

Wang, Jiaqi; Talmage, Edward; Lee, Hyunyoung; Welch, Jennifer L.

Improved time bounds for linearizable implementations of abstract data types. (English)

Zbl 1407.68313

Inf. Comput. 263, 1-30 (2018).

Summary: Linearizability is a well-known consistency condition for shared objects in concurrent systems. We focus on the problem of implementing linearizable objects of arbitrary data types in message-passing systems with bounded, but uncertain, message delay and bounded, but non-zero, clock skew. We present an algorithm that exploits axiomatic properties of different operations to reduce the running time of each operation below that obtainable with previously known algorithms. We also prove lower bounds on the time complexity of various kinds of operations, specified by the axioms they satisfy, resulting in reduced gaps in some cases and tight bounds in others.

MSC:

[68Q65](#) Abstract data types; algebraic specification

[68M14](#) Distributed systems

[68Q17](#) Computational difficulty of problems (lower bounds, completeness, difficulty of approximation, etc.)

Cited in 1 Document

Keywords:

[abstract data types](#); [distributed computing](#); [linearizability](#); [time bounds](#)

Software:

[Dynamo](#)

Full Text: [DOI](#)

References:

- [1] Herlihy, M.; Wing, J. M., Linearizability: a correctness condition for concurrent objects, ACM Trans. Program. Lang. Syst., 12, 3, 463-492, (1990)
- [2] Lamport, L., On interprocess communication. Part I: basic formalism, Distrib. Comput., 1, 2, 77-85, (1986) · [Zbl 0598.68022](#)
- [3] Lynch, N. A., Distributed Algorithms, (1996), Morgan Kaufmann · [Zbl 0877.68061](#)
- [4] Attiya, H.; Welch, J. L., Sequential consistency versus linearizability, ACM Trans. Comput. Syst., 12, 2, 91-122, (1994)
- [5] Lipton, R. J.; Sandberg, J. S., PRAM: A Scalable Shared Memory, (September 1988), Princeton University, Department of Computer Science, Tech. Rep. CS-TR-180-88
- [6] Mavronicolas, M.; Roth, D., Linearizable read/write objects, Theor. Comput. Sci., 220, 1, 267-319, (1999) · [Zbl 0916.68014](#)
- [7] Chaudhuri, S.; Gawlick, R.; Lynch, N. A., Designing algorithms for distributed systems with partially synchronized clocks, (Anderson, J.; Toueg, S., Proceedings of the Twelfth Annual ACM Symposium on Principles of Distributed Computing, Ithaca, New York, USA, August 15-18, 1993, (1993), ACM), 121-132 · [Zbl 1375.68195](#)
- [8] Weihl, W. E., Commutativity-based concurrency control for abstract data types, IEEE Trans. Comput., 37, 12, 1488-1505, (1988) · [Zbl 0674.68016](#)
- [9] Kosa, M. J., Time bounds for strong and hybrid consistency for arbitrary abstract data types, Chic. J. Theor. Comput. Sci., 1999, (1999) · [Zbl 0924.68081](#)
- [10] Lundelius, J.; Lynch, N. A., An upper and lower bound for clock synchronization, Inf. Control, 62, 2/3, 190-204, (1984) · [Zbl 0591.68023](#)
- [11] Wang, J.; Welch, J. L.; Lee, H., Brief announcement: time bounds for shared objects in partially synchronous systems, (Gavoille, C.; Fraigniaud, P., Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011, (2011), ACM), 347-348
- [12] Wang, J.; Talmage, E.; Lee, H.; Welch, J. L., Improved time bounds for linearizable implementations of abstract data types, (2014 IEEE 28th International Parallel and Distributed Processing Symposium, Phoenix, AZ, USA, May 19-23, 2014, (2014), IEEE Computer Society), 691-701
- [13] Biaz, S.; Welch, J. L., Closed form bounds for clock synchronization under simple uncertainty assumptions, Inf. Process. Lett.,

80, 3, 151-157, (2001) · [Zbl 1032.68512](#)

- [14] Fan, R.; Lynch, N. A., Gradient clock synchronization, *Distrib. Comput.*, 18, 4, 255-266, (2006) · [Zbl 1266.68049](#)
- [15] Lenzen, C.; Locher, T.; Wattenhofer, R., Clock synchronization with bounded global and local skew, (49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA, (2008), IEEE Computer Society), 509-518
- [16] Lenzen, C.; Locher, T.; Wattenhofer, R., Tight bounds for clock synchronization, *J. ACM*, 57, 2, (2010) · [Zbl 1327.68045](#)
- [17] Attiya, H.; Herzberg, A.; Rajsbaum, S., Optimal clock synchronization under different delay assumptions, *SIAM J. Comput.*, 25, 2, 369-389, (1996) · [Zbl 0844.68043](#)
- [18] Attiya, H.; Hay, D.; Welch, J. L., Optimal clock synchronization under energy constraints in wireless ad-hoc networks, (Anderson, J. H.; Prencipe, G.; Wattenhofer, R., Principles of Distributed Systems, 9th International Conference, OPODIS 2005, Pisa, Italy, December 12-14, 2005, Revised Selected Papers, Lecture Notes in Computer Science, vol. 3974, (2005), Springer), 221-234
- [19] Arjomandi, E.; Fischer, M. J.; Lynch, N. A., Efficiency of synchronous versus asynchronous distributed systems, *J. ACM*, 30, 3, 449-456, (1983) · [Zbl 0627.68025](#)
- [20] Attiya, H.; Mavronicolas, M., Efficiency of semisynchronous versus asynchronous networks, *Math. Syst. Theory*, 27, 6, 547-571, (1994) · [Zbl 0812.68078](#)
- [21] Rhee, I.; Welch, J. L., The impact of timing knowledge on the session problem, *SIAM J. Comput.*, 32, 4, 1007-1039, (2003) · [Zbl 1029.68073](#)
- [22] Dolev, D.; Halpern, J. Y.; Strong, H. R., On the possibility and impossibility of achieving clock synchronization, *J. Comput. Syst. Sci.*, 32, 2, 230-250, (1986) · [Zbl 0595.68029](#)
- [23] Fischer, M. J.; Lynch, N. A.; Merritt, M., Easy impossibility proofs for distributed consensus problems, *Distrib. Comput.*, 1, 1, 26-39, (1986) · [Zbl 0598.68024](#)
- [24] DeCandia, G.; Hastorun, D.; Jampani, M.; Kakulapati, G.; Lakshman, A.; Pilchin, A.; Sivasubramanian, S.; Voshall, P.; Vogels, W., Dynamo: amazon's highly available key-value store, (Bressoud, T. C.; Kaashoek, M. F., Proceedings of the 21st ACM Symposium on Operating Systems Principles 2007, SOSP 2007, Stevenson, Washington, USA, October 14-17, 2007, (2007), ACM), 205-220
- [25] Corbett, J. C.; Dean, J.; Epstein, M.; Fikes, A.; Frost, C.; Furman, J. J.; Ghemawat, S.; Gubarev, A.; Heiser, C.; Hochschild, P.; Hsieh, W. C.; Kanthak, S.; Kogan, E.; Li, H.; Lloyd, A.; Melnik, S.; Mwaure, D.; Nagle, D.; Quinlan, S.; Rao, R.; Rolig, L.; Saito, Y.; Szymaniak, M.; Taylor, C.; Wang, R.; Woodford, D., Spanner: google's globally-distributed database, (Thekkath, C.; Vahdat, A., 10th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2012, Hollywood, CA, USA, October 8-10, 2012, (2012), USENIX Association), 261-264
- [26] Li, C.; Porto, D.; Clement, A.; Gehrke, J.; Prego, N. M.; Rodrigues, R., Making geo-replicated systems fast as possible, consistent when necessary, (Thekkath, C.; Vahdat, A., 10th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2012, Hollywood, CA, USA, October 8-10, 2012, (2012), USENIX Association), 265-278
- [27] Vogels, W., Eventually consistent, *Commun. ACM*, 52, 1, 40-44, (2009)
- [28] Afek, Y.; Korland, G.; Yanovsky, E., Quasi-linearizability: relaxed consistency for improved concurrency, (Lu, C.; Masuzawa, T.; Mosbah, M., Principles of Distributed Systems - 14th International Conference, Proceedings, OPODIS 2010, Tozeur, Tunisia, December 14-17, 2010, Lecture Notes in Computer Science, vol. 6490, (2010), Springer), 395-410
- [29] Henzinger, T. A.; Kirsch, C. M.; Payer, H.; Sezgin, A.; Sokolova, A., Quantitative relaxation of concurrent data structures, (Giacobazzi, R.; Cousot, R., The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23-25, 2013, (2013), ACM), 317-328 · [Zbl 1301.68176](#)
- [30] Talmage, E.; Welch, J. L., Improving average performance by relaxing distributed data structures, (Kuhn, F., Distributed Computing - 28th International Symposium, Proceedings, DISC 2014, Austin, TX, USA, October 12-15, 2014, Lecture Notes in Computer Science, vol. 8784, (2014), Springer), 421-438
- [31] Talmage, E.; Welch, J. L., Anomalies and similarities among consensus numbers of relaxed queues, (Abadi, A. E.; Garbinato, B., International Conference on Networked Systems, NETYS 2017, Marrakech, Morocco, May 17-19, 2017, Lecture Notes in Computer Science, vol. 10299, (2017), Springer)
- [32] Shavit, N.; Taubenfeld, G., The computability of relaxed data structures: queues and stacks as examples, (Scheideler, C., Structural Information and Communication Complexity - 22nd International Colloquium, Post-Proceedings, SIROCCO 2015, Montserrat, Spain, July 14-16, 2015, Lecture Notes in Computer Science, vol. 9439, (2015), Springer), 414-428 · [Zbl 06527705](#)
- [33] Attiya, H.; Welch, J., *Distributed Computing*, (2004), John Wiley & Sons, Inc.

This reference list is based on information provided by the publisher or from digital mathematics libraries. Its items are heuristically matched to zbMATH identifiers and may contain data conversion errors. It attempts to reflect the references listed in the original paper as accurately as possible without claiming the completeness or perfect precision of the matching.