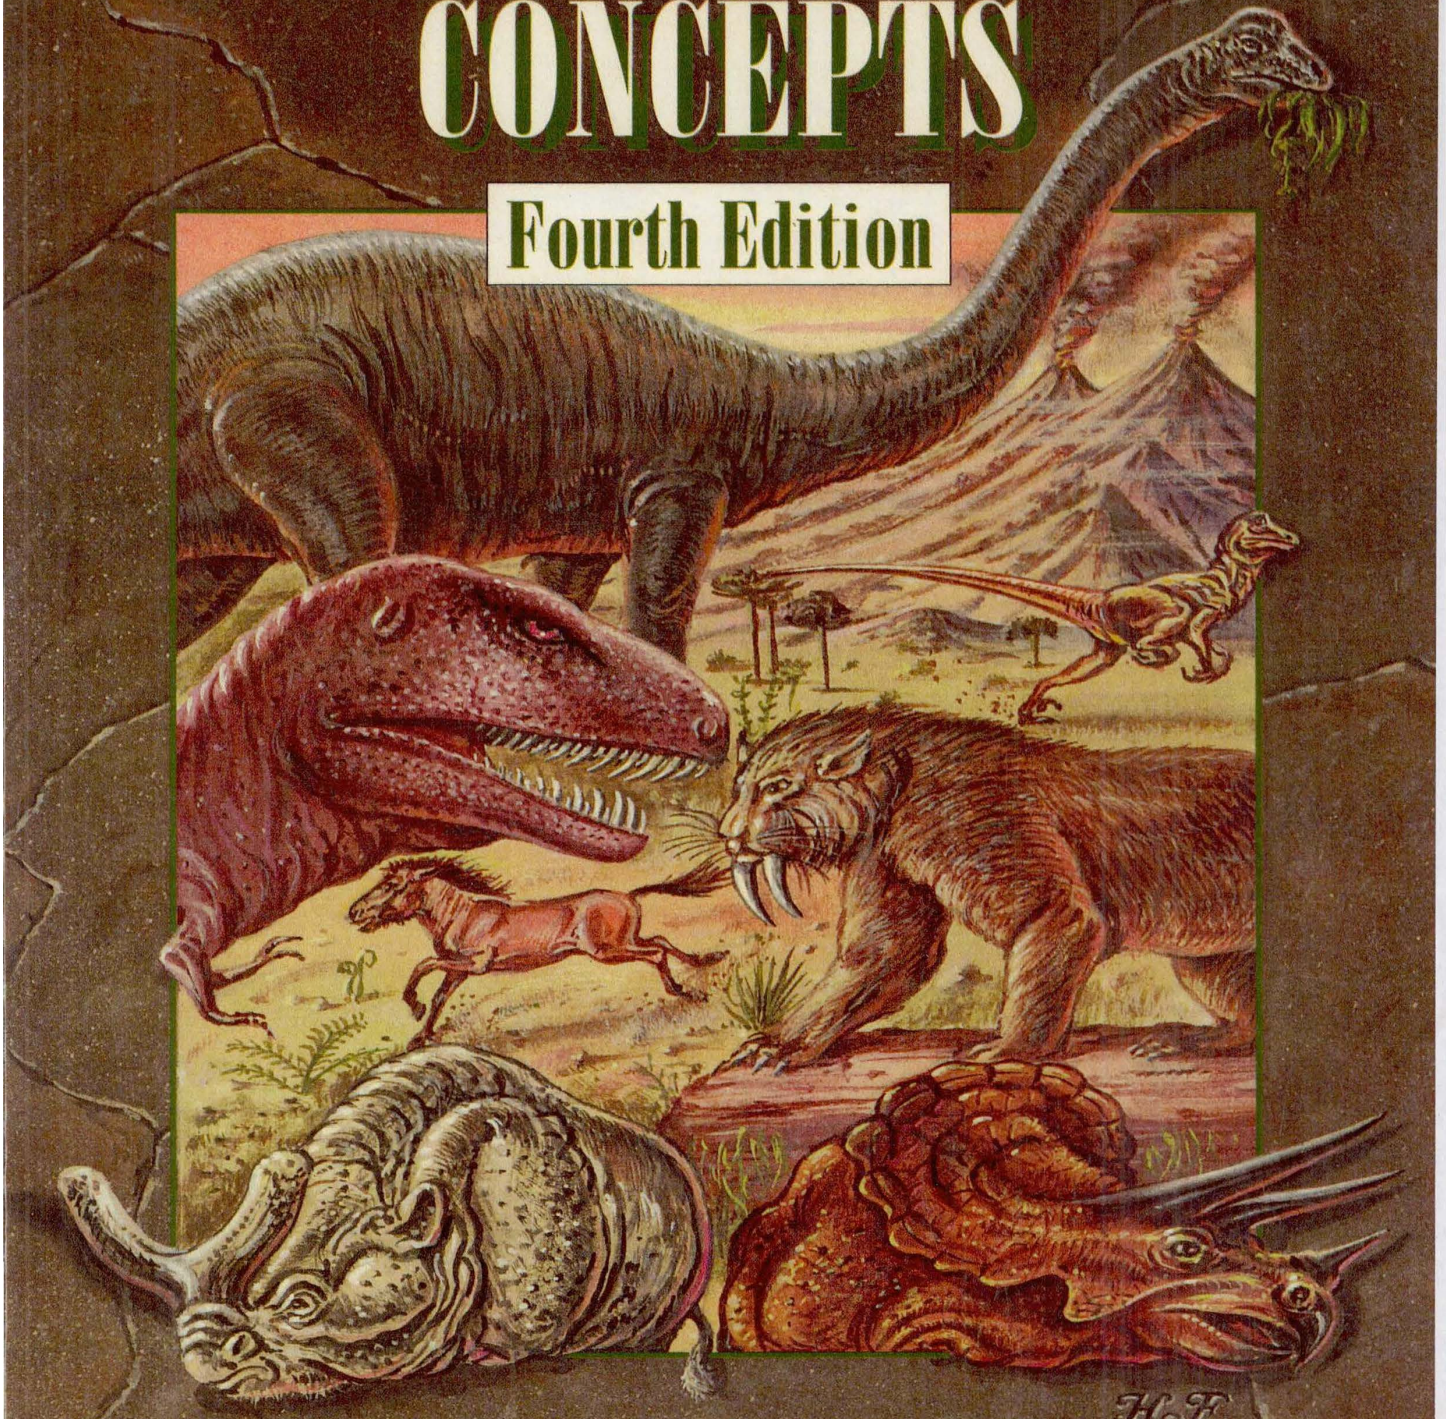
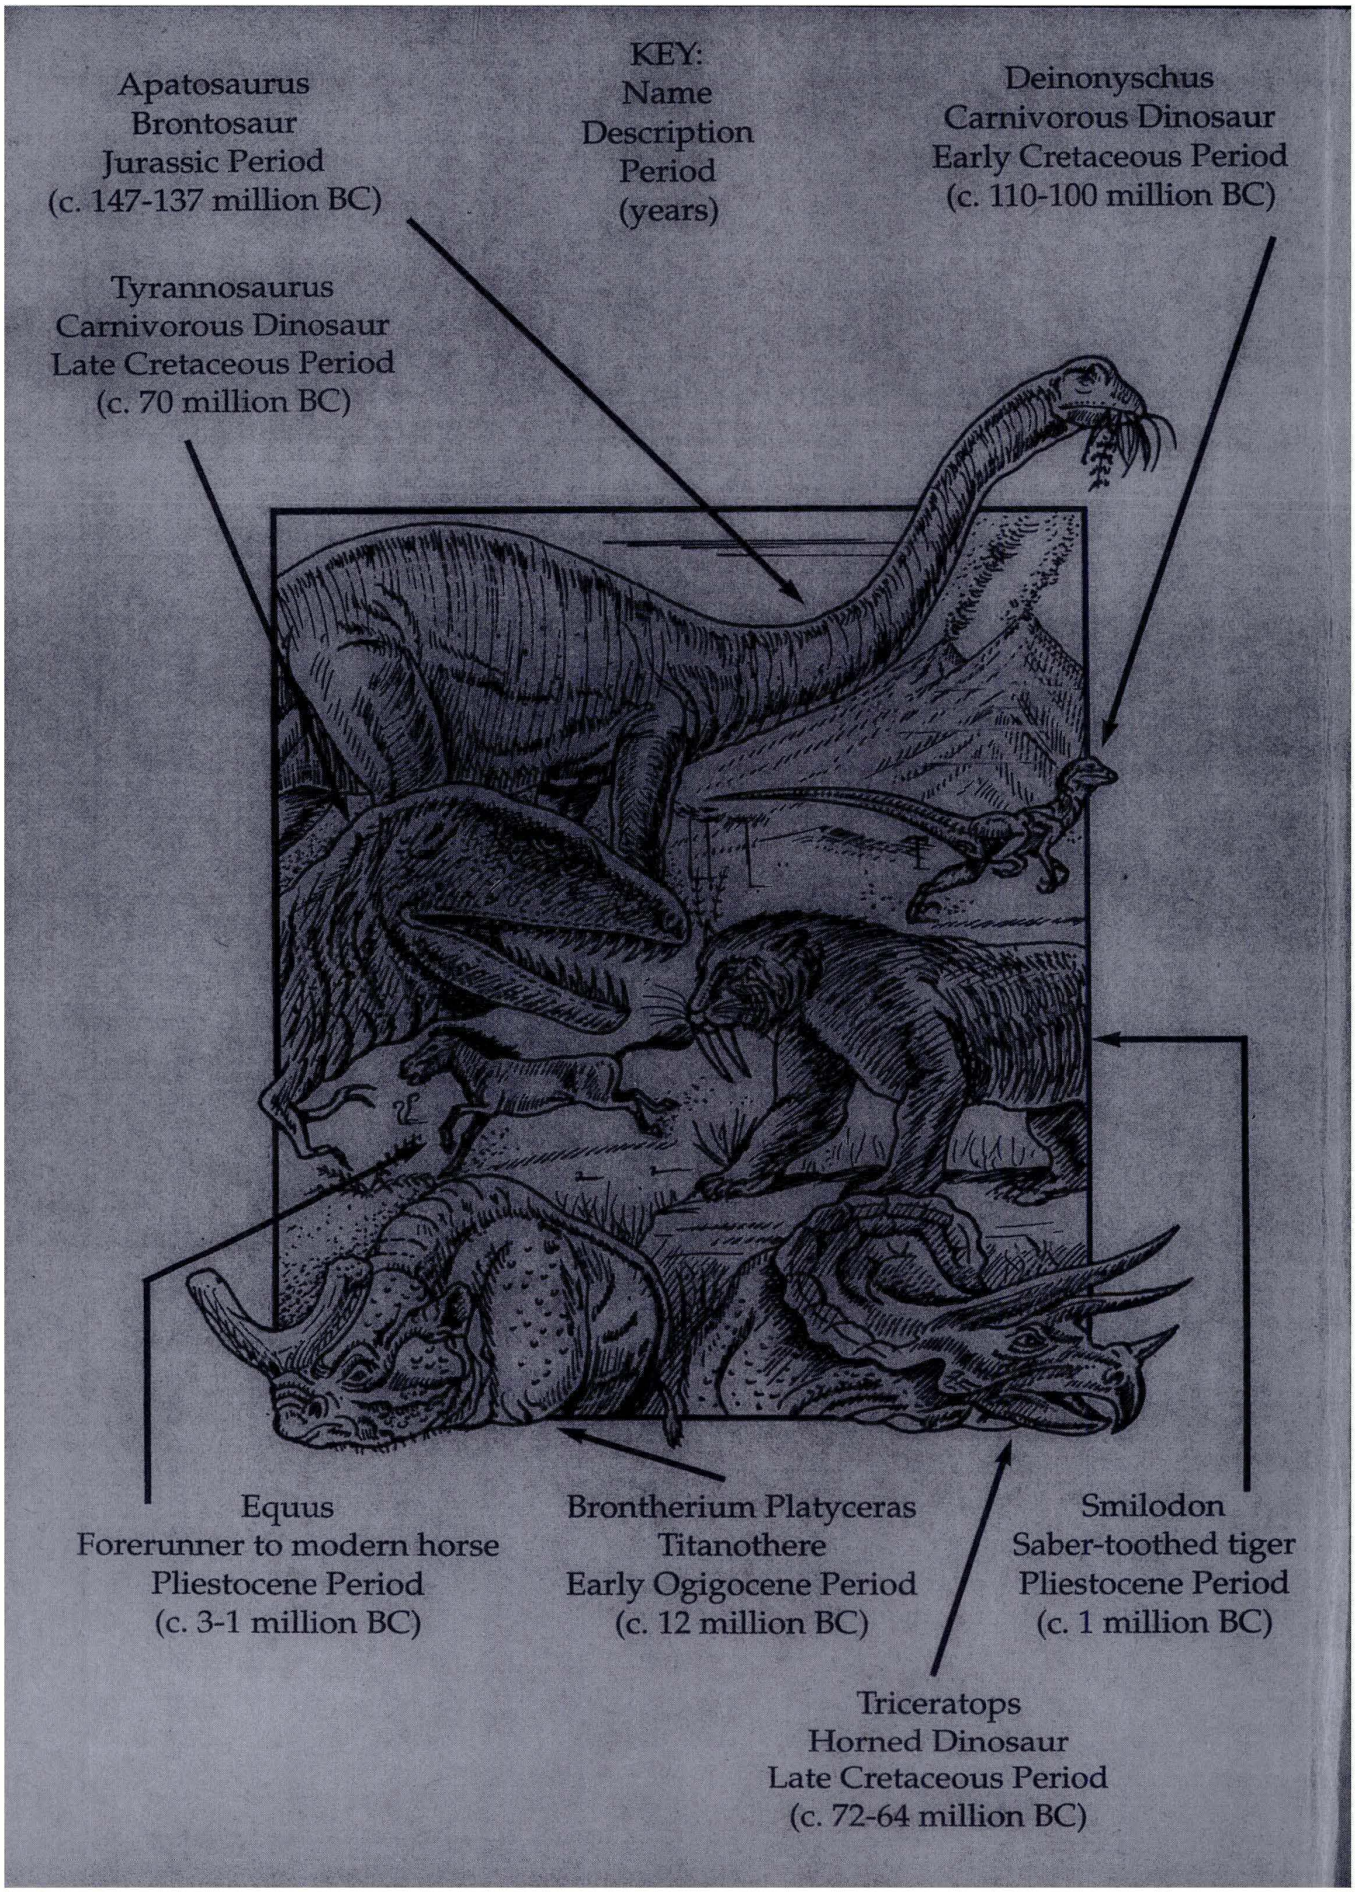
OPERATING SYSTEM CONCEPTS

Fourth Edition



SILBERSCHATZ
GALVIN

SILBERSCHATZ
GALVIN
OPERATING SYSTEM CONCEPTS
4-975-321
BookHolders.com


KEY:
Name
Description
Period
(years)
Apatosaurus
Brontosaur
Jurassic Period
(c. 147-137 million BC)
Deinonyschus
Carnivorous Dinosaur
Early Cretaceous Period
(c. 110-100 million BC)
Tyrannosaurus
Carnivorous Dinosaur
Late Cretaceous Period
(c. 70 million BC)
Equus
Forerunner to modern horse
Pliestocene Period
(c. 3-1 million BC)
Brontherium Platyceras
Titanothere
Early Ogigocene Period
(c. 12 million BC)
Smilodon
Saber-toothed tiger
Pliestocene Period
(c. 1 million BC)
Triceratops
Horned Dinosaur
Late Cretaceous Period
(c. 72-64 million BC)

FOURTH EDITION



OPERATING SYSTEM CONCEPTS

Abraham Silberschatz
University of Texas

Peter B. Galvin
Brown University

Addison-Wesley Publishing Company
Reading, Massachusetts • Menlo Park, California • New York
Don Mills, Ontario • Wokingham, England • Amsterdam • Bonn
Sydney • Singapore • Tokyo • Madrid • San Juan • Milan • Paris

Sponsoring Editor: *Deborah Lafferty*
Senior Editor: *Tom Stone*
Senior Production Supervisor: *Helen Wythe*
Marketing Manager: *Phyllis Cerys*
Technical Art Coordinator: *Susan London-Payne*
Cover and Endpaper Designer: *Howard S. Friedman*
Manufacturing Manager: *Roy Logan*

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Addison-Wesley was aware of a trademark claim, the designations have been printed in initial caps or all caps.

The procedures and applications presented in this book have been included for their instructional value. They have been tested with care but are not guaranteed for any particular purpose. The publisher does not offer any warranties or representations, nor does it accept any liabilities with respect to the programs or applications.

Library of Congress Cataloging-in-Publication Data

Silberschatz, Abraham.
Operating system concepts / Abraham Silberschatz, Peter B. Galvin.
p. cm.
Includes bibliographical references and index.
ISBN 0-201-50480-4
1. Operating systems (Computers) I. Galvin, Peter B. II. Title.
QA76.76.063S5583 1994
005.4'3--dc20 93-24415
CIP

Reprinted with corrections January, 1995

Reproduced by Addison-Wesley from camera-ready copy supplied by the authors.

Copyright © 1994 by Addison-Wesley Publishing Company, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America.

ISBN 0-201-50480-4
8 9 10-MA-99 98 97 96

To my parents, Wira and Mietek,
my wife, Haya,
and my children, Lemor, Sivan and Aaron.

Avi Silberschatz

To Carla and Gwendolyn.

Peter Galvin

CONTENTS

PART ONE ■ OVERVIEW

Chapter 1 Introduction

1.1 What Is an Operating System? 3
1.2 Early Systems 6
1.3 Simple Batch Systems 7
1.4 Multiprogrammed Batched Systems 13
1.5 Time-Sharing Systems 15
1.6 Personal-Computer Systems 17
1.7 Parallel Systems 20
1.8 Distributed Systems 22
1.9 Real-Time Systems 23
1.10 Summary 25
Exercises 26
Bibliographic Notes 27

Chapter 2 Computer-System Structures

2.1 Computer-System Operation 29
2.2 I/O Structure 32
2.3 Storage Structure 37
2.4 Storage Hierarchy 42
2.5 Hardware Protection 45
2.6 General-System Architecture 51
2.7 Summary 52
Exercises 53
Bibliographic Notes 55

Chapter 3 Operating-System Structures

3.1 System Components 57
3.2 Operating-System Services 63
3.3 System Calls 65
3.4 System Programs 74
3.5 System Structure 76
3.6 Virtual Machines 82

HULU EXHIBIT 1018
Page 7 of 22

3.7 System Design and Implementation 86
3.8 System Generation 89
3.9 Summary 90
Exercises 91
Bibliographic Notes 92

PART TWO ■ PROCESS MANAGEMENT

Chapter 4 Processes

4.1 Process Concept 97
4.2 Process Scheduling 100
4.3 Operation on Processes 105
4.4 Cooperating Processes 108
4.5 Threads 111
4.6 Interprocess Communication 116
4.7 Summary 126
Exercises 127
Bibliographic Notes 129

Chapter 5 CPU Scheduling

5.1 Basic Concepts 131
5.2 Scheduling Criteria 135
5.3 Scheduling Algorithms 137
5.4 Multiple-Processor Scheduling 149
5.5 Real-Time Scheduling 150
5.6 Algorithm Evaluation 152
5.7 Summary 158
Exercises 159
Bibliographic Notes 161

Chapter 6 Process Synchronization

6.1 Background 163
6.2 The Critical-Section Problem 165
6.3 Synchronization Hardware 172
6.4 Semaphores 175
6.5 Classical Problems of Synchronization 181
6.6 Critical Regions 186
6.7 Monitors 190
6.8 Synchronization in Solaris 2 198
6.9 Atomic Transactions 199
6.10 Summary 208
Exercises 210
Bibliographic Notes 214

Chapter 7 Deadlocks

7.1 System Model 217
7.2 Deadlock Characterization 219
7.3 Methods for Handling Deadlocks 223
7.4 Deadlock Prevention 224
7.5 Deadlock Avoidance 227
7.6 Deadlock Detection 234
7.7 Recovery from Deadlock 238

7.8 Combined Approach to Deadlock Handling 240
7.9 Summary 241
Bibliographic Notes 245
Exercises 242

PART THREE ■ STORAGE MANAGEMENT

Chapter 8 Memory Management

8.1 Background 249
8.2 Logical versus Physical Address Space 255
8.3 Swapping 256
8.4 Contiguous Allocation 259
8.5 Paging 267
8.6 Segmentation 283
8.7 Segmentation with Paging 290
8.8 Summary 294
Exercises 296
Bibliographic Notes 299

Chapter 9 Virtual Memory

9.1 Background 301
9.2 Demand Paging 303
9.3 Performance of Demand Paging 309
9.4 Page Replacement 312
9.5 Page-Replacement Algorithms 315
9.6 Allocation of Frames 326
9.7 Thrashing 329
9.8 Other Considerations 334
9.9 Demand Segmentation 341
9.10 Summary 342
Exercises 343
Bibliographic Notes 348

Chapter 10 File-System Interface

10.1 File Concept 349
10.2 Access Methods 358
10.3 Directory Structure 361
10.4 Protection 373
10.5 Consistency Semantics 378
10.6 Summary 379
Exercises 380
Bibliographic Notes 381

Chapter 11 File-System Implementation

11.1 File-System Structure 383
11.2 Allocation Methods 387
11.3 Free-Space Management 397
11.4 Directory Implementation 399
11.5 Efficiency and Performance 401
11.6 Recovery 403
11.7 Summary 405
Exercises 406
Bibliographic Notes 408

Chapter 12 Secondary-Storage Structure

12.1 Disk Structure 409
12.2 Disk Scheduling 410
12.3 Disk Management 417
12.4 Swap-Space Management 419
12.5 Disk Reliability 422
12.6 Stable-Storage Implementation 424
12.7 Summary 425
Exercises 426
Bibliographic Notes 427

PART FOUR ■ PROTECTION AND SECURITY

Chapter 13 Protection

13.1 Goals of Protection 431
13.2 Domain of Protection 432
13.3 Access Matrix 438
13.4 Implementation of Access Matrix 443
13.5 Revocation of Access Rights 446
13.6 Capability-Based Systems 448
13.7 Language-Based Protection 451
13.8 Summary 455
Exercises 455
Bibliographic Notes 457

Chapter 14 Security

14.1 The Security Problem 459
14.2 Authentication 461
14.3 Program Threats 464
14.4 System Threats 465
14.5 Threat Monitoring 469
14.6 Encryption 471
14.7 Summary 473
Exercises 473
Bibliographic Notes 474

PART FIVE ■ DISTRIBUTED SYSTEMS

Chapter 15 Network Structures

15.1 Background 479
15.2 Motivation 481
15.3 Topology 482
15.4 Network Types 488
15.5 Communication 491
15.6 Design Strategies 498
15.7 Networking Example 501
15.8 Summary 504
Exercises 504
Bibliographic Notes 505

Chapter 16 Distributed-System Structures

16.1 Network-Operating Systems 507
16.2 Distributed-Operating Systems 509
16.3 Remote Services 512
16.4 Robustness 517
16.5 Design Issues 519
16.6 Summary 521
Exercises 522
Bibliographic Notes 523

Chapter 17 Distributed-File Systems

17.1 Background 525
17.2 Naming and Transparency 527
17.3 Remote File Access 531
17.4 Stateful versus Stateless Service 536
17.5 File Replication 538
17.6 Example Systems 539
17.7 Summary 567
Exercises 568
Bibliographic Notes 569

Chapter 18 Distributed Coordination

18.1 Event Ordering 571
18.2 Mutual Exclusion 574
18.3 Atomicity 577
18.4 Concurrency Control 581
18.5 Deadlock Handling 586
18.6 Election Algorithms 595
18.7 Reaching Agreement 598
18.8 Summary 600
Exercises 601
Bibliographic Notes 602

PART SIX ■ CASE STUDIES

Chapter 19 The UNIX System

19.1 History 607
19.2 Design Principles 613
19.3 Programmer Interface 615
19.4 User Interface 623
19.5 Process Management 627
19.6 Memory Management 632
19.7 File System 636
19.8 I/O System 645
19.9 Interprocess Communication 649
19.10 Summary 655
Exercises 655
Bibliographic Notes 657

Chapter 20 The Mach System

20.1 History 659
20.2 Design Principles 661
20.3 System Components 662
20.4 Process Management 666
20.5 Interprocess Communication 673
20.6 Memory Management 679
20.7 Programmer Interface 685
20.8 Summary 686
Exercises 687
Bibliographic Notes 688

Chapter 21 Historical Perspective

21.1 Atlas 691
21.2 XDS-940 692
21.3 THE 693
21.4 RC 4000 694
21.5 CTSS 695
21.6 MULTICS 696
21.7 OS/360 696
21.8 Other Systems 698

Appendix The Nachos System

A.1 Overview 700
A.2 Nachos Software Structure 702
A.3 Sample Assignments 705
A.4 Information on Obtaining a Copy of Nachos 711
A.5 Conclusions 713
Bibliographic Notes 713

Bibliography 715

Credits 745

Index 747

HULU EXHIBIT 1018
Page 12 of 22

their performance is increasing. Because they are more rugged, they, like floppy disks, have the added advantage of removability.

2.3.4 Magnetic Tapes

Magnetic tape was used as an early secondary-storage media. Although it is relatively permanent, and can hold large numbers of data, magnetic tape is quite slow in comparison to the access time of main memory. Even more important, magnetic tape is limited to sequential access. Thus, it is unsuitable for providing the random access needed for most secondary-storage requirements. Tapes are used mainly for backup, for storage of infrequently used information, and as a medium for transferring information from one system to another.

A tape is kept in a spool, and is wound or rewound past a read–write head. Moving to the correct spot on a tape can take minutes rather than milliseconds; once positioned, however, tape drives can write data at densities and speeds approaching those of disk drives. Capacities vary depending on the length and width of the tape, and on the density at which the head can read and write. A tape drive is usually named by its width. Thus, there are 8-millimeter, 1/4-inch, and 1/2-inch (also known as 9-track) tape drives. The 8-millimeter tape drives have the highest density, due to the technology they use; they currently store 5 gigabytes of data (a gigabyte is one billion bytes) on a 350-foot tape.

2.4 ■ Storage Hierarchy

The wide variety of storage systems in a computer system can be organized in a hierarchy (Figure 2.6) according to their speed and their cost. The higher levels are expensive, but are fast. As we move down the hierarchy, the cost per bit decreases, whereas the access time increases. This tradeoff is reasonable; if a given storage system were both faster and less expensive than another — other properties being the same — then there would be no reason to use the slower, more expensive memory. In fact, many early storage devices, including paper tape and core memories, are relegated to museums now that magnetic tape and semiconductor memory have become faster and cheaper.

In addition to the speed and cost of the various storage systems, there is also the issue of storage volatility. Volatile storage loses its contents when the power to the device is removed. In the absence of expensive battery and generator backup systems, data must be written to nonvolatile storage for safekeeping. In the hierarchy shown in Figure 2.6, the storage system above disks is volatile, whereas the storage system below main memory is nonvolatile. The design of a complete memory system attempts to balance all these factors: It uses only as much expensive memory as

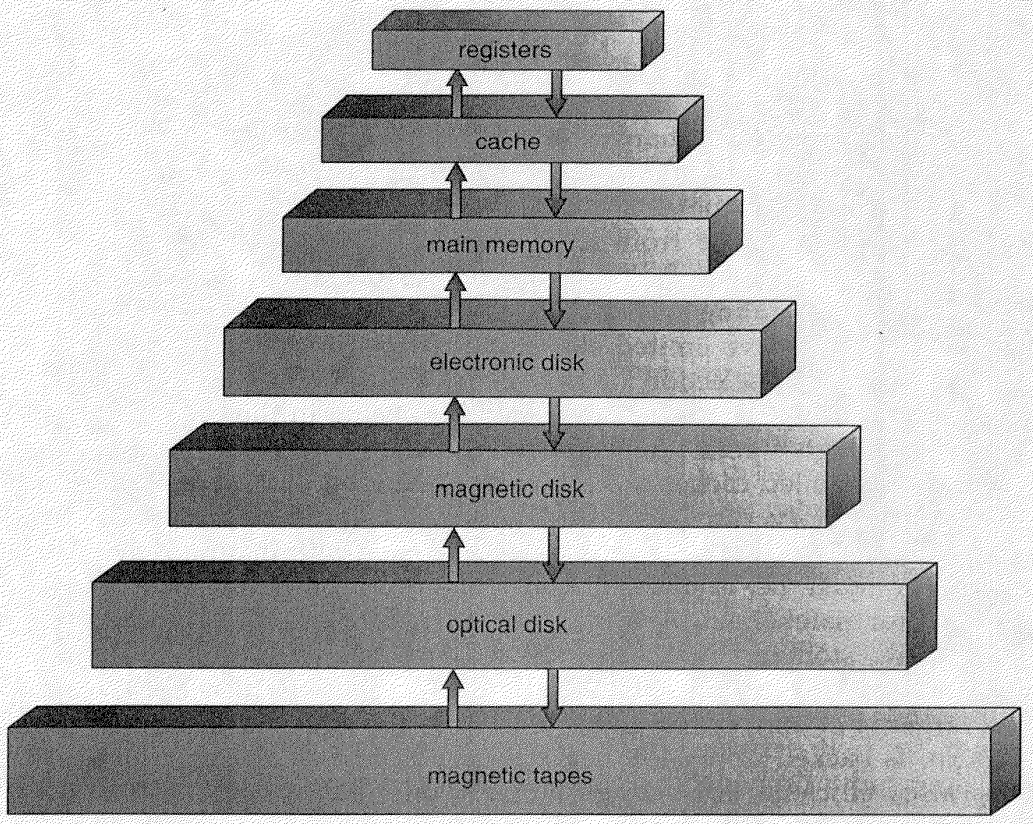


Figure 2.6 Storage-device hierarchy.

necessary, while providing as much inexpensive, nonvolatile, memory as possible. Caches can be installed to ameliorate performance differences where there is a large access-time or transfer-rate disparity between two components.

2.4.1 Caching

Caching is an important principle of computer systems, in both hardware and software. Information is normally kept in some storage system (such as main memory). As it is used, it is copied into a faster storage system, the cache, on a temporary basis. When we need a particular piece of information, we first check whether it is in the cache. If it is, we use the information directly from the cache; if it is not, we use the information from the main storage system, putting a copy in the cache under the assumption that there is a high probability that it will be needed again.

Extending this view, internal programmable registers, such as index registers and accumulators, are a high-speed cache for main memory. The

programmer (or compiler) implements the register-allocation and register-replacement algorithms to decide which information to keep in registers and which to keep in main memory. There are also caches that are implemented totally in hardware. For instance, most systems have an instruction cache to hold the next instructions expected to be executed. Without this cache, the CPU would have to wait several cycles while an instruction was fetched from main memory. We are not concerned with these hardware-only caches in this text, since they are outside of the control of the operating system.

Since caches have limited size, *cache management* is an important design problem. Careful selection of the cache size and of a replacement policy can result in 80 to 99 percent of all accesses being in the cache, resulting in extremely high performance. Various replacement algorithms for software-controlled caches are discussed in Chapter 9.

Main memory can be viewed as a fast *cache* for secondary memory, since data on secondary storage must be copied into main memory for use, and data must be in main memory before being moved to secondary storage for safekeeping. The file system itself, which must reside on nonvolatile storage, may have several levels of storage. At the highest level, we have the electronic (or RAM) disk storage, which is backed up by the larger, but slower, magnetic-disk storage. The magnetic-disk storage, in turn, is backed up by the larger, but slower, tape storage. Optical disks are also efficient high-capacity but low-cost storage media. When compared to magnetic tape, they have the drawback of higher cost, but offer much greater speed and convenience. Transfers between these two storage levels are generally requested explicitly, but some systems now automatically archive a file that has not been used for a long time (for example, 1 month), and then automatically fetch back the file to disk when it is next referenced.

The movement of information between levels of a storage hierarchy may be either explicit or implicit, depending on the hardware design and the controlling operating-system software. For instance, data transfer from cache to CPU and registers is usually a hardware function, with no operating-system intervention. On the other hand, transfer of data from disk to memory is usually controlled by the operating system.

2.4.2 Coherency and Consistency

In a hierarchical storage structure, the same datum may appear in different storage systems. For example, consider an integer A located in file B that is to be incremented by 1. Suppose that file B resides on magnetic disk. The increment operation proceeds by first issuing an I/O operation to copy the disk block on which A resides to main memory. This operation is followed by a possible copying of A to the cache, and by copying A to an internal register. Thus, the copy of A appears in several places. Once the

algorithm will never suffer from Belady's anomaly. An optimal page-replacement algorithm exists, and has been called OPT or MIN. It is simply

> Replace the page that will not be used
> for the longest period of time.

Use of this page-replacement algorithm guarantees the lowest possible page-fault rate for a fixed number of frames.

For example, on our sample reference string, the optimal page-replacement algorithm would yield nine page faults, as shown in Figure 9.10. The first three references cause faults that fill the three empty frames. The reference to page 2 replaces page 7, because 7 will not be used until reference 18, whereas page 0 will be used at 5, and page 1 at 14. The reference to page 3 replaces page 1, as page 1 will be the last of the three pages in memory to be referenced again. With only nine page faults, optimal replacement is much better than a FIFO algorithm, which had 15 faults. (If we ignore the first three, which all algorithms must suffer, then optimal replacement is twice as good as FIFO replacement.) In fact, no replacement algorithm can process this reference string in three frames with less than nine faults.

Unfortunately, the optimal page-replacement algorithm is difficult to implement, because it requires future knowledge of the reference string. (We encountered a similar situation with the SJF CPU-scheduling algorithm in Section 5.3.2.) As a result, the optimal algorithm is used mainly for comparison studies. For instance, it may be quite useful to know that, although a new algorithm is not optimal, it is within 12.3 percent of optimal at worst and within 4.7 percent on average.

9.5.3 LRU Algorithm

If the optimal algorithm is not feasible, perhaps an approximation to the optimal algorithm is possible. The key distinction between the FIFO and OPT algorithms (other than looking backward or forward in time) is that the

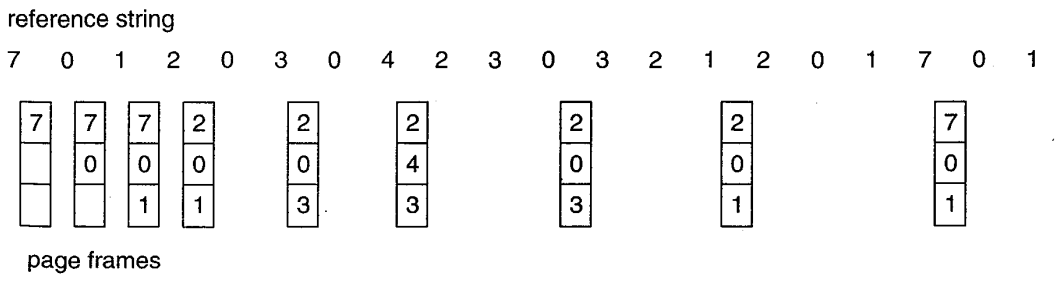


Figure 9.10 Optimal page-replacement algorithm.

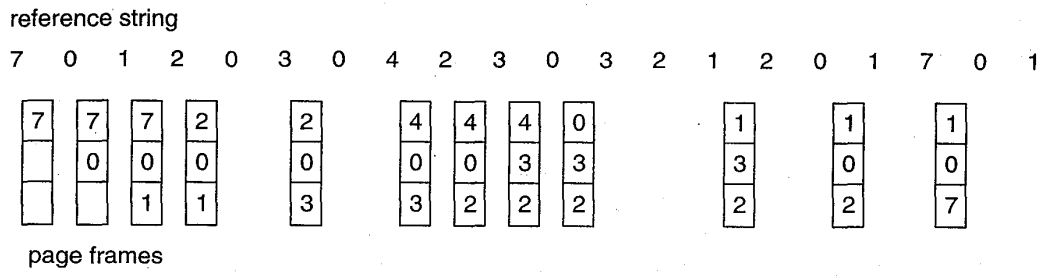


Figure 9.11 LRU page-replacement algorithm.

FIFO algorithm uses the time when a page was brought into memory; the OPT algorithm uses the time when a page is to be *used*. If we use the recent past as an approximation of the near future, then we will replace the page that *has not been used* for the longest period of time (Figure 9.11). This approach is the *least recently used (LRU)* algorithm.

LRU replacement associates with each page the time of that page's last use. When a page must be replaced, LRU chooses that page that has not been used for the longest period of time. This strategy is the optimal page-replacement algorithm looking backward in time, rather than forward. (Strangely, if we let S^R be the reverse of a reference string S, then the page-fault rate for the OPT algorithm on S is the same as the page-fault rate for the OPT algorithm on S^R. Similarly, the page-fault rate for the LRU algorithm on S is the same as the page-fault rate for the LRU algorithm on S^R.)

The result of applying LRU replacement to our example reference string is shown in Figure 9.11. The LRU algorithm produces 12 faults. Notice that the first five faults are the same as the optimal replacement. When the reference to page 4 occurs, however, LRU replacement sees that, of the three frames in memory, page 2 was used least recently. The most recently used page is page 0, and just before that page 3 was used. Thus, the LRU algorithm replaces page 2, not knowing that page 2 is about to be used. When it then faults for page 2, the LRU algorithm replaces page 3 since, of the three pages in memory {0, 3, 4}, page 3 is the least recently used. Despite these problems, LRU replacement with 12 faults is still much better than FIFO replacement with 15.

The LRU policy is often used as a page-replacement algorithm and is considered to be quite good. The major problem is *how* to implement LRU replacement. An LRU page-replacement algorithm may require substantial hardware assistance. The problem is to determine an order for the frames defined by the time of last use. Two implementations are feasible:

- **Counters:** In the simplest case, we associate with each page-table entry a time-of-use field, and add to the CPU a logical clock or counter. The

HULU EXHIBIT 1018
Page 17 of 22

clock is incremented for every memory reference. Whenever a reference to a page is made, the contents of the clock register are copied to the time-of-use field in the page table for that page. In this way, we always have the "time" of the last reference to each page. We replace the page with the smallest time value. This scheme requires a search of the page table to find the LRU page, and a write to memory (to the time-of-use field in the page table) for each memory access. The times must also be maintained when page tables are changed (due to CPU scheduling). Overflow of the clock must be considered.

- **Stack:** Another approach to implementing LRU replacement is to keep a *stack* of page numbers. Whenever a page is referenced, it is removed from the stack and put on the top. In this way, the top of the stack is always the most recently used page and the bottom is the LRU page (Figure 9.12). Because entries must be removed from the middle of the stack, it is best implemented by a doubly linked list, with a head and tail pointer. Removing a page and putting it on the top of the stack then requires changing six pointers at worst. Each update is a little more expensive, but there is no search for a replacement; the tail pointer points to the bottom of the stack, which is the LRU page. This approach is particularly appropriate for software or microcode implementations of LRU replacement.

Neither optimal replacement nor LRU replacement suffers from Belady's anomaly. There is a class of page-replacement algorithms, called *stack algorithms*, that can never exhibit Belady's anomaly. A stack algorithm is an algorithm for which it can be shown that the set of pages in memory for n

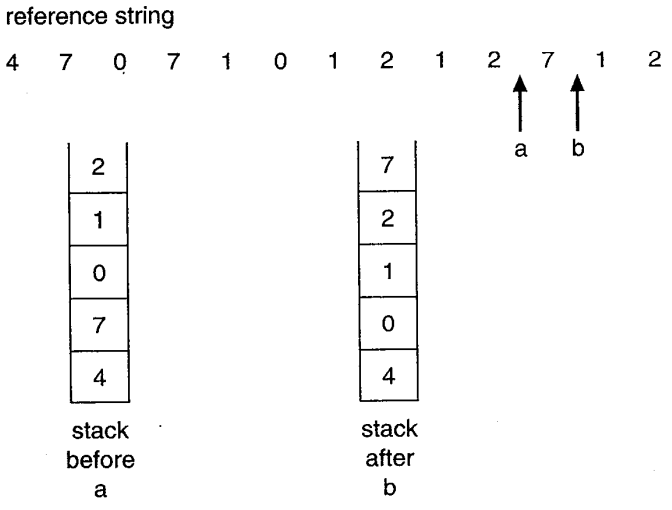


Figure 9.12 Use of a stack to record the most recent page references.

frames is always a *subset* of the set of pages that would be in memory with $n + 1$ frames. For LRU replacement, the set of pages in memory would be the n most recently referenced pages. If the number of frames is increased, these n pages will still be the most recently referenced and so will still be in memory.

Note that neither implementation of LRU would be conceivable without hardware assistance beyond the standard TLB registers. The updating of the clock fields or stack must be done for *every* memory reference. If we were to use an interrupt for every reference, to allow software to update such data structures, it would slow every memory reference by a factor of at least 10, hence slowing every user process by a factor of 10. Few systems could tolerate that level of overhead for memory management.

9.5.4 LRU Approximation Algorithms

Few systems provide sufficient hardware support for true LRU page replacement. Some systems provide no hardware support, and other page-replacement algorithms (such as a FIFO algorithm) must be used. Many systems provide some help, however, in the form of a *reference bit*. The reference bit for a page is set, by the hardware, whenever that page is referenced (either a read or a write to any byte in the page). Reference bits are associated with each entry in the page table.

Initially, all bits are cleared (to 0) by the operating system. As a user process executes, the bit associated with each page referenced is set (to 1) by the hardware. After some time, we can determine which pages have been used and which have not been used by examining the reference bits. We do not know the *order* of use, but we know which pages were used and which were not used. This partial ordering information leads to many page-replacement algorithms that approximate LRU replacement.

9.5.4.1 Additional-Reference-Bits Algorithm

We can gain additional ordering information by recording the reference bits at regular intervals. We can keep an 8-bit byte for each page in a table in memory. At regular intervals (say every 100 milliseconds), a timer interrupt transfers control to the operating system. The operating system shifts the reference bit for each page into the high-order bit of its 8-bit byte, shifting the other bits right 1 bit, discarding the low-order bit. These 8-bit shift registers contain the history of page use for the last eight time periods. If the shift register contains 00000000, then the page has not been used for eight time periods; a page that is used at least once each period would have a shift register value of 11111111.

A page with a history register value of 11000100 has been used more recently than has one with 01110111. If we interpret these 8-bit bytes as unsigned integers, the page with the lowest number is the LRU page, and it

can be replaced. Notice that the numbers are not guaranteed to be unique, however. We can either replace (swap out) all pages with the smallest value, or use a FIFO selection among them.

The number of bits of history can be varied, of course, and would be selected (depending on the hardware available) to make the updating as fast as possible. In the extreme case, the number can be reduced to zero, leaving only the reference bit itself. This algorithm is called the *second-chance* page-replacement algorithm.

9.5.4.2 Second-Chance Algorithm

The basic algorithm of second-chance replacement is a FIFO replacement algorithm. When a page has been selected, however, we inspect its reference bit. If the value is 0, we proceed to replace this page. If the reference bit is 1, however, we give that page a second chance and move on to select the next FIFO page. When a page gets a second chance, its reference bit is cleared and its arrival time is reset to the current time. Thus, a page that is given a second chance will not be replaced until all other pages are replaced (or given second chances). In addition, if a page is used often enough to keep its reference bit set, it will never be replaced.

One way to implement the second-chance (sometimes referred to as the clock) algorithm is as a circular queue. A pointer indicates which page is to be replaced next. When a frame is needed, the pointer advances until it finds a page with a 0 reference bit. As it advances, it clears the reference bits (Figure 9.13). Once a victim page is found, the page is replaced and the new page is inserted in the circular queue in that position. Notice that, in the worst case, when all bits are set, the pointer cycles through the whole queue, giving each page a second chance. It clears all the reference bits before selecting the next page for replacement. Second-chance replacement degenerates to FIFO replacement if all bits are set.

9.5.4.3 Enhanced Second-Chance Algorithm

The second-chance algorithm described above can be enhanced by considering both the reference bit and the modify bit (Section 9.4) as an ordered pair. With these 2 bits, we have the following four possible classes:

1. (0,0) neither recently used nor modified - best page to replace
2. (0,1) not recently used but modified - not quite as good, because the page will need to be written out before replacement
3. (1,0) recently used but clean - probably will be used again soon
4. (1,1) recently used and modified - probably will be used again, and write out will be needed before replacing it

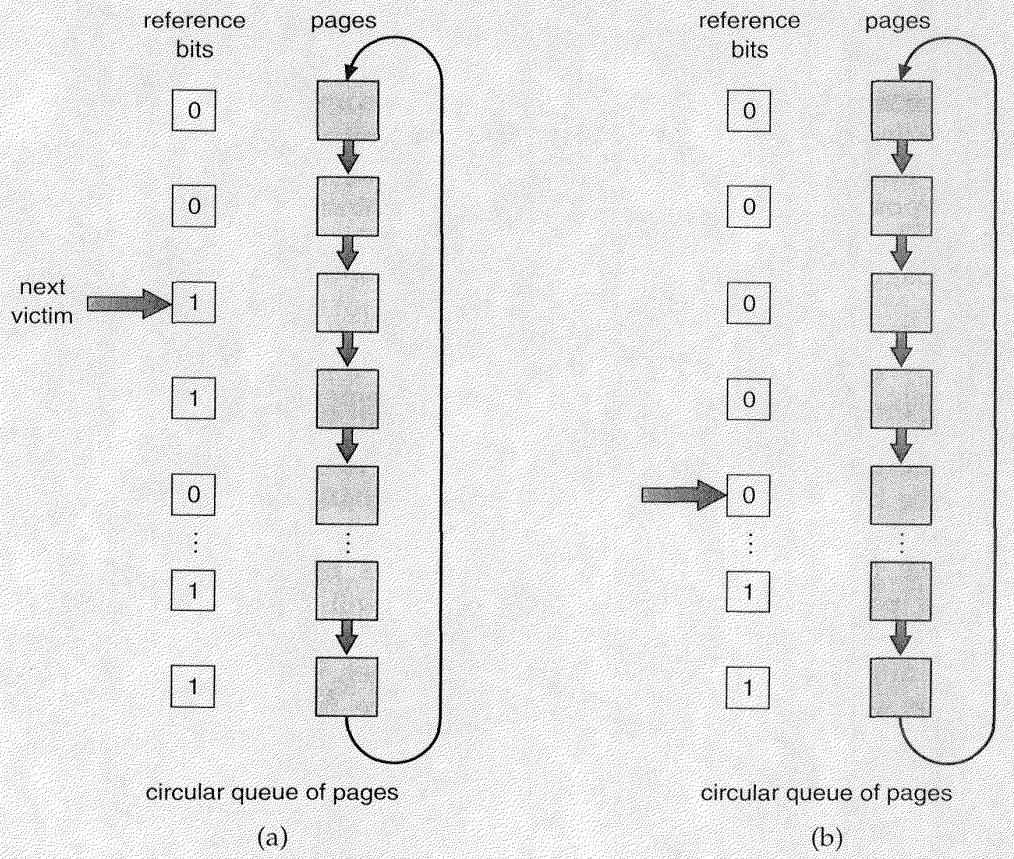


Figure 9.13 Second-chance (clock) page-replacement algorithm.

When page replacement is called for, each page is in one of these four classes. We use the same scheme as the clock algorithm, but instead of examining whether the page to which we are pointing has the reference bit set to 1, we examine the class to which that page belongs. We replace the first page encountered in the lowest nonempty class. Notice that we may have to scan the circular queue several times before we find a page to be replaced.

This algorithm is used in the Macintosh virtual-memory-management scheme. The major difference between this algorithm and the simpler clock algorithm is that here we give preference to those pages that have been modified to reduce the number of I/Os required.

9.5.5 Counting Algorithms

There are many other algorithms that can be used for page replacement. For example, we could keep a counter of the number of references that have been made to each page, and develop the following two schemes.

HULU EXHIBIT 1018
Page 21 of 22

- **LFU Algorithm:** The *least frequently used (LFU)* page-replacement algorithm requires that the page with the smallest count be replaced. The reason for this selection is that an actively used page should have a large reference count. This algorithm suffers from the situation in which a page is used heavily during the initial phase of a process, but then is never used again. Since it was used heavily, it has a large count and remains in memory even though it is no longer needed. One solution is to shift the counts right by 1 bit at regular intervals, forming an exponentially decaying average usage count.
- **MFU Algorithm:** The *most frequently used (MFU)* page-replacement algorithm is based on the argument that the page with the smallest count was probably just brought in and has yet to be used.

As you might expect, neither MFU nor LFU replacement is common. The implementation of these algorithms is fairly expensive, and they do not approximate OPT replacement very well.

9.5.6 Page Buffering Algorithm

Other procedures are often used in addition to a specific page-replacement algorithm. For example, systems commonly keep a *pool* of free frames. When a page fault occurs, a victim frame is chosen as before. However, the desired page is read into a free frame from the pool before the victim is written out. This procedure allows the process to restart as soon as possible, without waiting for the victim page to be written out. When the victim is later written out, its frame is added to the free-frame pool.

An expansion of this idea is to maintain a list of modified pages. Whenever the paging device is idle, a modified page is selected and is written to the disk. Its modify bit is then reset. This scheme increases the probability that a page will be clean when it is selected for replacement, and will not need to be written out.

Another modification is to keep a pool of free frames, but to remember which page was in each frame. Since the frame contents are not modified when a frame is written to the disk, the old page can be reused directly from the free-frame pool if it is needed before that frame is reused. No I/O is needed in this case. When a page fault occurs, we first check whether the desired page is in the free-frame pool. If it is not, we must select a free frame and read into it.

This technique is used in the VAX/VMS system, with a FIFO replacement algorithm. When the FIFO replacement algorithm mistakenly replaces a page that is still in active use, that page is quickly retrieved from the free-frame buffer, and no I/O is necessary. The free-frame buffer provides protection against the relatively poor, but simple, FIFO replacement algorithm. This method is necessary because the early versions of the VAX did not correctly implement the reference bit.