References

[1] Aapo Kyrola, Guy Blelloch, Carlos Guestrin. (2012) 'GraphChi: Large-Scale Graph Computation on Just a PC'.

[2] Abdul Quamar, Amol Deshpande, and Jimmy Lin. NScale: neighborhood-centric analytics on large graphs. Proceedings of the VLDB Endowment, 7(13):1673–1676, 2014.

[3] Abhishek Jindal and Steve Madden. GRAPHiQL: A graph intuitive query language for relational databases. In Big Data (Big Data), 2014 IEEE International Conference on, pages 441–450. IEEE, 2014.

[4] Adam Welc, Raghavan Raman, Zhe Wu, Sungpack Hong, Hassan Chafi, and Jay Banerjee. Graph analysis: do we have to reinvent the wheel? In First International Workshop on Graph Data Management Experiences and Systems, page 7. ACM, 2013.

[5] Alekh Jindal, Samuel Madden, Malu Castellanos, and Meichun Hsu. Graph Analytics using the Vertica Relational Database. arXiv preprint arXiv:1412.5263, 2014.

[6] AllegroGraph RDFStore Web 3.0's Database. http : / / franz . com / agraph / allegrograph/.

[7] Amitabha Roy, Ivo Mihailovic, Willy Zwaenepoel. (2013) 'X-Stream: Edge-centric Graph Processing using Streaming Partitions'.

[8] Apache Jena - Home. https://jena.apache.org.

[9] A Comprehensive List of Big Data Statistics. [Online]. Wikibon blog < http://wikibon.org/blog/big-data-statistics/ > Accessed 20 Sept 2015.

[10] Bin Shao, Haixun Wang, and Yatao Li. Trinity: A distributed graph engine on a memory cloud. In Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, pages 505–516. ACM, 2013.

[11] Chen.R, Weng.X, He.B, and Yang.M (2010) 'Large graph processing in the cloud', pp 1123–1126, Indianapolis, Indiana, USA,. ACM

[12] Dean.J and Ghemawat.S. (2004) 'Mapreduce: Simplified data processing on large clusters. In Proceedings of the 6th conference on Symposium on Opearting Systems Design & Implementation', Vol 6, pp 10–10.

[13] Eiko Yoneki, Amitabha Roy. (2013) 'Scale-up Graph Processing: A Storage-centric View'.

[14] Emil Eifrem. The NewWay to Access Super Fast Social Data. http://mashable.com/ 2012/09/26/graph-databases/.

[15] Facebook, 2014. http://facebook.com.

[16] Giraph - Welcome To Apache Giraph! http://giraph.apache.org.

[17] Gonzalez.J.E, Low.Y, Gu.H, Bickson.D, and Guestrin.C. (2012) 'Powergraph: distributed graph-parallel computation on natural graphs', pp 17–30.

[18] Google, 2014. http://google.com.

[19] Graph-tool: Efficient graph analysis. http://graph-tool.skewed.de.

[20] Grzegorz Malewicz, Matthew H Austern, Aart JC Bik, James C Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: a system for large-scale graph processing. In Proceedings of the 2010 ACM SIGMOD International Conference on Management of data, pages 135–146. ACM, 2010.

[21] Guozhang Wang, Wenlei Xie, Alan J Demers, and Johannes Gehrke. Asynchronous Large-Scale Graph Processing Made Easy. In CIDR, 2013.

[22] Ian Robinson, JimWebber, and Emil Eifrem. Graph databases. " O'Reilly Media, Inc.", 2013.

[23] Jiefeng Cheng, Qin Liu, Zhenguo Li, Wei Fan, John C.S. Lui, Cheng He (2015) 'VENUS: Vertex-Centric Streamlined Graph Computation on a Single PC'.IEEE.

[24] Jiewen Huang, Kartik Venkatraman, and Daniel J Abadi. Query optimization of distributed pattern matching. In Data Engineering (ICDE), 2014 IEEE 30th International Conference on, pages 64–75. IEEE, 2014.

[25] Jing Fan, Adalbert Gerald, Soosai Raj, and Jignesh M Patel. The case against specialized graph analytics engines. 2015.

[26] Kamran Najeebullah, Kifayat Ullah Khan, Waqas Nawaz and Young-Koo Lee. (2014) 'BiShard Parallel Processor: A Disk-Based Processing Engine for Billion-Scale Graphs'.

[27] Kyrola.A, Blelloch.G, and Guestrin.C. (2012) 'GraphChi: large-scale graph computation on just a PC', pp. 31–46.

[28] Lassila Ora and Swick Ralph. Resource Description Framework (RDF) Model and Syntax Specification. http : / / www. w3 . org / TR / 1999 / REC-rdf-syntax-19990222/

[29] Leskovec.J, Lang.K, Dasgupta.A, and Mahoney.M. (2009) 'Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. Internet Mathematics' pp 29–123.

[30] Low.Y, Gonzalez.J, Kyrola.A, Bickson.D, and Guestrin.C. (2012) 'Graphlab: A distributed framework for machine learning in the cloud'.

[31] Mahesh Lal (2015) 'Neo4j Graph Data Modeling'. Packt Publishing

[32] Malewicz.G, Austern.M. H, Bik.A. J, Dehnert.J, Horn.I, Leiser.N, and Czajkowski.G. (2010) 'Pregel: a system for large-scale graph processing'.

[33] Min-Soo Kim, Sangyeon Lee, Wook-Shin Han, Himchan Park, and Jeong-Hoon Lee (2015) 'DSP-CC-: I/O Efficient Parallel Computation of Connected Components in Billion-Scale Networks'.IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING.

[34] Neo4j, World's Leading Graph Database. [Online] < http://neo4j.com > Accessed 21 Dec 2015.

[35] Neo4j and Apache Spark - Neo4j Graph Database. http://neo4j.com/developer/apache-spark/#mazerunner.

[36] OrientDB - OrientDB Multi-Model NoSQL DatabaseOrientDB Multi-Model NoSQL Database. http://orientdb.com/orientdb/.

[37] Overview - NetworkX. http://networkx.github.io.

[38] Page.L, Brin.S, Motwani.R, and Winograd.T. (1999) 'The pagerank citation ranking: Bringing order to the web', Stanford InfoLab,

[39] Paul Erdos and Alfred Renyi. On the strength of connectedness of a random graph. Acta Mathematica Hungarica, 12(1-2):261–267, 1961.

[40] Pingpeng Yuan, Wenya Zhang, Changfeng Xie, Hai Jin, Ling Liu, Kisung Lee (2015) 'Fast Iterative Graph Computation: A Path Centric Approach'. International Conference for High Performance Computing, Networking, Storage and Analysis.IEEE.

[41] Power.R and Li. Piccolo.J (2010) 'Building fast, distributed programs with partitioned tables', pp 1–14.

[42] Santo Fortunato. Community detection in graphs. Physics Reports, 486(3):75–174, 2010.

[43] Raja Appuswamy, Christos Gkantsidis, Dushyanth Narayanan,Orion Hodson, and Antony Rowstron (2013) 'Scale-up vs Scale-out for Hadoop: Time to rethink?'. Cambridge University

[44] Raymond Cheng, Ji Hong, Aapo Kyrola, Youshan Miao, Xuetian Weng, Ming Wu, Fan Yang, Lidong Zhou, Feng Zhao, and Enhong Chen. Kineograph: taking the pulse of a fast-changing and connected world. In Proceedings of the 7th ACM european conference on Computer Systems, pages 85–98. ACM, 2012.

[45] Serge Abiteboul, Richard Hull, and Victor Vianu. Foundations of databases, volume 8. Addison-Wesley Reading, 1995

[46] Sherif Sakr, Sameh Elnikety, and Yuxiong He. G-SPARQL: a hybrid engine for querying large attributed graphs. In Proceedings of the 21st ACM international conference on Information and knowledge management, pages 335–344. ACM, 2012.

[47] Stanford Graph Analysis Project. http://snap.stanford.edu.

[48] Stardog: Enterprise Graph Database. http://stardog.com.

[49] Talbot.J, Yoo.R. M., and Kozyrakis.C. (2011) ‘Phoenix++: Modular MapReduce for Shared-Memory Systems’.

[50] Titan: Distributed Graph Database. http://thinkaurelius.github.io/titan/.

[51] Ulrik Brandes and Thomas Erlebach. Graph analysis: methodological foundations, volume 3418. Springer Science & Business Media, 2005.

[52] Wook-Shin Han, Sangyeon Lee, Kyungyeol Park, Jeong-Hoon Lee,Min-Soo Kim1, Jinha Kim, Hwanjo Yu. (2013) 'TurboGraph: A Fast Parallel Graph Engine Handling Billion-scale Graphs in a Single PC'.

[53] Xenarios.I, Salwinski.L, Duan.X. J., Higney.P, Kim.S.M., and Eisenberg.D. (2002) ‘The database of interacting proteins: a research tool for studying cellular networks of protein interactions. Nucleic acids research’, pp 303–305.

[54] Yan.X, Yu.P. S, and Han.J. (2004) ‘Graph indexing: A frequent structure-based approach’, pp 335–346.

[55] Yifang Jiang, Kai Chen, Yi Zhou, Diao Zhang, Qu Zhou, Jianhua He (2014) 'An Improved Memory Management Scheme for Large Scale Graph Computing Engine GraphChi'. 2014 IEEE International Conference on Big Data

[56] Yuanyuan Tian, Andrey Balmin, Severin Andreas Corsten, Shirish Tatikonda, and John McPherson. From think like a vertex to think like a graph. Proceedings of the VLDB Endowment, 7(3):193–204, 2013.

[57] Yucheng Low, Joseph E Gonzalez, Aapo Kyrola, Danny Bickson, Carlos E Guestrin, and Joseph Hellerstein. Graphlab: A new framework for parallel machine learning. arXiv preprint arXiv:1408.2041, 2014.

[58] Yuhanna Noel, Owens Leslie, and Elizabeth Cullen. Market Overview: Graph Databases. https://www.forrester.com/Market+Overview+Graph+Databases/ fulltext/-/E-res121473.

[59] Zaharia.M, Chowdhury.M, Franklin.M. J., Shenker.S, and Stoica.I. (2010) ‘Spark: Cluster computing with working sets. In HotCloud’.

[60] Zhao.P and Han.J. (2010) ‘On graph query optimization in large networks.PVLDB’, pp 340–351.

[61] Zhiyuan Lin, Minsuk Kahng, Kaeser Md. Sabrin, Duen Horng (Polo) Chau (2015) 'MMap: Fast Billion-Scale Graph Computation on a PC via Memory Mapping'. IEEE.