

- [36] M. Tamer Özsu and P Valduriez. *Principles of Distributed Database Systems*. 3rd ed. Springer, 2011. ISBN: 9781441988331.
- [37] Matthew Perron, Raul Castro Fernandez, David DeWitt, and Samuel Madden. “Starling: A Scalable Query Engine on Cloud Function Services.” In: *SIGMOD*. 2020.
- [38] Qifan Pu, U C Berkeley, Shivaram Venkataraman, Ion Stoica, U C Berkeley, and Implementation Nsdi. “Shuffling , Fast and Slow: Scalable Analytics on Serverless Infrastructure.” In: *NSDI*. 2019.
- [39] Wolf Rödiger, Tobias Mühlbauer, Alfons Kemper, and Thomas Neumann. “High-Speed Query Processing over High-Speed Networks.” In: *PVLDB* 9.4 (2015). DOI: [10.14778/2856318.2856319](https://doi.org/10.14778/2856318.2856319).
- [40] Josep Sampé, Gil Vernik, Marc Sánchez-Artigas, and Pedro García-López. “Serverless data analytics in the IBM cloud.” In: *Middleware Industry*. 2018. DOI: [10.1145/3284028.3284029](https://doi.org/10.1145/3284028.3284029).
- [41] Venkat Sowrirajan, Bharath Bhushan, and Mayank Ahuja. *Qubole offers Apache Spark on AWS Lambda*. 2017. URL: <https://www.qubole.com/blog/spark-on-aws-lambda/> (visited on 12/20/2019).
- [42] Transaction Processing Performance Council. *TPC Benchmark H (Revision 2.18)*. 2018.
- [43] Liang Wang, Mengyuan Li, Yinqian Zhang, Thomas Ristenpart, and Michael Swift. “Peeking Behind the Curtains of Serverless Platforms.” In: *USENIX ATC*. 2018.
- [44] Matei Zaharia et al. “Resilient Distributed Datasets: A Fault-tolerant Abstraction for In-memory Cluster Computing.” In: *NSDI*. 2012.