





































13. Kolev, B., Pau, R., Levchenko, O., Valduriez, P., Jiménez-Peris, R., Pereira, J.: Benchmarking polystores: the cloudmssql experience. In: *Big Data (Big Data)*, 2016 IEEE International Conference on. pp. 2574–2579. IEEE (2016)
14. Kolev, B., Valduriez, P., Bondiombouy, C., Jiménez-Peris, R., Pau, R., Pereira, J.: Cloudmssql: querying heterogeneous cloud data stores with a common language. *Distributed and parallel databases* **34**(4), 463–503 (2016)
15. LeFevre, J., Sankaranarayanan, J., Hacigumus, H., Tatemura, J., Polyzotis, N., Carey, M.J.: Miso: souping up big data query processing with a multistore system. In: *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. pp. 1591–1602. ACM (2014)
16. Leskovec, J., Sosič, R.: Snap: A general-purpose network analysis and graph-mining library. *ACM Transactions on Intelligent Systems and Technology (TIST)* **8**(1), 1 (2016)
17. Lu, J.: Towards benchmarking multi-model databases. In: *CIDR* (2017)
18. Lu, J., Holubová, I.: Multi-model data management: What’s new and what’s next? In: *EDBT*. pp. 602–605 (2017)
19. Lu, J., Liu, Z.H., Xu, P., Zhang, C.: Udbms: road to unification for multi-model data management. *arXiv preprint arXiv:1612.08050* (2016)
20. Palkar, S., Thomas, J.J., Shanbhag, A., Narayanan, D., Pirk, H., Schwab, M., Amarsinghe, S., Zaharia, M., InfoLab, S.: Weld: A common runtime for high performance data analytics. In: *Conference on Innovative Data Systems Research (CIDR)* (2017)
21. Simitsis, A., Wilkinson, K., Castellanos, M., Dajal, U.: Optimizing analytic data flows for multiple execution engines. In: *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. pp. 829–840. ACM (2012)
22. Stonebraker, M., Cetintemel, U.: ”one size fits all”: an idea whose time has come and gone. In: *Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on*. pp. 2–11. IEEE (2005)
23. Sun, N., Morris, J., Xu, J., Zhu, X., Xiao, Y.: Biginsight: A framework for big data-based banking customer analytics. *IBM Journal of Research and Development* **58**(5/6), 4–1 (2014)
24. Valduriez, P.: Parallel database systems: open problems and new issues. *Distributed and parallel Databases* **1**(2), 137–165 (1993)
25. Xu, C., Chen, Y., Liu, Q., Ren, Y., Min, Y., Gu, G.: A unified computation engine for big data analytics. In: *Big Data Computing (BDC), 2015 IEEE/ACM 2nd International Symposium on*. pp. 73–77. IEEE (2015)
26. Yu, K., Gadepally, V., Stonebraker, M.: Database engine integration and performance analysis of the bigdata polystore system. In: *High Performance Extreme Computing Conference (HPEC)*, 2017 IEEE. pp. 1–7. IEEE (2017)
27. Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., Stoica, I.: Spark: Cluster computing with working set. *HotCloud* **10**(10-10), 95 (2010)