

Errata and Updates For Database System Concepts, 6th Edition Silberschatz, Korth, and Sudarshan

Last updated: October 30, 2016

We list below errors, clarifications, and recent updates. NOTE: Some of the errata described here have been corrected in the 2nd and 3rd printing of the US edition of the book; these are tagged as “(1st pr.)”. If you own an international edition, note that these editions follow a different correction schedule, so your copy may still have errata that have been fixed in the US edition. Check your copy for the errata noted here, and ignore those that have been fixed in your copy.

Errata for Part 3: Data Storage and Querying, Chapters 10 to 13

CHAPTER 10

1. (1st pr.) Page 445, RAID level 3: The sentences “Since reads and writes of a byte are spread out over multiple disks, ... since every disk has to participate in every I/O request.” are technically correct, but somewhat confusing. Replace these sentences by:

“Since reads and writes of a byte are spread out over multiple disks, with N -way striping of data, the transfer rate is N times faster than a RAID level 1 organization using N -way striping, for reading or writing a single block. However, RAID level 3 has no transfer rate benefit for writes of multiple blocks. Further, RAID level 3 has a higher access time, and supports a lower number of I/O operations per second, since every disk has to participate in every I/O request.”

CHAPTER 11

1. (1st pr.) Page 486, first para of Section 11.3.1: After $K_i < K_j$ add “(we assume for now that there are no duplicate key values).”
2. (1st pr.) Page 487, 5th para: Delete the sentence: “We have shown instructor names abbreviated to 3 characters in order to depict the tree clearly; in reality, the tree nodes would contain the full names.”
3. (1st pr.) Page 487, last para: Delete the sentence: “As before, we have abbreviated instructor names only for clarity of presentation.”
4. (1st pr.) Page 488, Figure 11.9: The line from the left-most leaf to the record for Crick should come from the space between Califieri and Crick, not from the space after Crick.
(Reported by: Minhua Kang)
5. (1st pr.) Page 488, first para of Section 11.3.2: “Suppose that we wish to find records with a search-key value of V . Figure 11.11 presents pseudocode for a function `find()` to carry out this task.” →
“Suppose that we wish to find a record with a search-key value of V . Figure 11.11 presents pseudocode for a function `find()` to carry out this task, assuming there are no duplicates.”

³**Errors reported by:** Deepak Aggrawal, G. Aishwarya, Jameel Al-Aziz, Scot Anderson, Yahui Chang, David Chiu, Jonghoon Chun, Matt Cremeens, Dona Dungan, Pham Nguyen Duc Duong, Helena Galhardas, Eric Gossett, Ravindra Guravannavar, Leon Ho, Pranav Jain, Jevitha K. P., Cheqing Jin, Minhua Kang, Celine Kuttler, Daniel Sadoc Menasche, Thimas Nielsen, Linda Null, Judi Paige, Donnie Pinkston, Subhasish Saha, Vemireddy Satish, Shan Shimin, Stan Thomas, Cam Hong Tran, Duc Tran, Daniel Vieira, and a few others. Their help, and in particular that of Daniel Sadoc Menasche, is deeply appreciated. Also thanks to Juha Haaga for suggestions on improvements for future editions.

6. (1st pr.) Page 489, Figure 11.11, function `find()`: Change
 “/* Returns leaf node C and index i such that $C.P_i$ points to first record with search key value V */
 \rightarrow
 “/* Assumes no duplicate keys, and returns leaf node C and index i such that $C.P_i$ points to record with search key value V , if such a record exists */”
7. (1st pr.) Page 489, Figure 11.11, function `printAll()`: “Set $(L, i) = \text{find}(V)$;
 \rightarrow “Set $(L, i) = \text{findFirst}(V)$;
8. (1st pr.) Page 490, Para 2: Replace this entire paragraph starting with “If there is at most one record ..” with the following new paragraph:
 “The `find` function of Figure 11.11 needs to be modified to handle duplicate keys. With duplicate keys, for both leaf and internal nodes, if $i < j$, then $K_i < K_j$ may not hold, but certainly $K_i \leq K_j$ holds. Further, records in the subtree pointed to by P_i may contain values that are less than or equal to K_i (to understand why, consider two adjacent leaf nodes pointed to by P_i and P_{i+1} that both contain a duplicate key value v , in which case $K_i = v$). To fix this, we must modify the loop in the `find` function to set $C = C.P_i$, even if $V = C.K_i$. Further, the leaf node C reached thereby may contain only search keys less than V (even if V does exist in the tree); in this case the `find` procedure must set $C = \text{right sibling } C$, and recheck if C contains V . The modified `find` procedure, which we call `findFirst`, returns the first occurrence of value V in the tree.”
9. (1st pr.) Page 490, Para 3, Line 2:
 “The procedure first steps through the remaining keys in the node L , to find other records with search-key value V .”
 \rightarrow
 “The `printAll` procedure calls `findFirst` to find the node L with the first occurrence of V , and then steps through the remaining keys in the node L , to find other records with search-key value V .”
10. Page 490, Para 4: Change:
 “To execute such queries, we can create a procedure `printRange(L, U)`, whose body is the same as `printAll` except for these differences: `printRange` calls `find(L)`, instead of `find(V)`, and then steps through records as in procedure `printAll`, but with the stopping condition being that $L.K_i > U$, instead of $L.K_i > V$.”
 \rightarrow
 “To execute such queries, we can create a procedure `printRange(lb, ub)`, which does the following: it first traverses to a leaf in a manner similar to `find(lb)`; the leaf may or may not actually contain value lb . It then steps through records in a manner similar to `printall`, but only returns records with key values $L.K_i$ s.t. $lb \leq L.K_i \leq ub$, and stops when $L.K_i > ub$.”
 (Reported by: Deepak Aggrawal)
11. Page 494, Figure 11.15, procedure `insert_in_leaf`: Change
 “Let K_i be the highest value in L that is less than K ”
 \rightarrow
 “Let K_i be the highest value in L that is less than or equal to K ”
 and also change
 “Insert P, K into L just after $T.K_i$ ”
 \rightarrow
 “Insert P, K into L just after $L.K_i$ ”
 (Reported by: Donnie Pinkston)
12. Page 494, Figure 11.15, procedure `insert_in_parent`: Change:

Copy $T.P_1 \dots T.P_{\lfloor n/2 \rfloor}$ into P

Let $K'' = T.K_{\lceil n/2 \rceil}$
 Copy $T.P_{\lceil n/2 \rceil + 1} \dots T.P_{n+1}$ into P'

→

Copy $T.P_1 \dots T.P_{\lceil (n+1)/2 \rceil}$ into P
 Let $K'' = T.K_{\lceil (n+1)/2 \rceil}$
 Copy $T.P_{\lceil (n+1)/2 \rceil + 1} \dots T.P_{n+1}$ into P'

NOTE: The change does not affect correctness, both versions are correct (i.e. it does not matter whether we take $\lceil n/2 \rceil$ or $\lceil (n+1)/2 \rceil$), but this change was done to ensure the procedure is consistent with the examples in the book.

(Reported by: Helena Galhardas)

13. (1st pr.) Page 496, line 3: "... the leaf node containg "Mozart"..." → "... the leaf node containg "Gold" ...".
 (Reported by: Jayvant Anantpur)
14. Page 504, last 2 paragraphs:
 "11.13" → "11.09" in two places, and
 "... "Califieri", "Einstein", "Gold", ... " → "... , "Einstein", "Gold", ... "
 (Reported by: Deepak Aggrawal)
15. Page 505, Caption of Figure 11.21: "11.13" → "11.09"
 (Reported by: Deepak Aggrawal)
16. (1st pr.) Page 505, Figure 11.21: "... and soon for other records ..." →
 "... and so on for other records ...".
17. (1st pr.) Page 513, Para 2:
 "The form of hash structure that we have just described is sometimes referred to as closed hashing. Under an alternative approach called open hashing ..."

→

"The form of hash structure that we have just described is called closed addressing (or, less commonly, closed hashing). Under an alternative approach called open addressing (or, less commonly, open hashing) ..."

Also in the same para: change all occurrences of "open hashing" → "open addressing" and "closed hashing" → "closed addressing".

Why this change?: The form of hashing which we refer to as **closed hashing** is referred to more commonly as **closed addressing**, while the form of hashing we refer to as **open hashing** is referred to more commonly as **open addressing**.

But many sources also use the term "closed hashing" synonymously with "open addressing", and "open hashing" synonymously with "closed addressing", which is the exact opposite of our notation. To avoid this confusion, we suggest using the term closed addressing and open addressing, instead of closed hashing and open hashing.

(Reported by: Donnie Pinkston)

18. (1st pr.) Page 519, Figure 11.29: Change salary value of Srinivasan from 90000 → 65000.
19. Page 521, Figure 11.33: The value in the box above the top bucket (just above 15151) should be 1, not 2.
 (Reported by: Vemireddy Satish)
20. (1st pr.) Page 533, Question 11.11: "Outline the steps in ..." → "Assuming the availability of the above bitmap index on *salary*, and a bitmap index on *dept_name*, outline the steps in ...".

21. Page 534, Exercise 11.14 (b) iii: “... a n -page block” \rightarrow “... an n -block unit”.
22. (1st pr.) Page 534, exercise 11.19: “closed and open hashing” \rightarrow “closed and open addressing”.

CHAPTER 12

1. (1st pr.) Page 543, Figure 12.3: In row A3, “ $h_i * (t_T + t_S) + b * t_T$ ” \rightarrow “ $h_i * (t_T + t_S) + t_S + b * t_T$ ”
2. (1st pr.) Page 543, Figure 12.3: In row A5, “ $h_i * (t_T + t_S) + b * t_T$ ” \rightarrow “ $h_i * (t_T + t_S) + t_S + b * t_T$ ”
(Reported by: Daniel Sadoc Menasche, G. Aishwarya and Subhasish Saha)
3. Page 544, bullet for A5: “first tuple in the file that has a value of $A = v$ ” \rightarrow “first tuple in the file that has a value of $A \geq v$ ”.
4. (1st pr.) Page 549, paragraphs 1 and 2: Some of the formulae here assume runs are read in one block at a time, and others assume that runs are read in b_b blocks at a time, to reduce the number of seeks. To restore consistency, make the following changes:

- (a) (1st pr.) Replace the contents of Paragraph 1 starting from: “Since the number of runs decreases bny a factor of $M - 1$ in each merge pass ...” till the end of the paragraph by:

During the merge pass, reading in each run one block at a time leads to a large number of seeks; to reduce the number of seeks, a larger number of blocks, denoted b_b , are read or written at a time, requiring b_b buffer blocks to be allocated to each input run and to the output run. Then, $\lfloor M/b_b \rfloor - 1$ runs can be merged in each merge pass, decreasing the number of runs by a factor of $\lfloor M/b_b \rfloor - 1$. The total number of merge passes required is $\lceil \log_{\lfloor M/b_b \rfloor - 1}(b_r/M) \rceil$. Each of these passes reads every block of the relation once and writes it out once, with two exceptions. First, the final pass can produce the sorted output without writing its result to disk. Second, there may be runs that are not read in or written out during a pass—for example, if there are $\lfloor M/b_b \rfloor$ runs to be merged in a pass, $\lfloor M/b_b \rfloor - 1$ are read in and merged, and one run is not accessed during the pass. Ignoring the (relatively small) savings due to the latter effect, the total number of block transfers for external sorting of the relation is:

$$b_r(2\lceil \log_{\lfloor M/b_b \rfloor - 1}(b_r/M) \rceil + 1)$$

Applying this equation to the example in Figure 12.4, with b_b set to 1, we get a total of $12 * (4 + 1) = 60$ block transfers, as you can verify from the figure. Note that the above numbers do not include the cost of writing out the final result.

- (b) (1st pr.) Page 549: Replace the contents of Paragraph 2, starting from “During the merge phase ...” till the end of the paragraph by:

Each merge pass the requires around $\lceil b_r/b_b \rceil$ seeks for reading data.⁴ Although the output is written sequentially, if it is on the same disk as the input runs the head may have moved away between writes of consecutive blocks. Thus we would have to add a total of $2\lceil b_r/b_b \rceil$ seeks for each merge pass, except the final pass (since we assume the final result is not written back to disk).

$$2\lceil b_r/M \rceil + \lceil b_r/b_b \rceil(2\lceil \log_{\lfloor M/b_b \rfloor - 1}(b_r/M) \rceil - 1)$$

Applying this equation to the example in Figure 12.4, we get a total of $8 + 12 * (2 * 2 - 1) = 44$ disk seeks if we set the number of buffer blocks per run, b_b to 1.

(Reported by: Leon Ho)

5. Page 556, bullet 2: “1164” \rightarrow “1168” and “1264” \rightarrow “1268”.
(Reported by: Deepak Aggrawal)
6. (1st pr.) Page 561, Section 12.5.5.4, paragraph following the 2nd bullet: For the case with recursive partitioning, some formulae assume data is read in b_b blocks at a time, while others do not. For consistency, the following changes are required:

- (a) Replace the text in this paragraph starting from “Each pass reduces the size of each of the ...” by:

Again we assume that b_b blocks are allocated for buffering each partition. Each pass then reduces the size of each of the partitions by an expected factor of $\lfloor M/b_b \rfloor - 1$; and passes are repeated until each partition is of size at most M blocks. The expected number of passes required for partitioning s is therefore $\lceil \log_{\lfloor M/b_b \rfloor - 1}(b_s/M) \rceil$.

- (b) Replace the bullet immediately after the paragraph, starting with “Since, in each pass, every block ...” by:

Since, in each pass, every block of s is read in and written out, the total block transfers for partitioning of s is $2b_s \lceil \log_{\lfloor M/b_b \rfloor - 1}(b_s/M) \rceil$. The number of passes for partitioning of r is the same as the number of passes for partitioning of s , therefore the join is estimated to require:

$$2(b_r + b_s) \lceil \log_{\lfloor M/b_b \rfloor - 1}(b_s/M) \rceil + b_r + b_s$$

block transfers.

- (c) In the following bullet:

“ $2(\lceil b_r/b_b \rceil + \lceil b_s/b_b \rceil) \lceil \log_{M-1}(b_s) - 1 \rceil$ ”

→

“ $2(\lceil b_r/b_b \rceil + \lceil b_s/b_b \rceil) \lceil \log_{\lfloor M/b_b \rfloor - 1}(b_s/M) \rceil$ ”

7. (1st pr.) Page 562, first para: “There is enough memory to allocate 3 buffers for the input and each of the 5 outputs during partitioning ...”

→

“There is enough memory to allocate 3 buffers for the input and each of the 5 outputs during partitioning (that is, $b_b = 3$) ...”

8. (1st pr.) Page 566, Section 12.6.5, 2nd para: “*branch_name*” → “*dept_name*”.
(Reported by: Daniel Sadoc Menasche)

9. (1st pr.) Page 572, para 1: “Thus, the join of r_1 with s_0 , and s_0 with r_1 , ...” → “Thus, the join of r_1 with s_0 , and s_1 with r_0 , ...”.

(Reported by: G. Aishwarya and Subhasish Saha)

CHAPTER 13

1. Page 580, second line: “ten attributes” → “nine attributes”
(Reported by: Deepak Aggrawal)

2. (1st pr.) Page 584, Rule 3: At the end of the rule, add the line: “ where $L_1 \subseteq L_2 \subseteq \dots \subseteq L_n$.”
(Reported by: Daniel Sadoc Menasche)

3. (1st pr.) Page 585, Rule 12: At the end of the rule: “ The projection operation distributes over the union operation

$$\Pi_L(E_1 \cup E_2) = (\Pi_L(E_1)) \cup (\Pi_L(E_2))$$

”

add the condition: “provided E_1 and E_2 have the same schema.”

(Reported by: G. Aishwarya and Subhasish Saha)

4. Page 588, query just before Section 13.2.3: Add an extra “)” just after the relation *teaches*.

5. Page 589 1st para of 13.2.4:

“If an expression, say E_i , of any subexpression e_i of E_i (which could, as a special case, be E_i itself) matches one side of an equivalence rule, the optimizer generates a new expression where e_i is transformed to match the other side of the rule.”

→

“If a subexpression e_j of any expression $E_i \in EQ$ (as a special case, e_j could be E_i itself) matches one side of an equivalence rule, the optimizer generates a copy E_k of E_i , in which e_j is transformed to match the other side of the rule, and adds E_k to EQ .”

(Reported by: Deepak Aggrawal)

6. (1st pr.) Page 595, Line 4: “ $n(r)$ ” \rightarrow “ n_r ”

7. (1st pr.) Page 601-602: ”We have already seen equivalence rules with aggregation operation, and equivalence rules can also be created for outer joins.”

\rightarrow

“Equivalence rules can also be defined for the aggregation and outer join operations, as illustrated in Practice Exercises 13.1 and 13.2.”

(Reported by: Ravindra Guravannavar)