
Table of Contents

Foreword	ix
Preface	xi
I Introduction to Redundant Disk Array Architecture	1
1 A Case for Redundant Arrays of Inexpensive Disks (RAID)	3
<i>David Patterson, Garth Gibson, and Randy Katz</i>	
2 Disk System Architectures for High Performance Computing	15
<i>Randy H. katz, and Garth A. Gibson and D. Patterson</i>	
3 The Performance of Parity Placements in Disk Arrays	35
<i>Edward K. Lee and Randy H. Katz</i>	
4 A Performance Comparison of RAID-5 and Log-Structured Arrays	55
<i>Jai Menon</i>	
II Advanced Disk Array Architectures	65
5 Parity Logging Overcoming the Small Write Problem in Redundant Disk Arrays	67
<i>Daniel Stodolsky, Garth Gibson, and Mark Holland</i>	
6 Distributed RAID - A New Multiple Copy Algorithm	81
<i>Michael Stonebraker and Gerhard A. Schloss</i>	
7 The HP AutoRAID Hierarchical Storage System	90
<i>John Wilkes, Richard Golding, Carl Staelin, and Tim Sullivan</i>	
8 Scalable Distributed Log Structured Arrays	107
<i>Witold Litwin and Jai Menon</i>	
9 Comparison of Sparing Alternatives for Disk Arrays	117
<i>Jai Menon</i>	
10 Destage Algorithms for Disk Arrays with Non-Volatile Caches	129
<i>Anujan Varma and Quinn Jacobson</i>	
III Fault Tolerance Issues in Disk Arrays	147
11 Failure Correction Techniques for Large Disk Arrays	149
<i>Garth A. Gibson, Lisa Hellerstein, Richard M. Karp, Randy H. Katz and David A. Patterson</i>	

12	Tolerating Multiple Failures in RAID Architectures with Optimal Storage and Uniform Declustering	161
	<i>Guillermo A. Alvarez, Walter A. Burkhard, and Flaviu Cristian</i>	
13	Parity Declustering for Continuous Operation in Redundant Disk Arrays	173
	<i>Mark Holland and Garth A. Gibson</i>	
14	The EVENODD Code and its Generalization	187
	<i>Mario Blaum, Jim Brady, Jehoshua Bruck, Jai Menon, and Alexander Vardy</i>	
IV	Caching and Prefetching	209
15	RAPID-Cache — A Reliable and Inexpensive Write Cache for Disk I/O Systems	211
	<i>Yiming Hu, Qing Yang, and Tycho Nightingale</i>	
16	Informed Prefetching and Caching	224
	<i>R. Hugo Patterson, Garth A. Gibson, Eka Ginting, Daniel Stodolsky, and Jim Zelenka</i>	
17	Practical Prefetching Techniques for Multiprocessor File Systems	246
	<i>David Kotz and Carla Schlatter Ellis</i>	
18	Design Issues of a Cooperative Cache with no Coherence Problems	260
	<i>Toni Cortes, Sergi Girona and Jesús Labarta</i>	
19	Collective Buffering: Improving Parallel I/O Performance	272
	<i>Bill Nitzberg and Virginia Lo</i>	
V	Parallel File Systems	283
20	The Vesta Parallel File System	285
	<i>Peter F. Corbett and Dror G. Feitelson</i>	
21	The Zebra Striped Network File System	309
	<i>John H. Hartman and John K. Ousterhout</i>	
22	PPFS: A High Performance Portable Parallel File System	330
	<i>James V. Huber, Jr., Christopher L. Elford, Daniel A. Reed, Andrew A. Chien, and David S. Blumenthal</i>	
23	The Global File System	344
	<i>Steven R. Soltis, Thomas M. Ruwart, Grant M. Erickson, Kenneth W. Preslan, and Matthew T. O'Keefe</i>	
24	Serverless Network File Systems	364
	<i>Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe Matthews, David A. Patteson, Drew S. Roselli, and Randolph Y. Wang</i>	
VI	Parallel I/O Systems	387
25	Parallel I/O Subsystems in Massively Parallel Supercomputers	389
	<i>Dror G. Feitelson, Peter F. Corbett, Sandra Johnson Baylor, and Yarsun Hsu</i>	
26	RAID-II: A High-Bandwidth Network File Server	408
	<i>Ann L. Chervenak, Ken Shirriff, John H. Hartman, Ethan L. Miller, Srinivasan Seshan, Randy H. Katz, Ken Lutz, David A. Patterson, Edward K. Lee, Peter M. Chen, and Garth A. Gibson</i>	
27	Petal: Distributed Virtual Disks	420
	<i>Edward K. Lee and Chandramohan A. Thekkath</i>	

28	A Cost-Effective, High-Bandwidth Storage Architecture	431
	<i>Garth A. Gibson, David F. Nagle, Khalil Amiri, Jeff Butler, Fay W. Chang, Howard Gobioff, Charles Hardin, Erik Riedel, David Rochberg, and Jim Zelenka</i>	
29	RAID-x: A New Distributed Disk Array for I/O-Centric Cluster Computing	445
	<i>Kai Hwang, Hai Jin and Roy S. C. Ho</i>	
30	Designing a Self-Maintaining Storage System	453
	<i>Satoshi Asami, Nisha Talagala, and David A. Patterson</i>	
31	Modeling and Evaluation of Fibre Channel Storage Area Networks	464
	<i>Xavier Molero, Federico Silla, Vicente Santonja and JoséDuato</i>	
VII Parallel I/O Programming Paradigms		475
32	Overview of the MPI-IO Parallel I/O Interface	477
	<i>Peter Corbett, Dror Feitelson, Sam Fineberg, Yarsun Hsu, Bill Nitzberg, Jean-Pierre Prost, Marc Snir, Bernard Traversat, and Parkson Wong</i>	
33	Disk Resident Arrays: An Array-Oriented I/O Library for Out-of-Core Computations	488
	<i>Ian Foster and Jarek Nieplocha</i>	
34	Active Disks: Programming Model, Algorithms and Evaluation	499
	<i>Anurag Acharya, Mustafa Uysal, and Joel Saltz</i>	
35	Disk-directed I/O for MIMD Multiprocessors	513
	<i>David Kotz</i>	
VIII Parallel I/O Applications and Environments		537
36	Applications-Driven Parallel I/O	539
	<i>Nicholas P. Galbreath, William D. Gropp, and David M. Levine</i>	
37	Comparing Multimedia Storage Architectures	548
	<i>Benoit A. Gennart and Roger D. Hersch</i>	
38	High Availability in Clustered Multimedia Servers	555
	<i>Renu Tewari, Daniel M. Dias, Rajat Mukherjee, and Harrick M. Vin</i>	
39	An Architecture for a Scalable High-Performance Digital Library	566
	<i>R. Grossman, X. Qin, W. Xu, H. Hulen, and T. Tyler</i>	
40	I/O Requirements of Scientific Applications: An Evolutionary View	576
	<i>Evgenia Smirni, Ruth A. Aydt, Andrew A. Chien, Daniel A. Reed</i>	
41	Mitra: A Scalable Continuous Media Server	595
	<i>Shahram Ghandeharizadeh, Roger Zimmermann, Weifeng Shi, Reza Rejaie, Douglas J. Ierardi, and Ta-Wei Li</i>	
IX Emerging Technologies and Future Trends		615
42	An Introduction to the InfiniBand Architecture	617
	<i>Gregory F. Pfister</i>	
43	XML, Hyper-media, and Fortran I/O	633
	<i>Dror G. Feitelson and Tomer Klainer</i>	
44	I/O Programming Paradigms: Past and Future	645
	<i>Mahmut Taylan Kandemir and Alok Choudhary</i>	

45 Scientific Applications using Parallel I/O 655
Ron Oldfield and David Kotz

Index 667