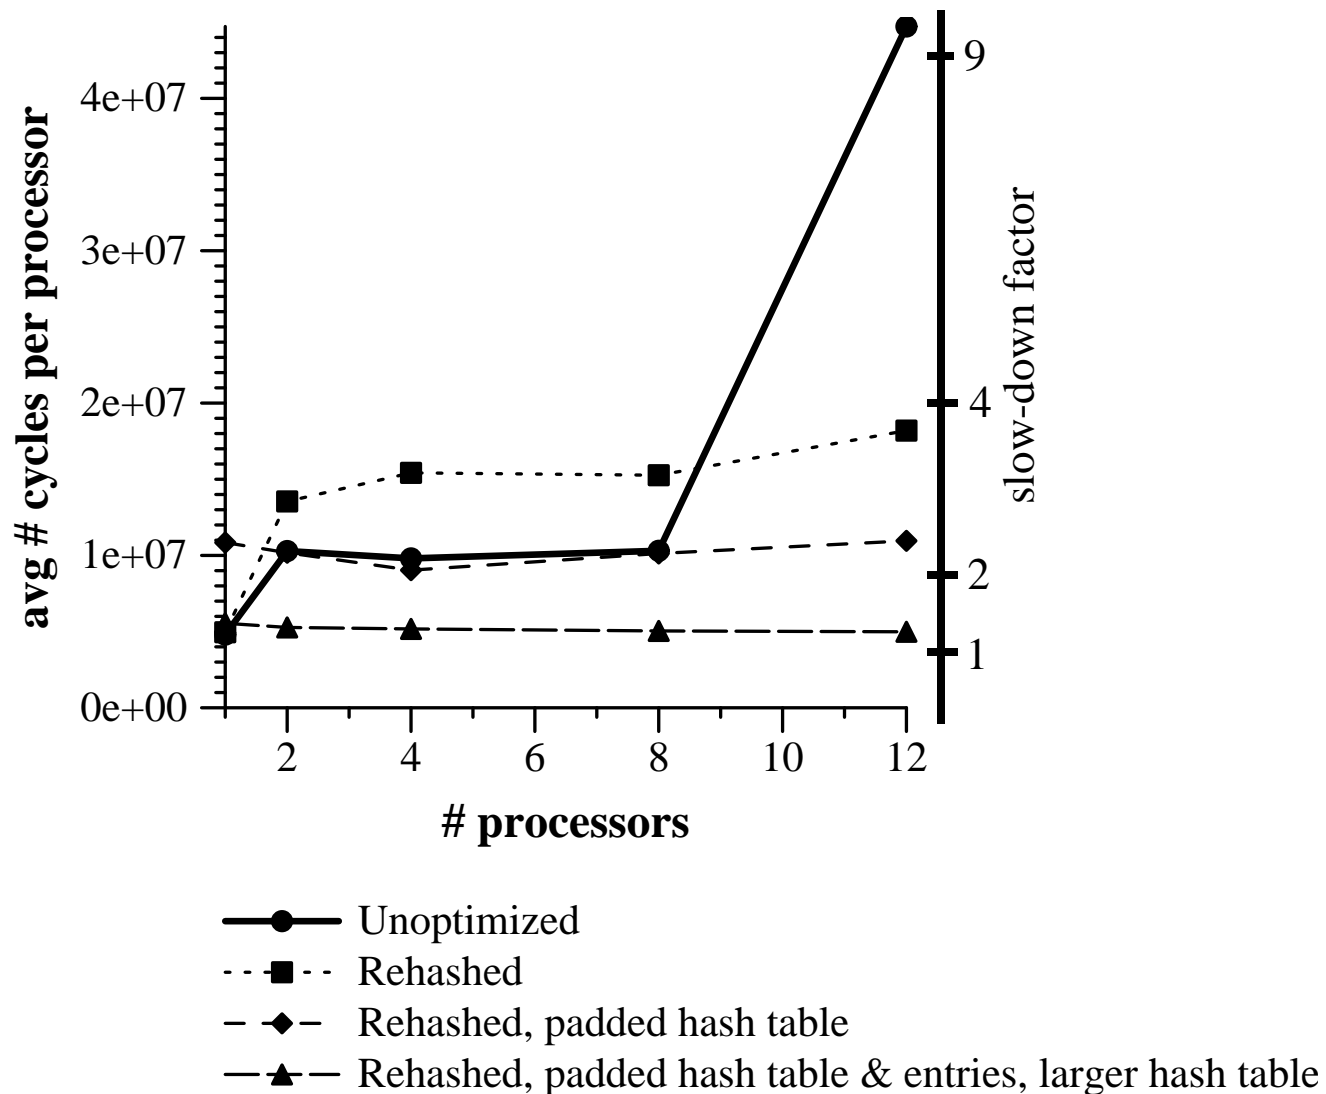
# Typical Organization of Cache



# Microbenchmarks - Read

Stress small meta-data cache


Read separate, non-shared 12MB file



# Microbenchmarks - Read

Stress large meta-data cache

Read separate, non-shared 4MB file



## Optimizations:

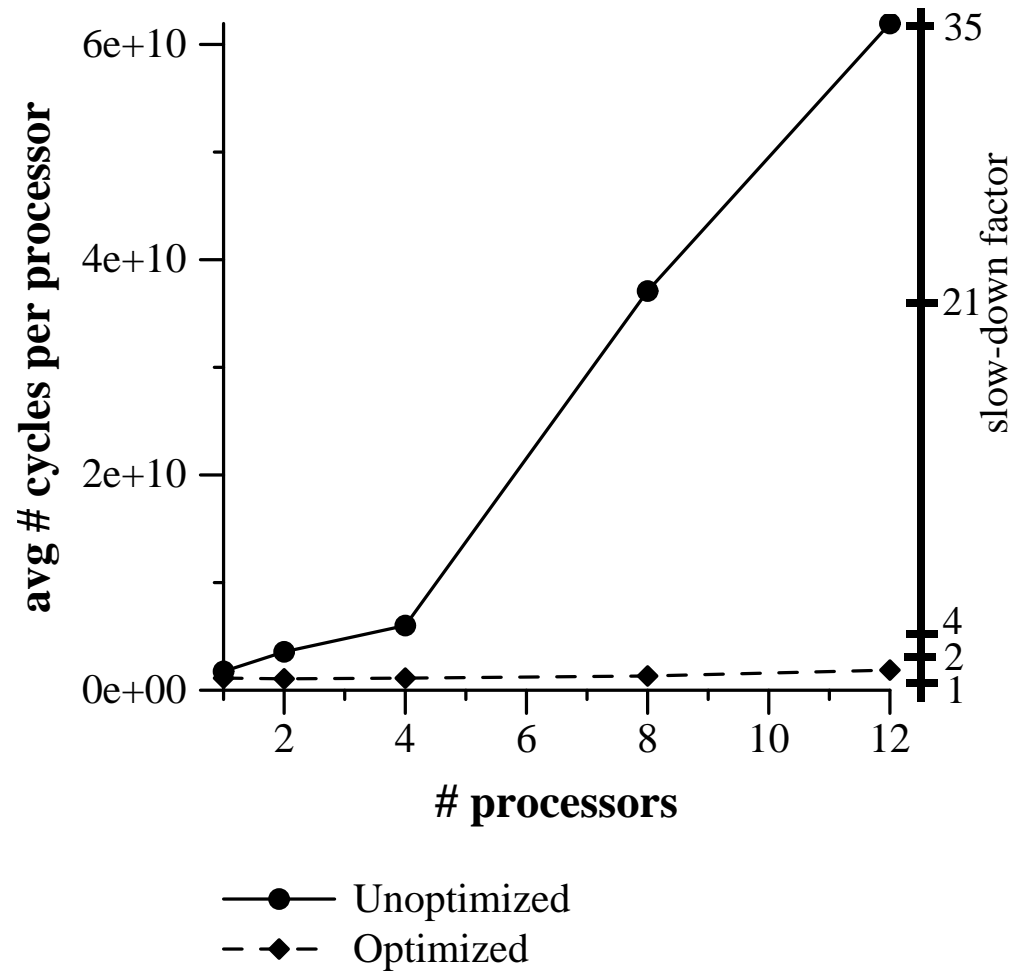
1. modified hash function
2. larger hash table
3. fine grain locks
4. padded hash table, cache entries, others
5. eliminated free list

- Unoptimized
- - ■ - - Global lock & optimizations 1, 2, 5
- - ◆ - - All optimizations

# Macrobenchmark

## Workload:

- sim. Web server
- World Cup 1998
- 4000 requests
- 14MB data transfer
- 37 secs



# Conclusions

## **Summary:**

Scalability of HFS has been improved

## **Take-Home Message:**

Achieving scalability is a non-trivial task.

Also requires attention to:

1. implementation details
2. hardware characteristics
3. impact of workloads

## **Contributions:**

1. Examined & improved FS scalability on K42
2. Provided guidance on achieving FS scalability
3. Enable further FS & OS research on K42

## **Future Work:**

1. Workloads with sharing
2. Benchmarks on hardware & validation of simulation results
3. Implement more distributed data structures