

15. A. Cuzzocrea and E. Bertino. Privacy preserving olap over distributed xml data: A theoretically-sound secure-multiparty-computation approach. *J. Comput. Syst. Sci.*, 77(6):965–987, 2011.
16. A. Cuzzocrea and S. Chakravarthy. Event-based lossy compression for effective and efficient olap over data streams. *Data Knowl. Eng.*, 69(7):678–708, 2010.
17. A. Cuzzocrea, R. Moussa, and G. Xu. Olap*: Effectively and efficiently supporting parallel olap over big data. In *MEDI*, pages 38–49, 2013.
18. A. Cuzzocrea, V. Russo, and D. Saccà. A robust sampling-based framework for privacy preserving olap. In *DaWaK*, pages 97–114, 2008.
19. A. Cuzzocrea and D. Saccà. Balancing accuracy and privacy of olap aggregations on data cubes. In *DOLAP*, pages 93–98, 2010.
20. A. Cuzzocrea, D. Saccà, and P. Serafino. A hierarchy-driven compression technique for advanced olap visualization of multidimensional data cubes. In *DaWaK*, pages 106–119, 2006.
21. A. Cuzzocrea, D. Saccà, and P. Serafino. Semantics-aware advanced olap visualization of multidimensional data cubes. *IJDWM*, 3(4):1–30, 2007.
22. A. Cuzzocrea, I.-Y. Song, and K. C. Davis. Analytics over large-scale multidimensional data: the big data revolution! In *DOLAP*, pages 101–104, 2011.
23. J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.
24. F. K. H. A. Dehne, T. Eavis, and A. Rau-Chaplin. The cgmcube project: Optimizing parallel data cube generation for rolap. *Distributed and Parallel Databases*, 19(1):29–62, 2006.
25. F. K. H. A. Dehne, Q. Kong, A. Rau-Chaplin, H. Zaboli, and R. Zhou. A distributed tree data structure for real-time olap on cloud architectures. In *BigData Conference*, pages 499–505, 2013.
26. A. G. Erdman, D. F. Keefe, and R. Schiestl. Grand challenge: Applying regulatory science and big data to improve medical device innovation. *IEEE Trans. Biomed. Engineering*, 60(3):700–706, 2013.
27. W. Fan and A. Bifet. Mining big data: current status, and forecast to the future. *SIGKDD Explorations*, 14(2):1–5, 2012.
28. N. Ferreira, J. Poco, H. T. Vo, J. Freire, and C. T. Silva. Visual exploration of big spatio-temporal urban data: A study of new york city taxi trips. *IEEE Trans. Vis. Comput. Graph.*, 19(12):2149–2158, 2013.
29. J. Gray, S. Chaudhuri, A. Bosworth, A. Layman, D. Reichart, M. Venkatrao, F. Pellow, and H. Pirahesh. Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub totals. *Data Min. Knowl. Discov.*, 1(1):29–53, 1997.
30. H. Herodotou, H. Lim, G. Luo, N. Borisov, L. Dong, F. B. Cetin, and S. Babu. Starfish: A self-tuning system for big data analytics. In *CIDR*, pages 261–272, 2011.
31. W. Hummer, B. Satzger, and S. Dustdar. Elastic stream processing in the cloud. *Wiley Interdisc. Rev.: Data Mining and Knowledge Discovery*, 3(5):333–345, 2013.
32. D. Jiang, B. C. Ooi, L. Shi, and S. Wu. The performance of mapreduce: An in-depth study. *PVLDB*, 3(1):472–483, 2010.
33. K. Kambatla, G. Kollias, V. Kumar, and A. Grama. Trends in big data analytics. *J. Parallel Distrib. Comput.*, 74(7):2561–2573, 2014.
34. U. Kang, L. Akoglu, and D. H. Chau. Big graph mining for the web and social media: algorithms, anomaly detection, and applications. In *WSDM*, pages 677–678, 2014.
35. U. Kang and C. Faloutsos. Big graph mining: algorithms and discoveries. *SIGKDD Explorations*, 14(2):29–36, 2012.
36. S. Khouri and L. Bellatreche. Dwobs: Data warehouse design from ontology-based sources. In *DASFAA (2)*, pages 438–441, 2011.

37. S. Khouri, L. Bellatreche, and N. Berkani. Modetl: A complete modeling and etl method for designing data warehouses from semantic databases. In *COMAD*, page 113, 2012.
38. H.-C. Kum, A. Krishnamurthy, A. Machanavajjhala, and S. C. Ahalt. Social genome: Putting big data to work for population informatics. *IEEE Computer*, 47(1):56–63, 2014.
39. D. Laney. 3D data management: Controlling data volume, velocity, and variety. Technical report, META Group, February 2001.
40. J. Lin and D. V. Ryaboy. Scaling big data mining infrastructure: the twitter experience. *SIGKDD Explorations*, 14(2):6–19, 2012.
41. Z. Lin, D. H. P. Chau, and U. Kang. Leveraging memory mapping for fast and scalable graph computation on a pc. In *BigData Conference*, pages 95–98, 2013.
42. F. J. Meng, X. Zhuo, B. Yang, J. M. Xu, P. Jin, A. Apte, and J. Wigglesworth. A generic framework for application configuration discovery with pluggable knowledge. In *IEEE CLOUD*, pages 236–243, 2013.
43. A. O’Driscoll, J. Dugelaite, and R. D. Sleator. ‘big data’, hadoop and cloud computing in genomics. *Journal of Biomedical Informatics*, 46(5):774–781, 2013.
44. M. Paoletti, G. Camiciottoli, E. Meoni, F. Bigazzi, L. Cestelli, M. Pistolesi, and C. Marchesi. Explorative data analysis techniques and unsupervised clustering methods to support clinical assessment of chronic obstructive pulmonary disease (copd) phenotypes. *Journal of Biomedical Informatics*, 42(6):1013–1021, 2009.
45. J. Pei. Some new progress in analyzing and mining uncertain and probabilistic data for big data analytics. In *RSFDGrC*, pages 38–45, 2013.
46. D. J. Power. Using ‘big data’ for analytics and decision support. *Journal of Decision Systems*, 23(2):222–228, 2014.
47. S. Sarawagi, R. Agrawal, and N. Megiddo. Discovery-driven exploration of olap data cubes. In *EDBT*, pages 168–182, 1998.
48. E. A. Sitaridi and K. A. Ross. Ameliorating memory contention of olap operators on gpu processors. In *DaMoN*, pages 39–47, 2012.
49. Y. Sun and J. Han. Mining heterogeneous information networks: a structural analysis approach. *SIGKDD Explorations*, 14(2):20–28, 2012.
50. A. Thusoo, J. S. Sarma, N. Jain, Z. Shao, P. Chakka, N. Zhang, S. Anthony, H. Liu, and R. Murthy. Hive - a petabyte scale data warehouse using hadoop. In *ICDE*, pages 996–1005, 2010.
51. A. N. Toosi, R. N. Calheiros, and R. Buyya. Interconnected cloud computing environments: Challenges, taxonomy, and survey. *ACM Comput. Surv.*, 47(1):7, 2014.
52. M. Weidner, J. Dees, and P. Sanders. Fast olap query execution in main memory on large data in a cluster. In *BigData Conference*, pages 518–524, 2013.
53. Y. Yuan, X. Lin, Q. Liu, W. Wang, J. X. Yu, and Q. Zhang. Efficient computation of the skyline cube. In *VLDB*, pages 241–252, 2005.
54. X. Zhang, C. Liu, S. Nepal, C. Yang, W. Dou, and J. Chen. Sac-frapp: a scalable and cost-effective framework for privacy preservation over big data on cloud. *Concurrency and Computation: Practice and Experience*, 25(18):2561–2576, 2013.