

References

- [1] Apache Jena. <http://jena.apache.org/>.
- [2] DBLP. <http://www.informatik.uni-trier.de/~ley/db/>.
- [3] Hadoop Distributed File System. <http://hadoop.apache.org/>.
- [4] Linked Open Data. <http://www.w3.org/wiki/SweoIG/TaskForces/CommunityProjects/LinkingOpenData>.
- [5] T. Akiba, Y. Iwata, and Y. Yoshida. Fast exact shortest-path distance queries on large networks by pruned landmark labeling. In *SIGMOD'13*, pages 349–360, 2013.
- [6] M. Atre, V. Chaoji, M. J. Zaki, and J. A. Hendler. Matrix "bit" loaded: a scalable lightweight join query processor for rdf data. In *WWW'10*, pages 41–50, 2010.
- [7] J. Cheng, X. Zeng, and J. X. Yu. Top-k graph pattern matching over large graphs. In *ICDE'13*, pages 1033–1044, 2013.
- [8] R. Fagin, A. Lotem, and M. Naor. Optimal aggregation algorithms for middleware. *J. Comput. Syst. Sci.*, pages 614–656, 2003.
- [9] W. Fan, J. Li, S. Ma, N. Tang, Y. Wu, and Y. Wu. Graph pattern matching: From intractable to polynomial time. *PVLDB*, pages 264–275, 2010.
- [10] W. Fan, X. Wang, and Y. Wu. Diversified top-k graph pattern matching. *PVLDB*, pages 1510–1521, 2013.
- [11] Y. Fang, K. C.-C. Chang, and H. W. Lauw. Roundtriprank: Graph-based proximity with importance and specificity. In *ICDE'13*, pages 613–624, 2013.
- [12] Google. "freebase data dumps". <https://developers.google.com/freebase/data>.
- [13] M. Gupta, J. Gao, X. Yan, H. Cam, and J. Han. Top-k interesting subgraph discovery in information networks. In *ICDE'14*, pages 820–831, 2014.
- [14] J. Huang, D. J. Abadi, and K. Ren. Scalable sparql querying of large rdf graphs. *PVLDB*, pages 1123–1134, 2011.
- [15] I. F. Ilyas, G. Beskales, and M. A. Soliman. A survey of top-k query processing techniques in relational database systems. *ACM Comput. Surv.*, 2008.
- [16] J. Jin, S. Khemmarat, L. Gao, and J. Luo. A distributed approach for top-k star queries on massive information networks. In *ICPADS'14*, pages 9–16, 2014.
- [17] J. Jin, S. Khemmarat, L. Gao, and J. Luo. Querying web-scale information networks through bounding matching scores. Technical report, 2015. <http://rio.ecs.umass.edu/mnilpub/papers/tech2015-jin.pdf>.
- [18] A. Khan, N. Li, X. Yan, Z. Guan, S. Chakraborty, and S. Tao. Neighborhood based fast graph search in large networks. In *SIGMOD'11*, pages 901–912, 2011.
- [19] A. Khan, Y. Wu, C. C. Aggarwal, and X. Yan. Nema: Fast graph search with label similarity. *PVLDB*, pages 181–192, 2013.
- [20] S. Khemmarat and L. Gao. Fast top-k path-based relevance query on massive graphs. In *ICDE'14*, pages 316–327, 2014.
- [21] Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. M. Hellerstein. Distributed graphlab: A framework for machine learning in the cloud. *PVLDB*, pages 716–727, 2012.
- [22] G. Malewicz, M. H. Austern, A. J. C. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: a system for large-scale graph processing. In *SIGMOD'10*, pages 135–146, 2010.
- [23] T. Neumann and G. Weikum. Rdf-3x: a risc-style engine for rdf. *PVLDB*, pages 647–659, 2008.
- [24] R. Power and J. Li. Piccolo: Building fast, distributed programs with partitioned tables. In *OSDI'10*, pages 293–306, 2010.
- [25] P. Ravindra and K. Anyanwu. Scalable processing of flexible graph pattern queries on the cloud. In *WWW'13 (Companion Volume)*, pages 169–170, 2013.
- [26] D. Shasha, J. T.-L. Wang, and R. Giugno. Algorithmics and applications of tree and graph searching. In *PODS'02*, pages 39–52, 2002.
- [27] Z. Sun, H. Wang, H. Wang, B. Shao, and J. Li. Efficient subgraph matching on billion node graphs. *PVLDB*, pages 788–799, 2012.
- [28] Y. Tian and J. M. Patel. Tale: A tool for approximate large graph matching. In *ICDE'08*, pages 963–972, 2008.
- [29] R. Verborgh, M. V. Sande, P. Colpaert, S. Coppens, E. Mannens, and R. V. de Walle. Web-scale querying through linked data fragments. In *LDOW'14*, 2014.
- [30] X. Yan, F. Zhu, P. S. Yu, and J. Han. Feature-based similarity search in graph structures. *ACM Trans. Database Syst.*, pages 1418–1453, 2006.
- [31] K. Zeng, J. Yang, H. Wang, B. Shao, and Z. Wang. A distributed graph engine for web scale rdf data. *PVLDB*, pages 265–276, 2013.
- [32] Y. Zhang, Q. Gao, L. Gao, and C. Wang. Priter: A distributed framework for prioritizing iterative computations. *IEEE Trans. Parallel Distrib. Syst.*, pages 1884–1893, 2013.