

---

# INDEX

---

- Abort, 3–5, 25–27, 89, 168, 179, 270, 321
- ABORT message; *see* Two phase commit
- Three phase commit
- Abort list, 176, 179, 321
- Abort record, 187, 321
- Aborted transaction, 4, 30, 321
- Access validation, 283–287, 289, 292
- ACP; *see* Atomic commitment protocol
- Active list, 176, 179, 321
- Active transaction, 4, 30, 321
- Adaplex, 162, 210
- After image, 175, 187, 321
- Aggressive scheduler, 47–49, 321
- Agrawal, R., 105, 210
- Alsberg, P.A., 307
- Ancestor, 316, 321
- Anderson, T., 24, 209
- Archive, 177f, 179ff, 207–209, 321–322
- Atomic commitment protocol, 218, 222–260,  
    268, 270, 284, 292ff, 322
- Atomicity, 1–2, 63–64, 171, 322
- Attar, R., 308
- Available copies algorithms, 281–294, 308,  
    322
- Available copy, 269, 322
- Avoid cascading aborts, 9, 34–37, 97, 125,  
    175, 223, 322
- Badal, D.Z., 138
- Balter, R., 105
- Bayer, R., 105, 210
- Beeri, C., 24, 104, 105
- Before image, 10, 175, 187, 190, 322
- Belford, G.G., 138
- Ben-Ari, M., 23
- Berard, P., 105
- Bernstein, P.A., 41, 104, 138, 164, 209, 307,  
    308
- Bertino, E., 104
- Bjork, L.A., 23, 209
- Blaustein, B.T., 308
- Blind write, 42, 322
- Block; *see* Disk page
- Blocking, 225–229, 234, 236, 240–241, 246,  
    255–256
- Blocking policy, 89–91, 105, 322
- Bos, L., 105
- Boyce, R.F., 23
- Breitwieser, H., 308
- Briatico, D., 210
- Brinch Hansen, P., 23
- Broadcasting, 235, 322
- B-tree, 194, 322
- B-tree locking, 99–105
- Buckley, G.N., 105
- Buffer; *see* Cache manager
- Cache consistent checkpointing, 184–185,  
    208, 322
- Cache directory, 173, 322
- Cache manager (CM), 18–20, 171, 173–174,  
    210, 322
- Careful replacement algorithm, 205, 323; *see*  
    *also* Shadowing
- Carey, M.J., 104, 105, 138
- Carlesi, C., 104
- Carlson, C.R., 105
- Casanova, M.A., 138

- Cascadeless; *see* Avoid cascading aborts
- Cascading abort, 6, 210; *see also* Avoid cascading aborts
- Centralized computer system, 17
- Certifier, 128–132, 138, 164, 223, 271, 323
- Certify locks, 156–159
- Chamberlin, D.D., 23
- Chan, A., 164, 210, 308
- Chandy, K.M., 104
- Checkpoint, 183–185, 188, 197, 323; *see also* Archive
- Checkpoint record, 187–190, 209, 323
- Cheng, W.K., 138
- Chesnais, A., 105
- Cheung, D., 260
- Chin, F., 260
- Chung, J., 210
- Ciuffoletti, A., 210
- Class, 121, 138, 311
- CM; *see* Cache manager
- Collmeyer, A.J., 104
- Commit, 3–5, 25–27, 30, 122–123, 125, 128, 179, 217–218, 223, 323; *see also* Atomic commitment protocol
- Commit consistent checkpoint, 184, 190, 207–208, 323
- Commit list, 161–163, 176, 179, 323
- COMMIT message; *see* Two phase commit; Three phase commit
- Commit record, 187, 323
- Committed database state, 169, 323
- Committed projection, 30, 36–37, 147, 169, 323
- Communication failure, 220–222, 225, 240–241, 256, 259, 269, 294–297, 301–302, 305, 323
- Communication topology, 237–239, 323
- Compatibility matrix, 38, 68–69, 74, 138, 323
- Complete history, 29, 323–324
- Complete multiversion history; *see* History, multiversion
- Component, 324; *see also* Partition
- Computer Corporation of America; *see* Adaplex; SDD-1
- Computer network, 218–219
- Concurrency control, 1–3, 11–17, 44
- Conflict based schedulers, 41, 324
- Conflict equivalence; *see* Equivalence
- Conflict graph, 121, 138, 139
- Conflict serializable (CSR) history, 40
- Conflicting operations, 28–29, 37–38, 48, 145, 149, 272
- Conflicting transactions, 33ff
- Connected graph, 313, 324
- Conservative scheduler, 47–49, 118–121, 324; *see also* Two phase locking, conservative; Timestamp ordering, conservative
- Conservative serialization graph testing, 123–125, 324
- Conservative TO; *see* Timestamp ordering, conservative
- Conservative 2PL; *see* Two phase locking, conservative
- Consistency preservation, 15, 167
- Conversion; *see* Lock conversion
- Cooper, E.C., 260
- Cooperative termination protocol, 228–229, 231, 233, 236–237, 260, 324; *see also* Termination protocol
- Coordinator, 224, 324; *see also* Atomic commitment protocol
- Copier transaction, 207ff, 279, 303, 324
- Cristian, F., 308
- Critical section, 16–17, 324
- CSR; *see* Conflict serializable
- Cycle, 315, 324
- Cyclic restart, 58, 84–86, 324
  
- Dag; *see* Directed acyclic graph
- Dag locking, 98–99, 324
- Data contention, 87, 91–93, 325
- Data items, 2, 325
- Data manager (DM), 20, 169–174, 268
- Data tree, 96, 325
- Database, 2, 324
- Database machines, 210
- Database management system (DBMS), 1, 2ff
- Database manager, 17ff, 325
- Database operations, 2, 325
- Database system (DBS), 2, 17–23, 325
- Davidson, S.B., 308
- Davies, C.T., 23
- Day, J.D., 307
- DBMS; *see* Database management system
- DBS; *see* Database system
- DB2, 87, 210
- DC-thrashing; *see* Thrashing; data contention
- Deadlock, 104–105, 325
  - centralized, 52–53, 56–58
  - distributed, 79–87, 126, 325
  - effect on performance, 88
  - phantom, 80–81, 331
  - see also* Timestamp-based deadlock prevention
- Decentralized 2PC protocol, 237–240, 325
- Decision; *see* Atomic commitment protocol
- DECISION-REQ message; *see* Cooperative termination protocol
- Decitre, P., 105
- Decrement operation, 37–38, 68
- Deferred output, 8, 325
- Deferred writing, 270–271, 325
- Degree 3 consistent, 11ff
- Delay, 20, 47–48
- Delayed commit, 186, 325
- Delayed writes, 166
- Descendant, 316, 325

- Devor, C., 105
- DeWitt, D.J., 210
- Digraph; *see* Directed graph
- Directed acyclic graph, 315–319, 325; *see also* History
- Directed graph, 314–319, 325
- Directory, 325
  - for replicated data, 268–269, 289–294, 305
  - for shadowing, 172–174, 201–205
- Directory-oriented available copies algorithm, 289–294, 305, 325; *see also* Available copies algorithms
- Dirty bit, 173–174, 182, 208, 325
- Disk page, 171–172
- Distributed computer system, 17
- Distributed database system (DDBS), 22–23, 326
- Distributed deadlock; *see* Deadlock, distributed
- Distributed scheduler, 125, 131–132; *see also* Two phase locking, distributed; Timestamp ordering, distributed
- Distributed transaction (DT) log, 224, 227, 229–235, 249–253, 255, 260, 326
- Distributed two phase locking, *see* Two phase locking, distributed
- DL; *see* Dag locking
- DM; *see* Data manager
- DT log; *see* Distributed transaction log
- Dubourdieu, D.J., 164, 210, 260
- Dump; *see* Archive
- Dwork, C., 260
  
- Eager, D.L., 308
- Edelberg, M., 104
- Edge, 313, 315
- Effelsberg, W., 210
- El Abbadi, A., 308
- Election protocol, 244, 245, 254–256, 259, 326
- Elhardt, K., 210
- Ellis, C.S., 105
- Equivalence, 326
  - of histories, 30–31
  - of MV histories, 145–146, 148–149
  - of RD histories, 273–274
  - see also* View equivalence
- Error detection, 168, 171, 222, 269
- Eswaran, K.P., 23, 104
- Exclude; *see* Directory-oriented available copies algorithm
- Explicit lock, 71–73
  
- Fail-stop, 219, 269, 326
- Failure, 167–168, 218–222, 269; *see also* Communication failure; Error detection; Site failure; System failure
- Failure-recovery serialization graph (FRSG), 280–281, 288, 309, 310, 326
  
- Fair scheduling, 62
- Fault tolerance; *see* Recovery
- Faults; *see* Failures
- Fetch, 19, 171, 174, 326
- File systems, 1
- Final write, 39, 273–274, 326
- First-in-first-out replacement strategy, 174
- Flush, 19, 171, 174, 176–178, 180, 327
- Force, 177ff; *see also* Redo
- Ford, R., 105
- Franaszek, P., 105
- FRSG; *see* Failure-recovery serialization graph
- Fuzzy checkpointing, 185, 209, 327
  
- Galler, B.I., 105
- Garbage Collection Rule, 178–179, 183, 200, 327
- Garbage collection of DT log, 231, 233; *see also* Distributed transaction log
- Garcia-Molina, H., 104, 210, 260, 308
- Garey, M.R., 42
- Gawlick, D., 104
- Gelenbe, E., 105
- Gifford, D.K., 308
- Gligor, V.D., 104
- Global deadlock detection; *see* Deadlock, distributed
- Global waits-for graph; *see* Waits-for graph, distributed
- Goodman, N., 104, 105, 138, 164, 209, 307, 308
- Graham, M.H., 24
- Granularity
  - of data items, 2, 171–172, 327
  - of locks, 69–70, 92–94
- Granularity curve, 93, 327
- Gray, J.N., 11ff, 17ff, 23, 24, 41, 104, 105, 210, 260
- Gray, R., 164
- Griffith, N., 24, 210
- Group commit, 186, 327
  
- Haas, L.M., 104
- Hadzilacos, T., 138, 164
- Hadzilacos, V., 41, 209, 210
- Haerder, T., 138, 177ff, 209, 210
- Hammer, M., 260
- Handshake, 22, 115, 120, 124, 137, 327; *see also* Redo Rule; Undo Rule
- Heller, H., 164
- Herlihy, M., 308
- Hierarchical deadlock detection, 105; *see also* Deadlock, distributed
- History, 28–30, 327
  - replicated data, 271–275, 332
  - for 2PL, 53–56
  - multiversion, 144–149, 328
- Holt, R.C., 23, 104
- Home site, 217–218, 223, 282, 304

- Horning, J.J., 24  
 Hot spot, 67–69, 94–95, 104, 327  
 Hunt, H.B., 104
- Ibaraki, T., 164  
 IBM; *see* System R; IMS  
 Idempotence, 179–180, 183, 190, 327; *see also* Restart  
 Immediate writing, 270, 327  
 Implicit lock, 71–73  
 IMS, 67, 105, 210  
 Include; *see* Directory-oriented available copies algorithm  
 Incomparability of sets, 35  
 Inconsistent retrieval, 12–14, 327  
 Increment operation, 37–38, 68  
 Indefinite postponement, 62  
 Independent failure modes, 206, 327  
 Independent recovery, 225, 229, 328  
 Index entry, 66, 328  
 Index locking, 66–67, 76–77, 328  
 INGRES, 87, 210  
 In-place updating, 172, 175, 328  
 Input statements, 3–4, 28  
 Insert, 66  
   in a B-tree, 101–103; *see also* Phantom problem  
 Integrated schedulers, 132–138, 328  
 Intention lock, 72, 328; *see also* Multi-granularity locking  
 Intentions list, 200–201, 270, 328  
 Interference between concurrent programs, 11–13, 328  
 Interleaved executions, 11, 44, 328  
 Isloor, S.S., 104
- Jipping, M., 105  
 Johnson, D.S., 42
- Kameda, T., 164, 260  
 Kaneko, A., 138  
 Kanellakis, P.C., 164  
 Kawazu, S., 104  
 Kedem, Z.M., 105  
 Kent, J., 210  
 Kersten, M., 138  
 Kiessling, W., 104  
 King, P.F., 104  
 Kinkade, D., 104  
 Korth, H.F., 104  
 Krenz, G., 105  
 Kung, H.T., 105, 138  
 Kwong, Y.S., 105
- Ladner, R.E., 105  
 Lai, M.Y., 138, 164  
 Lamson, B.W., 210, 260  
 Landes, O.E., 210  
 Landherr, G., 104
- Last committed value, 169, 328  
 Lausen, G., 164  
 Lavenberg, S.S., 105  
 Least recently used, 174, 328  
 Leblanc, Ph., 105  
 Lee, P.A., 24, 209  
 Lehman, P.L., 105  
 LeLann, G., 260  
 Leszak, M., 308  
 Lewis, P.M., 23, 41, 104, 260  
 Lin, W.K., 104, 105  
 Lindsay, B.G., 210, 260  
 Linear 2PC, 238–240, 328  
 Links for B-tree locking, 102–105  
 Liskov, B., 24  
 Live-lock; *see* Cyclic restart  
 Livny, M., 105  
 Lock table, 61–63, 107  
 Lock conversion, 53, 74–75, 101–102, 157, 324  
 Lock coupling, 96, 105, 328; *see also* Tree locking  
 Lock escalation, 75–76, 328; *see also* Multi-granularity locking  
 Lock instance graph, 71, 77, 96ff, 328  
 Lock manager, 60–64, 107, 328  
 Lock type graph, 70–71, 76–77, 328  
 Locked point, 106, 328  
 Locking, *see* Two phase locking  
 Locking granularity; *see* Granularity, of locks; Multigranularity locking  
 Log sequence number, 186, 192–195, 210, 328  
 Logging, 175–179, 186–195, 207–209, 213–214, 224, 297; *see also* Distributed transaction log  
 Logical logging, 176, 191–194, 213–214, 328  
 Lomet, D.B., 24, 104  
 Lorie, R.A., 104, 209, 210  
 Lost updates, 12–14, 328  
 LSN; *see* Log sequence number  
 Lynch, N.S., 24
- McLean, G., Jr., 138  
 Macri, P.M., 104  
 MADMAN, 87  
 Majority consensus, 256–258, 307, 329; *see also* Quorum consensus  
 Manber, U., 105  
 Marsland, T., 104  
 Media failures, 19–20, 168, 171, 177, 179ff, 206–209, 329  
 Menasce, D.A., 104, 138, 210  
 Messages, 6, 24, 220, 329  
 MGL; *see* Multigranularity locking  
 Might-write lock, 102  
 Miller, J.A., 210  
 Millstein, R.E., 138, 307

- Mirroring, 206–207, 329
- Misra, J., 104
- Missing writes algorithm, 301–304, 308, 329
- Missing writes validation, 283–287, 289, 292
- Mitra, D., 105
- Mitrani, I., 105
- Mixed integrated schedulers, 133, 135–138, 160–164, 329; *see also* Integrated schedulers
- Mohan, C., 260
- Monitor, 62–63
- Moore, J.E., 24
- Morris, R.J.T., 138, 105
- Moss, J.E.B., 24
- MPL; *see* Multiprogramming level
- Mueller, E., 24
- Multigranularity locking (MGL), 69–77, 104, 107–108, 138, 329
- Multiprogramming level (MPL), 89, 329
- Multiversion (MV) history; *see* History, multiversion
- Multiversion scheduler, 143–164, 197–198, 200ff, 329
- Multiversion serialization graph (MVSG), 146, 149–153, 329–330; *see also* Serialization graph
- Multiversion timestamp ordering (MVTO); *see* Timestamp ordering, multiversion
- Muntz, R.R., 104
- Munz, R., 105
- Mutual exclusion, 16–17, 330
- MV history; *see* History, multiversion
- MVSG; *see* Multiversion serialization graph
- MVTO; *see* Timestamp ordering, multiversion
- Nakanishi, T., 138
- Nested transactions, 24, 101ff
- Network partition, 220–221, 294, 305, 308, 311, 330; *see also* Communication failure
- Network topology, 241
- Node, 313–315
- Nolte, J., 104, 105
- Non-blocking property, 241–244, 247–248; *see also* Uncertainty period; Blocking
- No-undo/no-redo algorithm, 156ff, 201–205
- No-undo/redo algorithm, 198–201
- Null operation, 120
- Obermarck, R., 104, 105
- 1SR; *see* One-copy serializability
- Ong, K.S., 210
- One-copy history; *see* History
- One-copy serializability, 330
  - in MV histories, 150–153, 155–156, 159
  - in RD histories, 266–267, 274–281, 285–288, 292–293, 296–297, 301, 307, 308
- Operating systems, 23, 88, 97, 168
- Operations, 2
- Optimistic scheduler; *see* Certifier
- Ordering operations, 21–22; *see also* Handshake
- Ordering transactions, 15–16
- Output statement, 4, 28
- Page; *see* Disk page
- Papadimitriou, C.H., 41, 42, 104, 164
- Partial data item logging, 186–191, 330
- Partial failures, 218–219, 330
- Partial order, 26–28, 319, 330; *see also* History
- Participant; *see* Atomic commitment protocol
- Partition, 313–314, 330; *see also* Network partition
- Path, 313–315, 330
- Path pushing; *see* Deadlock, distributed
- Peinl, P., 104
- Penultimate checkpoint, 185, 189, 193, 331
- Performance
  - of aggressive vs. conservative schedulers, 48
  - of recovery, 210
  - of three phase commit, 240, 255
  - of tree locking, 97
  - of two phase commit, 234–240
  - see also* Two phase locking, performance of
- Performance failures; *see* Timeout failures
- Peterson, R.J., 210
- Phantom deadlocks; *see* Deadlock, phantom
- Phantom problem, 64–67, 104, 331
- Physical logging, 176, 194, 210, 331
- Piggybacking, 80, 298, 331
- Pin, 174, 182, 331
- Popek, G., 24
- Potier, D., 105
- PRE-COMMIT message; *see* Three phase commit
- Predeclaration, 49, 58, 119, 121, 123, 331
- Predicate locking, 67, 331
- Prefix commit-closed, 36–37, 44, 150ff, 331
- Prefix of a partial order, 319, 331
- Preserves reflexive reads-from relationships, 147, 155, 159, 273, 331
- Preserve database consistency; *see* Consistency preservation
- Primary copy, 271, 331
- Prime Computer, 162, 164, 198, 210, 260
- Programming language, 4
- Pure integrated schedulers, 133, 135, 331; *see also* Integrated schedulers
- Pure restart policy, 89–91, 105, 331
- Purging versions, 144, 162–163
- Putzolu, G.R., 104
- QC; *see* Quorum consensus algorithm
- Queries, 5, 104, 160–164, 296–297, 331
- Queues, 22

- Quorum, 260, 295–296, 298–304, 331  
 Quorum consensus (QC) algorithm, 296–304, 331
- Ramarao, K.V.S., 260  
 Rappaport, R.L., 210  
 RC-thrashing; *see* Thrashing  
 RDSDG; *see* Replicated data serialization graph  
 Read, 2, 26, 28, 179, 217, 266, 304, 332  
 Read lock, 50, 60, 94, 98–99, 234, 332  
 Read order, 276–277, 301, 303–304, 332  
 Reads-from relationship, 7, 34, 332  
   in MV history, 148–149  
   in RD history, 272–275  
   *see also* View equivalence  
 Readset, 48, 121, 126, 127, 129  
 Ready, 119, 124  
 Record level locking, 194–195  
 Recoverability, 6–11, 34–37, 125, 130, 153, 175, 210, 223, 268, 273, 332  
 Recovery, 1, 6  
   centralized, 167–209  
   distributed, 217–260  
   *see also* Recoverability  
 Recovery block, 24  
 Recovery manager (RM), 19–20, 171, 174–180, 332  
 Recovery procedure; *see* Restart  
 Redo, 177–178, 192, 194  
 Redo Rule, 178, 183, 186, 197, 200, 205, 223, 332  
 Reed, D.P., 24, 164  
 Reiser, A., 164  
 Reject, 20, 47–48  
 Replacement strategy, 174, 176, 180, 332  
 Replicated data, 138, 265–308  
 Replicated data history; *see* History, replicated data  
 Replicated data serialization graph (RDSDG), 276–277, 301, 303–304, 309, 333  
 Resource contention, 87–93, 105, 333  
 Restart, 169–170, 175, 177–184, 188–196, 198, 199, 204, 219, 234, 333  
 Reuter, A., 104, 177ff, 209, 210  
 Ries, D., 104, 105  
 RM; *see* Recovery manager  
 RM-Abort, 175, 204, 181, 190, 198, 199  
 RM-Commit, 175, 181, 182, 196–197, 199, 201, 202, 204, 205  
 RM-Read, 174, 181, 198–199, 201, 204, 205  
 RM-Write, 174, 181, 182, 186–187, 193, 198, 204  
 Robinson, J.T., 105, 138  
 Root, 315  
 Rosenkrantz, D.J., 23, 41, 104, 164, 260  
 Rothnie, J.B., Jr., 138, 307  
 Rounds, 234–236, 238, 239, 255  
 rw phased, 141  
 rw synchronization, 133, 333  
 Ryu, I.K., 105
- Samadi, B., 105  
 Sargeant, G., 308  
 Scheduler, 20, 21, 28, 333; *see also* Multi-version scheduler; Serialization graph testing; Timestamp ordering; Two phase locking  
 Scheifler, R., 24  
 Schkolnick, M., 105  
 Schlageter, G., 138  
 Schultz, T., 105  
 SDD-1, 138  
 Sector, 210  
 Semaphores, 63, 175  
 Sequential operations, 2  
 Sequin, J., 308  
 Serial history, 13, 31, 150, 274, 333  
 Serializability, 11–17; *see also* Serializability theory  
 Serializability Theorem, 32–34, 114, 129, 333  
 Serializability theory, 25–41, 308  
   for MV histories, 144–153, 158–159, 164  
   in proving correctness of 2PL, 53–56  
   for RD histories, 271–281, 286–288, 299–301  
 Serialization graph (SG), 32–33, 42, 57ff, 136, 333–334  
   in proving TO correct, 114  
   in proving 2PL correct, 54–56, 129  
   for RD histories, 273–275, 303–304  
   and tree locking, 96–97  
   *see also* Replicated data serialization graph; Failure-recovery serialization graph; Stored serialization graph  
 Serialization graph testing (SGT), 121–128, 133, 138, 334  
 Sevcik, K.C., 308  
 SG; *see* Serialization graph  
 SGT; *see* Serialization graph testing  
 Shadow version algorithm, 205, 334; *see also* Shadowing  
 Shadowing, 172, 174, 192, 201–205, 208–209, 334  
 Shapiro, R.M., 138, 307  
 Shasha, D., 105  
 Shattuck, S.H., 104  
 Shipman, D.W., 41, 104, 138, 260, 307  
 Shrivastava, S.K., 24, 209  
 Shum, A.W., 105  
 Siewiorek, D.P., 24, 209  
 Silberschatz, A., 105, 164  
 Simoncini, L., 210  
 Sirius-Delta, 260  
 Site failure, 219–220, 226, 240, 241, 254, 269, 305, 334; *see also* Failure  
 Site quorum, 295–297, 334; *see also* Quorum

- Sites, 22
- Skeen, D., 260, 308
- Source, 315, 334
- Spirakis, P.G., 105
- Spontaneous abort, 80–81
- Spooler, 309
- SSG; *see* Stored serialization graph
- Stable checkpoint; *see* Checkpoint
- Stable database, 172, 177, 334
- Stable-LSN, 187, 188, 190, 334
- Stable storage, 18, 168, 171–174, 178–180, 184, 202, 334
- Start, 3, 334
- State, 2
- Static 2PL; *see* Two phase locking, conservative
- Steal, 177ff; *see also* Undo
- Stearns, R.E., 23, 41, 104, 164, 260
- Stonebraker, M., 104, 105, 210, 307
- Stored serialization graph (SSG), 122–128, 130–131, 334
- Strickland, J.P., 210
- Strict, 9–11, 34–37, 175, 335
- Strict TO; *see* Timestamp ordering, strict
- Strict 2PL; *see* Two phase locking, strict
- Sturgis, H., 210, 260
- Subgraph, 313, 335
- Suri, R., 104, 105
- Swarz, R.S., 209
- Synapse, 210
- System failure, 19, 168, 177, 179, 182, 206, 335; *see also* Failure
- System R, 87, 105, 192, 210
  
- Tandem Computers, 56
- Tay, Y.C., 104, 105
- Tebra, H., 138
- Terminal I/O, 8, 16ff
- Termination protocol, 228, 244–248, 252–260, 263, 335; *see also* Cooperative termination protocol
- Termination rule, 244–245, 258
- Thanos, C., 104
- Thomas, R.H., 138, 307
- Thomas' Write Rule (TWR), 134–138, 160, 335
- Thomasian, A., 105
- Thrashing, 87–88, 91–93, 105, 325, 332, 335
- Three phase commit (3PC), 240–260, 335
- Timeout, 56, 222, 231, 235, 269, 335
- Timeout action, 335
- for three phase commit, 243–246
- for two phase commit, 227–229, 231–233
- Timeout failures, 222, 306ff, 335
- Timestamp, 85–87, 117–118, 136–137, 160, 335
- Timestamp interval, 154
- Timestamp ordering (TO), 114–121, 123, 131, 133, 335
- Timestamp ordering (*continued*)
  - conservative, 118–121, 324
  - distributed, 136–138
  - multiversion, 153–156, 160–161, 164, 330
  - strict, 116–117, 335
- Timestamp table, 117–118
- Timestamp-based deadlock prevention, 84–87, 115, 139, 164
- Tirri, T., 104
- TL, *see* Tree locking
- TM; *see* Transaction manager
- TO; *see* Timestamp ordering
- TO rule, 114, 116, 120, 131, 335
- Topological sort, 33, 317, 335
- Total failure, 219, 226, 240, 245, 249–255, 294, 335
- Toueg, S., 308
- Traiger, I.L., 23, 210
- Transaction, 1–6, 25–28, 146–148, 335–336
- Transaction class; *see* Class
- Transaction failure; *see* Abort; Failure
- Translation function, 146–147, 271
- Transaction identifier, 3, 61, 83, 162, 176
- Transaction manager (TM), 17ff, 21, 23, 268–269, 336
- Transaction processing systems, 1
- Transitive closure, 127, 317–319, 336
- Tree locking (TL), 95–104, 164, 336
- Two phase commit (2PC), 132, 223, 226–240, 260, 336
- Two phase locking, 16ff, 49–95, 117, 123, 133, 275, 336
  - conservative, 58–59, 91, 105, 119, 324
  - distributed, 77–78, 118, 136–138
  - implementation of, 60–64
  - multiversion, 156–159
  - performance of, 69–70, 87–95, 104, 109–111
  - strict, 59–60, 78, 135–138, 160–161, 281–282, 285, 286, 290ff, 335; *see also* Certifier; Deadlock, distributed
- Two phase rule, 51, 55, 78, 105, 117, 336
- Two version two phase locking; *see* Two phase locking, multiversion
- TWR; *see* Thomas' Write Rule
  
- Ullman, J.D., 104
- Uncertainty period, 225, 227–229, 231, 241–245, 247–249, 336
- Uncommitted, 4
- Undeliverable message, 221–222
- Undirected graph, 313–314, 336
- Undo, 6, 177–178, 191–192, 194–195
- Undo Rule, 177, 183, 186, 197, 200, 205, 337
- Undo/no-redo algorithm, 177, 196–198
- Undo/redo algorithm, 177, 180–195
- Uninterpreted, 27–28, 337
- Unpin, 174, 337

- Update record, 187, 191, 192, 337
- Updater, 5, 104, 160–164, 337
- Validation protocol; *see* Missing writes validation; Access validation
- Value of a data item, 2
- Venn diagram, 35, 42
- Verhofstad, J.S.M., 209, 210
- Version; *see* Multiversion scheduler
- Version number, 298, 300, 303, 307, 337
- Version order, 151–153, 155, 160–161, 337
- Version order edge, 152–153, 159, 337
- Victim selection, 57–58, 80
- View; *see* Virtual partition algorithm
- View equivalence, 38–41, 145–146, 148, 273, 337
- View graph, 43
- View serializability, 39–41, 337
- Virtual partition (VP) algorithm, 296–297, 304–307, 308, 337
- Volatile storage, 18, 168, 173, 337
- Vote; *see* Atomic commitment protocol
- VOTE-REQ message; *see* Two phase commit; Three phase commit
- VP; *see* Virtual partition algorithm
- VSR; *see* View serializability
- Wait-die deadlock prevention; *see* Timestamp-based deadlock prevention
- Waits-for graph (WFG), 56–57, 79–82, 105–126, 337
- Weihl, W.E., 164
- Weinberger, P.J., 104, 105
- WFG; *see* Waits-for graph
- Wiederhold, G., 308
- Wilkinson, W.K., 138, 164
- Wilnes, P., 308
- Wong, K.C., 41, 104
- Wong, W.S., 104, 105, 138
- Wood, D., 105
- Workload, 92–95, 325
- Workspace model, 141, 327
- Wound-wait; *see* Timestamp-based deadlock prevention
- Wright, D., 308
- Write, 2, 26, 28, 171–172, 179, 217, 266, 269–272, 337
- Write ahead log protocol, 177ff; *see also* Undo Rule
- Write-all approach, 266–267, 299, 304, 337
- Write-all-available approach, 267–268, 276, 281, 337
- Write lock, 50, 60, 337
- Write order, 276–277, 301, 303, 337
- Writerset, 48, 121, 126, 127, 129, 337
- ww synchronization, 133, 337
- Yannakakis, M., 41, 138
- Yao, S.B., 105