

**Robust and Efficient Data Management for a Distributed Hash Table**

by

Josh Cates

Submitted to the Department of Electrical Engineering and Computer Science  
in partial fulfillment of the requirements for the degree of

Master of Engineering in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2003

©2003 Josh Cates. All rights reserved.

The author hereby grants M.I.T. permission to reproduce and distributed publicly  
paper and electronic copies of this thesis and to grant others the right to do so.

Author .....

Department of Electrical Engineering and Computer Science  
16 May, 2003

Certified by .....

*M. Frans Kaashoek*  
M. Frans Kaashoek  
Professor of Computer Science and Engineering  
Thesis Supervisor


Certified by .....

Robert Morris  
Assistant Professor of Computer Science and Engineering  
Thesis Supervisor

Accepted by .....

Arthur C. Smith  
Chairman, Department Committee on Graduate Students

**BARKER**





# Robust and Efficient Data Management for a Distributed Hash Table

by

Josh Cates

Submitted to the Department of Electrical Engineering and Computer Science  
on 16 May, 2003, in partial fulfillment of the  
requirements for the degree of  
Master of Engineering in Computer Science and Engineering

## Abstract

This thesis presents a new design and implementation of the DHash distributed hash table based on erasure encoding. This design is both more robust and more efficient than the previous replication-based implementation [15].

DHash uses erasure coding to store each block as a set of fragments. Erasure coding increases availability while saving storage and communication costs compared to a replication based design. DHash combines Chord's synthetic coordinates with the the set of fragments to implement server selection on block retrieval.

DHash enhances robustness by implementing efficient fragment maintenance protocols. These protocols restore missing or misplaced fragments caused by hosts joining and leaving the system.

Experiments with a 270-node DHash system running on the PlanetLab [1] and RON [4] testbeds show that the changes to DHash increase the rate at which the system can fetch data by a factor of six, and decrease the latency of a single fetch by more than a factor of two. The maintenance protocols ensure that DHash is robust without penalizing performance. Even up to large database size, the per host memory footprint is less than 10 MB and the per host network bandwidth is under 2 KB/sec over a wide range of system half-lives.

Thesis Supervisor: M. Frans Kaashoek  
Title: Professor of Computer Science and Engineering

Thesis Supervisor: Robert Morris  
Title: Assistant Professor of Computer Science and Engineering



“Of the gladdest moments in human life, methinks, is the departure upon a distant journey into unknown lands. Shaking off with one mighty effort, the fetters of Habit, the leaden weight of Routine, the cloak of many cares and the slavery of Home, one feels once more happy.”

Sir Richard Burton - Journal Entry - 1856

# Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

## Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

## Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

## Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

## API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

## LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

## FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

## E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.