

A File System for Mobile Computing

Carl Downing Tait

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY
1993

© 1993

Carl Downing Tait

All Rights Reserved

ABSTRACT

A File System for Mobile Computing

Carl Downing Tait

Distributed file systems are a fundamental structure of distributed computing, and much attention has been focused on their design. But one important new design point has not yet been explored thoroughly: how to support mobile clients. Portable workstations are becoming increasingly common, and people who use these machines should not be forced to accept inferior performance or usability. This dissertation argues that two ideas — efficient variable-consistency replication and intelligent file prefetching — lead to a file system that supports mobile clients particularly well. Algorithms, implementations, and analyses are presented to support this assertion.

Contents

Table of Contents	i
List of Figures	iv
List of Tables	v
Acknowledgements	vi
1 Introduction	1
1.1 Thesis	2
1.2 Sub-Theses	3
1.2.1 Efficient Variable-Consistency Replication	3
1.2.2 Intelligent File Prefetching	5
1.3 Outline of the Dissertation	5
2 Background	6
2.1 Introduction	6
2.2 Naming	7
2.3 Replication	11
2.4 Caching	18
2.5 Summary	20
3 Efficient Variable-Consistency Replication	22
3.1 Introduction	22
3.2 Operation in the Absence of Failures	24
3.2.1 Dual-Read-Call Interface	26
3.2.2 Currency Tokens	28
3.2.3 Further Details	35
3.2.4 Client-Primary Attachment	37
3.2.5 Filesystem-Secondary Attachment	39
3.2.6 Connection at a Distant Site	41
3.2.7 Mobility at a Distant Site	42
3.3 Failure Recovery	42
3.3.1 Secondary Server Failure	43

3.3.2	Primary Server Failure	43
3.3.3	Client Failure	44
3.3.4	Reaction to Partition	44
3.3.5	Resolution of Conflicting Updates	45
3.3.6	Accommodation of Wandering Users	48
3.3.7	Semantics	48
3.4	Experimental Results	49
3.4.1	Experiments	49
3.4.2	Results	50
3.5	Prototype Implementation	54
3.5.1	Overview	55
3.5.2	Source Code	59
3.5.3	Restrictions	62
3.6	Comparison with Related Work	62
3.6.1	Performance	62
3.6.2	Resiliency	63
3.7	Conclusion	65
4	Intelligent File Prefetching	67
4.1	Introduction	67
4.2	The Algorithm	69
4.2.1	Data Structure — The Working Forest	69
4.2.2	Saving and Loading Trees	71
4.2.3	Tree Splitting	75
4.2.4	Common Prefix Trees	75
4.2.5	Cycle Trees	77
4.2.6	Prefetch Confidence	77
4.2.7	Garbage Collection	79
4.3	Simulation Results	80
4.3.1	Trace Data	80
4.3.2	Simulation Methodology	80
4.3.3	Results	82
4.3.4	Limitations	86
4.4	Implementation	88
4.4.1	Design Overview	89
4.4.2	NFS Overview	91
4.4.3	New Files	93
4.4.4	Modifications to Existing Files	95
4.4.5	Differences between the Simulation and the Implementation	98
4.4.6	Evaluation	100
4.4.7	Discussion	104
4.5	Related Work	105

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.