

Storage and File Structure

Solutions to Practice Exercises

11.1 This arrangement has the problem that P_i and B_{4i-3} are on the same disk. So if that disk fails, reconstruction of B_{4i-3} is not possible, since data and parity are both lost.

- 11.2 a. To ensure atomicity, a block write operation is carried out as follows:
- i. Write the information onto the first physical block.
 - ii. When the first write completes successfully, write the same information onto the second physical block.
 - iii. The output is declared completed only after the second write completes successfully.

During recovery, each pair of physical blocks is examined. If both are identical and there is no detectable partial-write, then no further actions are necessary. If one block has been partially rewritten, then we replace its contents with the contents of the other block. If there has been no partial-write, but they differ in content, then we replace the contents of the first block with the contents of the second, or vice versa. This recovery procedure ensures that a write to stable storage either succeeds completely (that is, updates both copies) or results in no change.

The requirement of comparing every corresponding pair of blocks during recovery is expensive to meet. We can reduce the cost greatly by keeping track of block writes that are in progress, using a small amount of non-volatile RAM. On recovery, only blocks for which writes were in progress need to be compared.

- b. The idea is similar here. For any block write, the information block is written first followed by the corresponding parity block. At the time of

recovery, each set consisting of the n^{th} block of each of the disks is considered. If none of the blocks in the set have been partially-written, and the parity block contents are consistent with the contents of the information blocks, then no further action need be taken. If any block has been partially-written, it's contents are reconstructed using the other blocks. If no block has been partially-written, but the parity block contents do not agree with the information block contents, the parity block's contents are reconstructed.

- 11.3**
- a. MRU is preferable to LRU where $R_1 \bowtie R_2$ is computed by using a nested-loop processing strategy where each tuple in R_2 must be compared to each block in R_1 . After the first tuple of R_2 is processed, the next needed block is the first one in R_1 . However, since it is the least recently used, the LRU buffer management strategy would replace that block if a new block was needed by the system.
 - b. LRU is preferable to MRU where $R_1 \bowtie R_2$ is computed by sorting the relations by join values and then comparing the values by proceeding through the relations. Due to duplicate join values, it may be necessary to “back-up” in one of the relations. This “backing-up” could cross a block boundary into the most recently used block, which would have been replaced by a system using MRU buffer management, if a new block was needed.

Under MRU, some unused blocks may remain in memory forever. In practice, MRU can be used only in special situations like that of the nested-loop strategy discussed in example 0.a

- 11.4**
- a. Although moving record 6 to the space for 5, and moving record 7 to the space for 6, is the most straightforward approach, it requires moving the most records, and involves the most accesses.
 - b. Moving record 7 to the space for 5 moves fewer records, but destroys any ordering in the file.
 - c. Marking the space for 5 as deleted preserves ordering and moves no records, but requires additional overhead to keep track of all of the free space in the file. This method may lead to too many “holes” in the file, which if not compacted from time to time, will affect performance because of reduced availability of contiguous free records.

11.5 (We use “ $\uparrow i$ ” to denote a pointer to record “ i ”.)
The original file of Figure 11.8.

header				$\uparrow 1$
record 0	A-102	Perryridge	400	
record 1				$\uparrow 4$
record 2	A-215	Mianus	700	
record 3	A-101	Downtown	500	
record 4		Perryridge		$\uparrow 6$
record 5	A-201		900	
record 6				
record 7	A-110	Downtown	600	
record 8	A-218	Perryridge	700	

a. The file after **insert** (Brighton, A-323, 1600).

header				$\uparrow 4$
record 0	A-102	Perryridge	400	
record 1	A-323	Brighton	1600	
record 2	A-215	Mianus	700	
record 3	A-101	Downtown	500	
record 4				$\uparrow 6$
record 5	A-201	Perryridge	900	
record 6				
record 7	A-110	Downtown	600	
record 8	A-218	Perryridge	700	

b. The file after **delete** record 2.

header				$\uparrow 2$
record 0	A-102	Perryridge	400	
record 1	A-323	Brighton	1600	
record 2				$\uparrow 4$
record 3	A-101	Downtown	500	
record 4				$\uparrow 6$
record 5	A-201	Perryridge	900	
record 6				
record 7	A-110	Downtown	600	
record 8	A-218	Perryridge	700	

The free record chain could have alternatively been from the header to 4, from 4 to 2, and finally from 2 to 6.

c. The file after **insert** (Brighton, A-626, 2000).

header				↑ 4
record 0	A-102	Perryridge	400	
record 1	A-323	Brighton	1600	
record 2	A-626	Brighton	2000	
record 3	A-101	Downtown	500	
record 4				↑ 6
record 5	A-201	Perryridge	900	
record 6				
record 7	A-110	Downtown	600	
record 8	A-218	Perryridge	700	

11.6 Instance of relations:

course relation

course_name	room	instructor
Pascal	CS-101	Calvin, B
C	CS-102	Calvin, B
Lisp	CS-102	Kess, J

c_1
 c_2
 c_3

course_name	student_name	grade
Pascal	Carper, D	A
Pascal	Merrick, L	A
Pascal	Mitchell, N	B
Pascal	Bliss, A	C
Pascal	Hames, G	C
C	Nile, M	A
C	Mitchell, N	B
C	Carper, D	A
C	Hurly, I	B
C	Hames, G	A
Lisp	Bliss, A	C
Lisp	Hurly, I	B
Lisp	Nile, M	D
Lisp	Stars, R	A
Lisp	Carper, D	A

Block 0 contains: $c_1, e_1, e_2, e_3, e_4,$ and e_5

Block 1 contains: c_2, e_6, e_7, e_8, e_9 and e_{10}

Block 2 contains: $c_3, e_{11}, e_{12}, e_{13}, e_{14},$ and e_{15}

- 11.7**
- a.** Everytime a record is inserted/deleted, check if the usage of the block has changed levels. In that case, update the corresponding bits. Note that we don't need to access the bitmaps at all unless the usage crosses a boundary, so in most of the cases there is no overhead.
 - b.** When free space for a large record or a set of records is sought, then multiple free list entries may have to be scanned before finding a proper sized one, so overheads are much higher. With bitmaps, one page of bitmap can store free info for many pages, so I/O spent for finding free space is minimal. Similarly, when a whole block or a large part of it is deleted, bitmap technique is more convenient for updating free space information.