

As with virtually all books, some bugs were discovered after printing. These errors have been corrected in subsequent printings (impressions). The corrections recorded to date are listed below. To apply them to your copy, you need to work out which impression you have and then apply the corrections appropriate to that impression.

We are currently (December 1997) on the Fifth Impression. To discover the impression number of your copy, look on the copyright page – *page iv* (reverse of the title page). If you have the:

Fifth impression you will find this on the copyright page:

- 1st and 2nd impressions (1994)
- 3rd impression (1995)
- 4th impression (1995)
- 5th impression (1998)

We have not yet recorded any corrections to the fifth impression.

Fourth impression you will find this on the copyright page:

- 1st and 2nd impressions (1994)
- 3rd impression (1995)
- 4th impression (1995)

We have recorded 26 corrections to the fourth impression. You should apply only the **Fourth impression errata** (see Section 4).

Third impression you will find this on the copyright page:

- 1st and 2nd impressions (1994)
- 3rd impression (1995)

We have recorded 56 corrections to the third impression. You should apply the **Third impression errata** (see Section 3) and the **Fourth impression errata** (see Section 4).

Second impression you will find this on the copyright page:

- Printed 1994 (twice)

We have recorded 90 corrections to the second impression. You should apply the **Second impression errata** (see Section 2 below), the **Third impression errata** (see Section 3 below) and the **Fourth impression errata** (see Section 4).

A small number of copies of the Second Impression also have a severely damaged index. The damage resulted in the printing of page 638 as a replica of page 636 and page 642 as a replica of page 640. If your copy of the book doesn't have a page 638 that is a copy of page 636, then you can ignore these corrections. If your book does have the index problem, then you should download and print the PostScript files in the directory:

- <ftp://ftp.dcs.qmw.ac.uk/distrib/dsbook/index.corrections/>

First impression you will find this on the copyright page:

- First printed 1994

We have recorded 132 corrections to the first impression. You should apply all of the corrections listed in Sections 1, 2, 3 and 4 below. In two places, the corrections are too complex to apply manually, so we have supplied PostScript files for printing replacement pages. These are:

- p147-8.ps
- pp505-8.ps

NB. Some cosmetic changes were also made in the later impressions, particularly to the punctuation, layout and figures. For simplicity, we have not included them in this list.

1. First impression errata

<i>Page</i>	<i>existing version</i>	<i>correction</i>	<i>date</i>
ix	line 9: ESS	SS	14-10-94
xv	FOREWARD	FOREWORD	30-6-94
21	para 5: As an illustration of the presence of network transparency,...	As an illustration of the absence of network transparency,...	15-7-94
79	line 1	change 46 to 64	27-9-94
106	line 7: ;ocation transparency	location transparency	30-6-94
121	Figure 4.10	add arrowhead to horizontal line (c.f. Figure 4.11)	30-6-94
124	Exercise 4.6	delete second sentence	15-7-94
147	Figure 5.10 (or use attached corrected pages 147-8)	change caption to: Outline of C program for server in ANSA testbench.	27-9-94
147	Figure 5.10 lines 15 and 19 (or use attached corrected pages 147-8)	programs for FileReadWrite_Read and FileReadWrite_Write	27-9-94
147	Figure 5.10 lines 12, 16	replace “ _” by “_” (i.e. remove leading space)	27-9-94
147	last para, line 2: (or use attached corrected pages 147-8)	Replace by: As well as specifying the name of the interface, the <i>Export</i> procedure can specify a context in the Trader database, the desired attributes of the service and the number of requests that may be queued at the service (five in our example).	27-9-94
147	line -3 to -1 (or use attached corrected pages 147-8)	Replace by The names of the server procedures consist of the interface name followed by an underline and the name of the corresponding procedure in the interface definition. For example, the service procedures corresponding to Write and Read in Figure 5.7 are FileReadWrite_Write and FileReadWrite_Read as in Figure 5.10.	27-9-94
148	last 2 sentences of para 1 (or use attached corrected pages 147-8)	For example, the FileReadWrite_Read procedure has one result whose type is declared as Data *.	27-9-94
199	Para 4, line 6, sentence beginning “Most current file services...”	Replace sentence with: Most current file services follow modern UNIX standards in providing advisory or mandatory file- or record-level locking.	14-10-94
225	Para 3, line 5, a feature of UNIX System V (but are not in UNIX BSD 4.3).	a feature of most recent versions of UNIX.	14-10-94
246	line -7 and line -5	add “)” at end of each line.	14-10-94
250	Ex. 8.2, last sentence	What constraints should an underlying filing system obey to be suitable for the implementation of NFS servers?	11-8-94 27-9-94
258	line 16 : <i>.gene.parallel</i>	<i>.gene..parallel</i>	27-9-94
259	line 28 : <i>.cs.</i>	<i>.cs</i>	27-9-94
262	Figure 9.4: <i>dibbLab</i>	<i>distrib</i>	27-9-94
265	line 8: <i>.cs.DistribLab.gene</i>	<i>.cs.distrib.gene</i>	27-9-94
	line 11: two or more	they	27-9-94
269	line 5: <i>dcs,</i>	<i>com,</i>	27-9-94

307	line 8: If the arrived identifier is smaller, then it substitutes its own identifier in the message before forwarding it; but it makes the substitution only if it is not already a <i>participant</i> .	If the arrived identifier is smaller and the receiver is not a <i>participant</i> then it substitutes its own identifier in the message and forwards it; but it does not forward the message if it is already a <i>participant</i> .	1-12-94
312	line 6: later in full	delete words	22-1-95
324	line -10: two rounds of messages are required	three messages are sent between the front end and each replica manager	22-1-95
332	line 10: <i>inval</i>	<i>executed</i>	22-1-95
334	line -12: piggy-packing	piggy-backing	22-1-95
350	Exercise 11.6, line 1: debits	credits	6-11-94
394	line -12: (to some other transaction)	(to some other committed transaction)	25-11-94
395	line -12	Add footnote: Forward validation should allow for the fact that read sets of active transactions may change during validation and writing.	4-12-94
395	line 2: StartTn = T2	StartTn + 1 = T2	25-11-94
397	line -5: Tj > the maximum write timestamp of the data item	Tj > the maximum write timestamp of the committed data item	25-11-94
397	line -7: Tj >	Tj >=	25-11-94
400	last sentence	However, if the other transaction had read the data item, the one with the late write will fail due to the read timestamp on the data item.	29-11-94
413	Fig 14.3	Message numbers: exchange: 3 with 5 and 4 with 6	29-11-94
419	line -4: but the fate of T21	but the fate of T12	29-11-94
430	line 12	change <T-> -> U> to <T-> U>	6-11-94
436	last line	change T writes to U writes	6-12-94
440	before line 7, insert:	Locks are applied with the same granularity as version numbers.	19-12-94
460	line 8: 2N	3N	1- 2-95
465	Fig 15.5 (b)	reverse direction of horizontal arrow	9-11-94
477	Section numbers not in sequence	numbers 16.5-16.7	1- 2-95
495	line 17: (... Section 16.1)	(... Section 16.5)	3-12-94
512	Table, message 2, Notes	Delete first sentence: The authentication server...plain text.	9-11-94
535	Figure 17.10 (b)	remove 3 horizontal arrows, move 'owner' to process A	1- 2-95
564	Figure 18.9, row 5: External pager directs kernel to write or flush page	External pager directs kernel to write page/set page's access permissions	15-11-94
628	Insert reference	[Mattern 1988] Mattern, F.(1989). Virtual time and global states of distributed systems. In <i>Parallel and Distributed Algorithms</i> , (M. Cosnard et al eds.), North-Holland, Amsterdam, pp. 215-226.	3-12-94
638	If your copy has the same index entries on pp 636 and 638, see the README.FIRST file in		1.2.95
-642	ftp directory 'index.corrections', or contact Addison-Wesley for a reprinted index.		

3. Third impression errata

<i>Page</i>	<i>existing version</i>	<i>correction</i>	<i>date</i>
xiv	Support for instructors and readers section revised to show Web page as below: URL for the book's home page: http://www.dcs.qmw.ac.uk/research/distrib/book.html		28.9.95
133	line 10 after table: This known as	This is known as	24-3-95
142	line 7 to 8	delete sentence "The number of retries..."	24-3-95
145	last line: anInterface	aFileServerRef	24-3-95
258	line 22: <i>directories</i>	directories	5/10/95
267	line -8: it	the server	5/10/95
268	line 14: previous section	above	5/10/95
269	line -13: different countries	countries	5/10/95
318	line 6: perform	request	5/10/95
322	line -6: request managers	replica managers	5/10/95
359	line 7: and	which	24-3-95
359	line -3: Figure 12.2:	Figure 12.1.	24-3-95
366	Figure 12.7 caption: An inconsistent retrievals solution.	A solution to the problem of inconsistent retrievals.	24-3-95
367	Replace last para with: Premature writes [] Consider another implication of the possibility that a transaction may abort. This one is related to the interaction between <i>Write</i> operations on the same data item belonging to different transactions. For an illustration, we introduce to the Bank service, a new operation <i>SetBalance(Name, Amount)</i> which sets the balance of account <i>Name</i> to <i>Amount</i> . We consider two <i>SetBalance</i> transactions T and U on account A, as shown in Figure 12.9. Before the transactions, the balance of account A was \$100.		24-3-95
368	Replace first 3 paras with: The two executions are serially equivalent, with T setting the balance to \$3 and U setting it to \$5. If the transaction U aborts and T commits, the balance should be \$3. Some database systems implement the action of Abort by restoring 'before images' of all the <i>Writes</i> of a transaction. In our example, A is \$100 initially, which is the 'before image' of T's <i>Write</i> , similarly \$3 is the 'before image' of U's <i>Write</i> . Thus if U aborts, we get the correct balance of \$3. Now consider the case when U commits and then T aborts. The balance should be \$5 but as the 'before image' of T's <i>Write</i> is \$100, we get the wrong balance of \$100. Similarly if T aborts and then U aborts, the 'before image' of U's <i>Write</i> is \$3 and we get the wrong balance of \$3 – the balance should of course revert to \$100. To ensure correct results, Write operations must be delayed until earlier transactions that updated the same data items have either committed or aborted.		24-3-95
368	Figure 12.9. Replace by: Figure 12.9 Over-writing uncommitted values.		24-3-95

Transaction T:		Transaction U:	
<i>Bank\$SetBalance(A, 3)</i>		<i>Bank\$SetBalance(A, 5)</i>	
<i>A.Write(3)</i>	\$3	<i>A.Write(5)</i>	\$5

381	lines 25-28. Replace by:	To enforce (1) a request for a write lock on a data item is delayed by the presence of a read lock belonging to another transaction. To enforce (2) a request for either a read lock or a write lock on a data item is delayed by the presence of a write lock belonging to another transaction.	24-3-95
390	Figure 13.12. Under column Commit in rows Write and Commit: "Wait"	"_"	24-3-95
399	line 13: ...commits or aborts	commits or aborts then re-apply the Read rule	24-3-95
399	last line: (see Section 12.3)	(see Section 12.4)	24-3-95
402	line 6 and line 10: less than	less than or equal to	24-3-95
403	line 5: read timestamp T5	read timestamp T5 of the version with the write timestamp T3	24-3-95
415	Figure 14.4: HaveCommitted(Trans)	HaveCommitted(Trans, Worker)	24-3-95
420	line -10: committed	provisionally committed	24-3-95
433	line -4: Figure 14.18	Figure 14.18 (there can be different numbers of replicas of the various data items).	24-3-95
440	First line of table, last col.: Example 2	Example 3	24-3-95
440	Last line of table penultimate column: 0.03	0.0101	24-3-95
442	line 7: which one vote	which has one vote	24-3-95
465	lines 1 and 2	exchange {yes, yes, yes} and {yes, yes, no}	24-3-95
496	Figure 16.8	invert the request and reply arrows	30-3-95
497	line 6: on the right.	at the bottom of the page.	30-3-95
629	Before line 5 [Myer 1988]: insert missing reference (which is cited on page 212)		28-9-95
	[Mullender 1985]	Mullender, S.J. (1985). <i>Principles of Distributed Operating System Design</i> , Mathematisch Centrum, Amsterdam (Doctoral Thesis)	

4. Fourth impression errata

<i>Page</i>	<i>existing version</i>	<i>correction</i>	<i>date</i>
93	line 10: proceeding recursively from left to right	proceeding recursively from right to left	1/11/97
111	line 10: Exercise 4.18	Exercise 4.19	24/1/97
171	line 10: $(0.75 \times 8 + 0.25 \times 0) = 2$	$(0.75 \times 0 + 0.25 \times 8) = 2$	16-11-95
171	line -11 to -9: replace sentence: These two threads...	The reader should check that two threads can process 444 requests per second and three or more threads, bounded by the input-output time, can process 500 requests per second.	27-2-96
233	Lines 17-18: We shall consider ... Andrew File System.	Delete these lines.	13-11-97
240	Figure 8.8: Definition of <i>Store()</i>	Updates the attributes and (optionally) the contents of the specified file.	16-11-95
246	line -14: the first, beginning $s \neq \emptyset$ refers to disconnected operation and the second deals with all situations in which the AVSG is not empty.	the first, beginning $s \neq \emptyset$ refers to all situations in which the AVSG is not empty and the second deals with disconnected operation.	16-11-95
247	line -9: probe message to the all servers	probe message to all the servers	13/11/97
247	line -11: shrinking of an AVSG (due to an inaccessible server becoming accessible)	shrinking of an AVSG (due to a server becoming inaccessible)	13/11/97
317	line 7: If the clients and replica managers...	If the clients and front ends...	7-2-96
327	line 6: it can be shown that $u \leq v$	it can be shown that $u < v$	24-11-95
411	Figure 14.1a	show T to left of client	20-8-96
414	line 17: client	coordinator	20-8-96
421	line -1: ...the fate of the entire transaction.	...the fate of its parent transaction.	27-2-96
422	line 15 : transactions during the atomic commit protocol.	transactions during the atomic commit protocol, although an aborted transaction releases its locks after phase 1 of the protocol.	20-8-96
442	line -7: it keeps on trying.	it keeps on trying until it can create a new virtual partition.	13-11-97
459	line 14: example, T and then T, because parent	However, parent and child transactions cannot run in parallel and must therefore access the data item in turn, for example, T and then T1 and then T11 and then T1 and then T.	28/11/96
462	line -11 to -6: But it sometimes executes a remote procedure more than once. If the service operations are not idempotent, response failures can occur. Thus, at-least-once masks omission failures but may convert them into response failures. At-most-once provides a service without response failures – an example of hiding the omission failures altogether. The latter has more expensive storage requirements.	As it sometimes executes a remote procedure more than once, the service operations should be idempotent. At-most-once provides a service that masks both omission failures (by retransmitting request messages) and response failures due to non-idempotent operations (by filtering duplicates and re-transmission of replies).	13-11-97
467	line 25: causal ordering	causal as well as total ordering	13-11-97
479	line 9: Section 16.1	Section 16.5	27-2-96

Errata list	Distributed Systems: Concepts and Design - Edition 2	December 5, 1997	E8
489	line -13: so blocks will be of size $2^7 = 128$	since $2^7 = 128$	13-11/97
490	line -13: Section 16.2	Section 16.6	27-2-96
491	line -3: Section 16.2	Section 16.6	27-2-96
495	line 7: to message 3, so that it becomes:	to the ticket in message 2, so that it becomes:	7-2-96
495	line 7: B decrypts this message	On receipt of message 3, B decrypts the ticket	7-2-96
583	Figure 18.17, key: Task	Actor	16-11-95
