

Bibliography

- [1] Intrusion as (Anti) social Communication : Characterization and Detection. *Kdd '12*, pages 886–894, 2012.
- [2] Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, 2015.
- [3] Alvarez Abdul-Rahman. The pgp trust model. In *EDI-Forum: the Journal of Electronic Commerce*, volume 10, pages 27–31, 1997.
- [4] Charu C Aggarwal. *Data streams: models and algorithms*, volume 31. Springer Science & Business Media, 2007.
- [5] Charu C Aggarwal, Jiawei Han, Jianyong Wang, and Philip S Yu. A framework for clustering evolving data streams. In *Proceedings of the 29th international conference on Very large data bases-Volume 29*, pages 81–92. VLDB Endowment, 2003.
- [6] Tero Aittokallio and Benno Schwikowski. Graph-based methods for analysing networks in cell biology. *Briefings in bioinformatics*, 7(3):243–255, 2006.
- [7] István Albert and Réka Albert. Conserved network motifs allow protein–protein interaction prediction. *Bioinformatics*, 20(18):3346–3352, 2004.
- [8] Réka Albert, István Albert, and Gary L Nakarado. Structural vulnerability of the north american power grid. *Physical review E*, 69(2):025103, 2004.
- [9] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *nature*, 406(6794):378–382, 2000.
- [10] Berni J Alder and TE Wainwright. Studies in molecular dynamics. i. general method. *The Journal of Chemical Physics*, 31(2):459–466, 1959.
- [11] Noga Alon, Phuong Dao, Iman Hajirasouliha, Fereydoun Hormozdiari, and S Cenk Sahinalp. Biomolecular network motif counting and discovery by color coding. *Bioinformatics*, 24(13):i241–i249, 2008.
- [12] Luis A Nunes Amaral, Antonio Scala, Marc Barthelemy, and H Eugene Stanley. Classes of small-world networks. *Proceedings of the National Academy of Sciences*, 97(21):11149–11152, 2000.
- [13] Reid Andersen, Fan Chung, and Kevin Lang. Local graph partitioning using pagerank vectors. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 475–486. IEEE, 2006.
- [14] James P Anderson. Computer security threat monitoring and surveillance. Technical report, Technical report, James P. Anderson Company, Fort Washington, Pennsylvania, 1980.
- [15] Renzo Angles. A comparison of current graph database models. In *Data Engineering Workshops (ICDEW), 2012 IEEE 28th International Conference on*, pages 171–177. IEEE, 2012.
- [16] Renzo Angles and Claudio Gutierrez. Survey of graph database models. *ACM Computing Surveys (CSUR)*, 40(1):1, 2008.
- [17] AR Atilgan, SR Durell, RL Jernigan, MC Demirel, O Keskin, and I Bahar. Anisotropy of fluctuation dynamics of proteins with an elastic network model. *Biophysical journal*, 80(1):505–515, 2001.
- [18] Ching Avery. Giraph: Large-scale graph processing infrastructure on hadoop. *Proceedings of the Hadoop Summit. Santa Clara*, 11, 2011.

- [19] Brian Babcock, Shivnath Babu, Mayur Datar, Rajeev Motwani, and Jennifer Widom. Models and issues in data stream systems. In *Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 1–16. ACM, 2002.
- [20] David A Bader, Jonathan Berry, Adam Amos-Binks, Daniel Chavarría-Miranda, Charles Hastings, Kamesh Madduri, and Steven C Poulos. Stinger: Spatio-temporal interaction networks and graphs (sting) extensible representation. *Georgia Institute of Technology, Tech. Rep*, 2009.
- [21] David A Bader, Shiva Kintali, Kamesh Madduri, and Milena Mihail. Approximating betweenness centrality. In *International Workshop on Algorithms and Models for the Web-Graph*, pages 124–137. Springer, 2007.
- [22] David A Bader and Kamesh Madduri. Snap, small-world network analysis and partitioning: an open-source parallel graph framework for the exploration of large-scale networks. In *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pages 1–12. IEEE, 2008.
- [23] Guillaume Bagan, Angela Bonifati, Radu Ciucanu, George Fletcher, Aurélien Lemay, and Nicky Advokaat. Generating flexible workloads for graph databases. *Proceedings of the VLDB Endowment (PVLDB)*, 2016.
- [24] James P Bagrow and Erik M Bollt. Local method for detecting communities. *Physical Review E*, 72(4):046108, 2005.
- [25] Alexandru T Balaban. Applications of graph theory in chemistry. *Journal of Chemical Information and Computer Sciences*, 25(3):334–343, 1985.
- [26] Gareth Ball, Paul Aljabar, Sally Zebari, Nora Tusor, Tomoki Arichi, Nazakat Merchant, Emma C Robinson, Enitan Ogundipe, Daniel Rueckert, A David Edwards, et al. Rich-club organization of the newborn human brain. *Proceedings of the National Academy of Sciences*, 111(20):7456–7461, 2014.
- [27] Ziv Bar-Yossef, Ravi Kumar, and D Sivakumar. Reductions in streaming algorithms, with an application to counting triangles in graphs. In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 623–632. Society for Industrial and Applied Mathematics, 2002.
- [28] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [29] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [30] Albert-Laszlo Barabási, Hawoong Jeong, Zoltan Néda, Erzsebet Ravasz, Andras Schubert, and Tamas Vicsek. Evolution of the social network of scientific collaborations. *Physica A: Statistical mechanics and its applications*, 311(3):590–614, 2002.
- [31] Adrian Barbu and S-C Zhu. Graph partition by swendsen-wang cuts. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 320–327. IEEE, 2003.
- [32] Alain Barrat, Marc Barthelemy, Romualdo Pastor-Satorras, and Alessandro Vespignani. The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(11):3747–3752, 2004.
- [33] Marc Barthelemy. Betweenness centrality in large complex networks. *The European Physical Journal B-Condensed Matter and Complex Systems*, 38(2):163–168, 2004.
- [34] Mathieu Bastian, Sebastien Heymann, Mathieu Jacomy, et al. Gephi: an open source software for exploring and manipulating networks. *ICWSM*, 8:361–362, 2009.
- [35] Vladimir Batagelj and Andrej Mrvar. Pajek-program for large network analysis. *Connections*, 21(2):47–57, 1998.
- [36] Vladimir Batagelj and Andrej Mrvar. Pajekanalysis and visualization of large networks. In *Graph drawing software*, pages 77–103. Springer, 2004.
- [37] Luca Becchetti, Paolo Boldi, Carlos Castillo, and Aristides Gionis. Efficient semi-streaming algorithms for local triangle counting in massive graphs. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 16–24. ACM, 2008.

- [38] Wayne J Becktel and John A Schellman. Protein stability curves. *Biopolymers*, 26(11):1859–1877, 1987.
- [39] Richard Bellman. On a routing problem. *Quarterly of applied mathematics*, pages 87–90, 1958.
- [40] Herman JC Berendsen, JPM van Postma, Wilfred F van Gunsteren, ARHJ DiNola, and JR Haak. Molecular dynamics with coupling to an external bath. *The Journal of chemical physics*, 81(8):3684–3690, 1984.
- [41] Elisabetta Bergamini, Henning Meyerhenke, and Christian L Staudt. Approximating betweenness centrality in large evolving networks. *arXiv preprint arXiv:1409.6241*, 2014.
- [42] Tanya Y Berger-Wolf and Jared Saia. A framework for analysis of dynamic social networks. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 523–528. ACM, 2006.
- [43] Betim Berjani and Thorsten Strufe. A recommendation system for spots in location-based online social networks. In *Proceedings of the 4th Workshop on Social Network Systems*, page 4. ACM, 2011.
- [44] Aaron Bernstein and Liam Roditty. Improved dynamic algorithms for maintaining approximate shortest paths under deletions. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 1355–1365. SIAM, 2011.
- [45] Chris Biemann, Lachezar Krumov, Stefanie Roos, and Karsten Weihe. Network motifs are a powerful tool for semantic distinction. In *Towards a Theoretical Framework for Analyzing Complex Linguistic Networks*, pages 83–105. Springer, 2016.
- [46] Chris Biemann, Stefanie Roos, and Karsten Weihe. Quantifying semantics using complex network analysis. In *In COLING12*. Citeseer, 2012.
- [47] Zsolt Bikadi and Eszter Hazai. Application of the pm6 semi-empirical method to modeling proteins enhances docking accuracy of autodock. *Journal of cheminformatics*, 1(1):1, 2009.
- [48] RH Bisseling and WF McColl. Scientific computing on bulk synchronous parallel architectures. 1994.
- [49] Daniel K Blandford, Guy E Blelloch, and Ian A Kash. Compact representations of separable graphs. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 679–688. Society for Industrial and Applied Mathematics, 2003.
- [50] Daniel K Blandford, Guy E Blelloch, and Ian A Kash. An experimental analysis of a compact graph representation. 2004.
- [51] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008.
- [52] Stefano Boccaletti, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. Complex networks: Structure and dynamics. *Physics reports*, 424(4):175–308, 2006.
- [53] Phillip Bonacich. Some unique properties of eigenvector centrality. *Social networks*, 29(4):555–564, 2007.
- [54] Stephen P Borgatti, Ajay Mehra, Daniel J Brass, and Giuseppe Labianca. Network analysis in the social sciences. *science*, 323(5916):892–895, 2009.
- [55] Dhruba Borthakur. The hadoop distributed file system: Architecture and design. *Hadoop Project Website*, 11(2007):21, 2007.
- [56] Ulrik Brandes. A faster algorithm for betweenness centrality*. *Journal of mathematical sociology*, 25(2):163–177, 2001.
- [57] Tom Britton, Maria Deijfen, and Anders Martin-Löf. Generating simple random graphs with prescribed degree distribution. *Journal of Statistical Physics*, 124(6):1377–1397, 2006.
- [58] Yingyi Bu. Pregelix: dataflow-based big graph analytics. In *Proceedings of the 4th annual Symposium on Cloud Computing*, page 54. ACM, 2013.
- [59] Buriol et al. Counting triangles in data streams. In *ACM SIGMOD-SIGACT*, 2006.

- [60] D Ewen Cameron, Caleb J Bashor, and James J Collins. A brief history of synthetic biology. *Nature Reviews Microbiology*, 12(5):381–390, 2014.
- [61] Germano Caronni. Walking the web of trust. In *Enabling Technologies: Infrastructure for Collaborative Enterprises, 2000.(WET ICE 2000). Proceedings. IEEE 9th International Workshops on*, pages 153–158. IEEE, 2000.
- [62] Samarjit Chakraborty and Somenath Biswas. Approximation algorithms for 3-d common substructure identification in drug and protein molecules. In *Workshop on Algorithms and Data Structures*, pages 253–264. Springer, 1999.
- [63] Timothy M Chan. All-pairs shortest paths with real weights in $o(n^3/\log n)$ time. In *Workshop on Algorithms and Data Structures*, pages 318–324. Springer, 2005.
- [64] Jin Chen, Wynne Hsu, Mong Li Lee, and See-Kiong Ng. Nemofinder: Dissecting genome-wide protein-protein interactions with meso-scale network motifs. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 106–115. ACM, 2006.
- [65] 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.
- [66] Avery Ching. Scaling apache giraph to a trillion edges. *Facebook Engineering blog*, page 25, 2013.
- [67] Avery Ching. Giraph: Production-grade graph processing infrastructure for trillion edge graphs. *ATPESC, ser. ATPESC*, 14, 2014.
- [68] Marek Ciglan, Alex Averbuch, and Ladiav Hluchy. Benchmarking traversal operations over graph databases. In *Data Engineering Workshops (ICDEW), 2012 IEEE 28th International Conference on*, pages 186–189. IEEE, 2012.
- [69] Benoit Claise. Cisco Systems NetFlow Services Export Version 9. RFC 3954, October 2015.
- [70] G Marius Clore, Ettore Appella, Masaki Yamada, Kouji Matsushima, and Angela M Gronenborn. Three-dimensional structure of interleukin 8 in solution. *Biochemistry*, 29(7):1689–1696, 1990.
- [71] Edgar F Codd. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387, 1970.
- [72] Recep Colak, Fereydoun Hormozdiari, Flavia Moser, Alexander Schönhuth, J Holman, Martin Ester, and Süleyman Cenk Sahinalp. Dense graphlet statistics of protein interaction and random networks. In *Pacific Symposium on Biocomputing*, volume 548, pages 178–189. Citeseer, 2009.
- [73] Vittoria Colizza, Alessandro Flammini, M Angeles Serrano, and Alessandro Vespignani. Detecting rich-club ordering in complex networks. *Nature physics*, 2(2):110–115, 2006.
- [74] Guusje Collin, Olaf Sporns, René CW Mandl, and Martijn P van den Heuvel. Structural and functional aspects relating to cost and benefit of rich club organization in the human cerebral cortex. *Cerebral cortex*, 24(9):2258–2267, 2014.
- [75] Michael Collins. *A protocol graph based anomaly detection system*. ProQuest, 2008.
- [76] Thomas H Cormen. *Introduction to algorithms*. MIT press, 2009.
- [77] Evangelos A Coutsiias, Chaok Seok, and Ken A Dill. Using quaternions to calculate rmsd. *Journal of computational chemistry*, 25(15):1849–1857, 2004.
- [78] Ciprian G Crismaru, Gjalt G Wybenga, Wiktor Szymanski, Hein J Wijma, Bian Wu, Sebastian Bartsch, Stefaan de Wildeman, Gerrit J Poelarends, Ben L Feringa, Bauke W Dijkstra, et al. Biochemical properties and crystal structure of a β -phenylalanine aminotransferase from variovorax paradoxus. *Applied and environmental microbiology*, 79(1):185–195, 2013.
- [79] Paolo Crucitti, Vito Latora, and Massimo Marchiori. A topological analysis of the italian electric power grid. *Physica A: Statistical mechanics and its applications*, 338(1):92–97, 2004.
- [80] Paolo Crucitti, Vito Latora, Massimo Marchiori, and Andrea Rapisarda. Error and attack tolerance of complex networks. *Physica A: Statistical Mechanics and its Applications*, 340(1):388–394, 2004.

- [81] Leonardo Dagum and Ramesh Menon. Openmp: an industry standard api for shared-memory programming. *IEEE computational science and engineering*, 5(1):46–55, 1998.
- [82] Xavier Daura, Bernhard Jaun, Dieter Seebach, Wilfred F Van Gunsteren, and Alan E Mark. Reversible peptide folding in solution by molecular dynamics simulation. *Journal of molecular biology*, 280(5):925–932, 1998.
- [83] Wouter De Nooy, Andrej Mrvar, and Vladimir Batagelj. *Exploratory social network analysis with Pajek*, volume 27. Cambridge University Press, 2011.
- [84] Mattias De Wael, Stefan Marr, Joeri De Koster, Jennifer B Sartor, and Wolfgang De Meuter. Just-in-time data structures. In *2015 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!)*, pages 61–75. ACM, 2015.
- [85] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- [86] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: a flexible data processing tool. *Communications of the ACM*, 53(1):72–77, 2010.
- [87] Eliezer Dekel, David Nassimi, and Sartaj Sahni. Parallel matrix and graph algorithms. *SIAM Journal on computing*, 10(4):657–675, 1981.
- [88] Camil Demetrescu, Irene Finocchi, and Giuseppe F Italiano. Dynamic graphs. *Handbook on Data Structures and Applications*. Boca Raton, FL, USA: Chapman & Hall/CRC, 2005.
- [89] Camil Demetrescu and Giuseppe F Italiano. Fully dynamic transitive closure: Breaking through the $O(n^2)$ barrier. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 381–389. IEEE, 2000.
- [90] Camil Demetrescu and Giuseppe F Italiano. A new approach to dynamic all pairs shortest paths. *Journal of the ACM (JACM)*, 51(6):968–992, 2004.
- [91] Dorothy E Denning. An intrusion-detection model. *Software Engineering, IEEE Transactions on*, (2):222–232, 1987.
- [92] Edsger W Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.
- [93] Edsger Wybe Dijkstra, Edsger Wybe Dijkstra, Edsger Wybe Dijkstra, Etats-Unis Informaticien, and Edsger Wybe Dijkstra. *A discipline of programming*, volume 1. prentice-hall Englewood Cliffs, 1976.
- [94] Chris Ding, Xiaofeng He, Hongyuan Zha, Ming Gu, and Horst Simon. Spectral min-max cut for graph partitioning and data clustering. 2001.
- [95] Chris HQ Ding, Xiaofeng He, Hongyuan Zha, Ming Gu, and Horst D Simon. A min-max cut algorithm for graph partitioning and data clustering. In *Data Mining, 2001. ICDM 2001, Proceedings IEEE International Conference on*, pages 107–114. IEEE, 2001.
- [96] Mikhail Dmitriev. Profiling java applications using code hotswapping and dynamic call graph revelation. In *ACM SIGSOFT Software Engineering Notes*, volume 29, pages 139–150. ACM, 2004.
- [97] David Dominguez-Sal, Norbert Martinez-Bazan, Victor Muntés-Mulero, Pere Baleta, and Josep Lluís Larriba-Pey. A discussion on the design of graph database benchmarks. In *Technology Conference on Performance Evaluation and Benchmarking*, pages 25–40. Springer, 2010.
- [98] David Dominguez-Sal, P Urbón-Bayes, Aleix Giménez-Vanó, Sergio Gómez-Villamor, Norbert Martínez-Bazan, and Josep-Lluís Larriba-Pey. Survey of graph database performance on the hpc scalable graph analysis benchmark. In *International Conference on Web-Age Information Management*, pages 37–48. Springer, 2010.
- [99] SN Dorogovtsev, JFF Mendes, and AN Samukhin. Size-dependent degree distribution of a scale-free growing network. *Physical Review E*, 63(6):062101, 2001.
- [100] Antoine Dutot, Frédéric Guinand, Damien Olivier, and Yoann Pigné. Graphstream: A tool for bridging the gap between complex systems and dynamic graphs. In *Emergent Properties in Natural and Artificial Complex Systems. Satellite Conference within the 4th European Conference on Complex Systems (ECCS'2007)*, 2007.

- [101] David Ediger and David A Bader. Investigating graph algorithms in the bsp model on the cray xmt. In *Parallel and Distributed Processing Symposium Workshops & PhD Forum (IPDPSW), 2013 IEEE 27th International*, pages 1638–1645. IEEE, 2013.
- [102] David Ediger, Rob McColl, Jason Riedy, and David A Bader. Stinger: High performance data structure for streaming graphs. In *High Performance Extreme Computing (HPEC), 2012 IEEE Conference on*, pages 1–5. IEEE, 2012.
- [103] John Ellson, Emden Gansner, Lefteris Koutsofios, Stephen C North, and Gordon Woodhull. Graphviz open source graph drawing tools. In *International Symposium on Graph Drawing*, pages 483–484. Springer, 2001.
- [104] Daniel L Ensign, Peter M Kasson, and Vijay S Pande. Heterogeneity even at the speed limit of folding: large-scale molecular dynamics study of a fast-folding variant of the villin headpiece. *Journal of molecular biology*, 374(3):806–816, 2007.
- [105] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(17-61):43, 1960.
- [106] Paul Erds and Alfréd Rényi. On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.
- [107] Denis J Evans and Brad Lee Holian. The nose–hoover thermostat. *The Journal of chemical physics*, 83(8):4069–4074, 1985.
- [108] Shimon Even. *Graph algorithms*. Cambridge University Press, 2011.
- [109] Stephan Ewen, Kostas Tzoumas, Moritz Kaufmann, and Volker Markl. Spinning fast iterative data flows. *Proceedings of the VLDB Endowment*, 5(11):1268–1279, 2012.
- [110] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM computer communication review*, volume 29, pages 251–262. ACM, 1999.
- [111] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. On graph problems in a semi-streaming model. *Theoretical Computer Science*, 348(2):207–216, 2005.
- [112] Dima Feldman and Yuval Shavitt. Automatic large scale generation of internet pop level maps. In *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, pages 1–6. IEEE, 2008.
- [113] Dima Feldman, Yuval Shavitt, and Noa Zilberman. A structural approach for pop geo-location. *Computer Networks*, 56(3):1029–1040, 2012.
- [114] Alan R Fersht, Andreas Matouschek, and Luis Serrano. The folding of an enzyme: I. theory of protein engineering analysis of stability and pathway of protein folding. *Journal of molecular biology*, 224(3):771–782, 1992.
- [115] Robert W Floyd. Algorithm 97: shortest path. *Communications of the ACM*, 5(6):345, 1962.
- [116] Santo Fortunato. Community detection in graphs. *Physics reports*, 486(3):75–174, 2010.
- [117] Daniele Frigioni, Alberto Marchetti-Spaccamela, and Umberto Nanni. Semidynamic algorithms for maintaining single-source shortest path trees. *Algorithmica*, 22(3):250–274, 1998.
- [118] Daniele Frigioni, Alberto Marchetti-Spaccamela, and Umberto Nanni. Fully dynamic algorithms for maintaining shortest paths trees. *Journal of Algorithms*, 34(2):251–281, 2000.
- [119] Zhang Fuhao and Liu Jiping. An algorithm of shortest path based on dijkstra for huge data. In *Fuzzy Systems and Knowledge Discovery, 2009. FSKD'09. Sixth International Conference on*, volume 4, pages 244–247. IEEE, 2009.
- [120] Zvi Galil and Oded Margalit. All pairs shortest paths for graphs with small integer length edges. *Journal of Computer and System Sciences*, 54(2):243–254, 1997.
- [121] Alexandros V Gerbessiotis and Leslie G Valiant. Direct bulk-synchronous parallel algorithms. *Journal of parallel and distributed computing*, 22(2):251–267, 1994.

- [122] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The google file system. In *ACM SIGOPS operating systems review*, volume 37, pages 29–43. ACM, 2003.
- [123] Chris Giannella, Jiawei Han, Jian Pei, Xifeng Yan, and Philip S Yu. Mining frequent patterns in data streams at multiple time granularities. *Next generation data mining*, 212:191–212, 2003.
- [124] Michelle Girvan and Mark EJ Newman. Community structure in social and biological networks. *Proceedings of the national academy of sciences*, 99(12):7821–7826, 2002.
- [125] Jennifer Golbeck and James Hendler. Accuracy of metrics for inferring trust and reputation in semantic web-based social networks. In *Engineering knowledge in the age of the semantic web*, pages 116–131. Springer, 2004.
- [126] Jennifer Golbeck, Bijan Parsia, and James Hendler. *Trust networks on the semantic web*. Springer, 2003.
- [127] Andrew V Goldberg. *Efficient graph algorithms for sequential and parallel computers*. PhD thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 1987.
- [128] Mira Gonen and Yuval Shavitt. Approximating the number of network motifs. *Internet Mathematics*, 6(3):349–372, 2009.
- [129] Joseph E Gonzalez, Yucheng Low, Haijie Gu, Danny Bickson, and Carlos Guestrin. Powergraph: Distributed graph-parallel computation on natural graphs. In *Presented as part of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, pages 17–30, 2012.
- [130] Joseph E Gonzalez, Reynold S Xin, Ankur Dave, Daniel Crankshaw, Michael J Franklin, and Ion Stoica. Graphx: Graph processing in a distributed dataflow framework. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pages 599–613, 2014.
- [131] Oded Green, Robert McColl, and David A Bader. A fast algorithm for streaming betweenness centrality. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*, pages 11–20. IEEE, 2012.
- [132] Derek Greene, Donal Doyle, and Pdraig Cunningham. Tracking the evolution of communities in dynamic social networks. In *Advances in social networks analysis and mining (ASONAM), 2010 international conference on*, pages 176–183. IEEE, 2010.
- [133] Douglas Gregor and Andrew Lumsdaine. Lifting sequential graph algorithms for distributed-memory parallel computation. In *ACM SIGPLAN Notices*, volume 40, pages 423–437. ACM, 2005.
- [134] Sudipto Guha, Nina Mishra, Rajeev Motwani, and Liadan O’Callaghan. Clustering data streams. In *Foundations of computer science, 2000. proceedings. 41st annual symposium on*, pages 359–366. IEEE, 2000.
- [135] Roger Guimera, Stefano Mossa, Adrian Turttschi, and LA Nunes Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities’ global roles. *Proceedings of the National Academy of Sciences*, 102(22):7794–7799, 2005.
- [136] David Hales and Stefano Arteconi. Motifs in evolving cooperative networks look like protein structure networks. *Networks and Heterogeneous Media*, 3(2):239, 2008.
- [137] Minyang Han, Khuzaima Daudjee, Khaled Ammar, M Tamer Özsu, Xingfang Wang, and Tianqi Jin. An experimental comparison of pregel-like graph processing systems. *Proceedings of the VLDB Endowment*, 7(12):1047–1058, 2014.
- [138] Wentao Han, Youshan Miao, Kaiwei Li, Ming Wu, Fan Yang, Lidong Zhou, Vijayan Prabhakaran, Wenguang Chen, and Enhong Chen. Chronos: a graph engine for temporal graph analysis. In *Proceedings of the Ninth European Conference on Computer Systems*, page 1. ACM, 2014.
- [139] Pawan Harish and PJ Narayanan. Accelerating large graph algorithms on the gpu using cuda. In *International Conference on High-Performance Computing*, pages 197–208. Springer, 2007.
- [140] Michael Hay, Chao Li, Gerome Miklau, and David Jensen. Accurate estimation of the degree distribution of private networks. In *2009 Ninth IEEE International Conference on Data Mining*, pages 169–178. IEEE, 2009.

- [141] Dorothea Helmer, Ina Rink, James AR Dalton, Kevin Brahm, Marina Jöst, Tobias M Nargang, Witali Blum, Parvesh Wadhvani, Gerald Brenner-Weiss, Bastian E Rapp, et al. Rational design of a peptide capture agent for cxcl8 based on a model of the cxcl8: Cxcr1 complex. *RSC Advances*, 5(33):25657–25668, 2015.
- [142] M Rauch Henzinger, Prabhakar Raghavan, and Sridhar Rajagopalan. Computing on data streams. Technical report, Technical Note 1998-011, Digital Systems Research Center, Palo Alto, CA, 1998.
- [143] Andreas Hofer, Stefanie Roos, and Thorsten Strufe. Greedy embedding, routing and content addressing for darknets. In *Networked Systems (NetSys), 2013 Conference on*, pages 43–50. IEEE, 2013.
- [144] Petter Holme, Beom Jun Kim, Chang No Yoon, and Seung Kee Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109, 2002.
- [145] John Hopcroft and Robert Tarjan. Algorithm 447: efficient algorithms for graph manipulation. *Communications of the ACM*, 16(6):372–378, 1973.
- [146] Wenjie Hu, Yihua Liao, and V Rao Vemuri. Robust anomaly detection using support vector machines. In *Proceedings of the international conference on machine learning*, pages 282–289, 2003.
- [147] Pan Hui, Eiko Yoneki, Shu Yan Chan, and Jon Crowcroft. Distributed community detection in delay tolerant networks. In *Proceedings of 2nd ACM/IEEE international workshop on Mobility in the evolving internet architecture*, page 7. ACM, 2007.
- [148] E Gail Hutchinson and Janet M Thornton. Promotifa program to identify and analyze structural motifs in proteins. *Protein Science*, 5(2):212–220, 1996.
- [149] Anand Padmanabha Iyer, Li Erran Li, Tathagata Das, and Ion Stoica. Time-evolving graph processing at scale. In *Proceedings of the Fourth International Workshop on Graph Data Management Experiences and Systems*, page 5. ACM, 2016.
- [150] Matthew O Jackson et al. *Social and economic networks*, volume 3. Princeton university press Princeton, 2008.
- [151] Van Jacobson, Craig Leres, and S McCanne. The tcpdump manual page. *Lawrence Berkeley Laboratory, Berkeley, CA*, 1989.
- [152] Sushil Jajodia, Steven Noel, and Brian OBerry. Topological analysis of network attack vulnerability. In *Managing Cyber Threats*, pages 247–266. Springer, 2005.
- [153] Dusanka Janezic, Ante Milicevic, Sonja Nikolic, and Nenad Trinajstic. *Graph-theoretical matrices in chemistry*. CRC Press, 2015.
- [154] Madhav Jha, Comandur Seshadhri, and Ali Pinar. A space efficient streaming algorithm for triangle counting using the birthday paradox. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 589–597. ACM, 2013.
- [155] Salim Jouili and Valentin Vansteenberghe. An empirical comparison of graph databases. In *Social Computing (SocialCom), 2013 International Conference on*, pages 708–715. IEEE, 2013.
- [156] Hossein Jowhari and Mohammad Ghodsi. New streaming algorithms for counting triangles in graphs. In *International Computing and Combinatorics Conference*, pages 710–716. Springer, 2005.
- [157] Maliackal Poulo Joy, Amy Brock, Donald E Ingber, and Sui Huang. High-betweenness proteins in the yeast protein interaction network. *BioMed Research International*, 2005(2):96–103, 2005.
- [158] Changhee Jung, Silvius Rus, Brian P Railing, Nathan Clark, and Santosh Pande. Brainy: effective selection of data structures. In *ACM SIGPLAN Notices*, volume 46, pages 86–97. ACM, 2011.
- [159] David Jurgens and Tsai-Ching Lu. Temporal motifs reveal the dynamics of editor interactions in wikipedia. In *ICWSM*, 2012.
- [160] S Kalir, J McClure, K Pabbaraju, C Southward, M Ronen, S Leibler, MG Surette, and U Alon. Ordering genes in a flagella pathway by analysis of expression kinetics from living bacteria. *Science*, 292(5524):2080–2083, 2001.

- [161] Panos Kalnis, Karim Awara, Hani Jamjoom, and Zuhair Khayyat. Mizan: Optimizing graph mining in large parallel systems. Technical report, King Abdullah University of Science and Technology, 2012.
- [162] Daniel M Kane, Kurt Mehlhorn, Thomas Sauerwald, and He Sun. Counting arbitrary subgraphs in data streams. In *Automata, Languages, and Programming*, pages 598–609. Springer, 2012.
- [163] George Karypis and V Kumar. Metis: Family of multilevel partitioning algorithms. *wwwusers.cs.umn.edu/~karypis/memis/main.shtml*, 1995.
- [164] George Karypis and Vipin Kumar. Metis–unstructured graph partitioning and sparse matrix ordering system, version 2.0. 1995.
- [165] George Karypis and Vipin Kumar. Multilevel algorithms for multi-constraint graph partitioning. In *Proceedings of the 1998 ACM/IEEE conference on Supercomputing*, pages 1–13. IEEE Computer Society, 1998.
- [166] George Karypis and Vipin Kumar. Multilevel k-way partitioning scheme for irregular graphs. *Journal of Parallel and Distributed computing*, 48(1):96–129, 1998.
- [167] George Karypis and Vipin Kumar. A parallel algorithm for multilevel graph partitioning and sparse matrix ordering. *Journal of Parallel and Distributed Computing*, 48(1):71–95, 1998.
- [168] George Karypis, Kirk Schloegel, and Vipin Kumar. Parmetis: Parallel graph partitioning and sparse matrix ordering library. *Version 1.0, Dept. of Computer Science, University of Minnesota*, 1997.
- [169] Zahra RM Kashani, Hayedeh Ahrabian, Elahe Elahi, Abbas Nowzari-Dalini, Elnaz S Ansari, Sahar Asadi, Shahin Mohammadi, Falk Schreiber, and Ali Masoudi-Nejad. Kavosh: a new algorithm for finding network motifs. *BMC bioinformatics*, 10(1):318, 2009.
- [170] Nadav Kashtan, Shalev Itzkovitz, Ron Milo, and Uri Alon. Mfinder tool guide. *Department of Molecular Cell Biology and Computer Science and Applied Math., Weizmann Inst. of Science, Rehovot Israel, technical report*, 2002.
- [171] Nadav Kashtan, Shalev Itzkovitz, Ron Milo, and Uri Alon. Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. *Bioinformatics*, 20(11):1746–1758, 2004.
- [172] Gary J Katz and Joseph T Kider Jr. All-pairs shortest-paths for large graphs on the gpu. In *Proceedings of the 23rd ACM SIGGRAPH/EUROGRAPHICS symposium on Graphics hardware*, pages 47–55. Eurographics Association, 2008.
- [173] Richard A Kemmerer and Giovanni Vigna. Intrusion detection: A brief history and overview (supplement to computer magazine). *Computer*, (4):27–30, 2002.
- [174] Zuhair Khayyat, Karim Awara, Amani Alonazi, Hani Jamjoom, Dan Williams, and Panos Kalnis. Mizan: a system for dynamic load balancing in large-scale graph processing. In *Proceedings of the 8th ACM European Conference on Computer Systems*, pages 169–182. ACM, 2013.
- [175] Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William G Griswold. An overview of aspectj. In *European Conference on Object-Oriented Programming*, pages 327–354. Springer, 2001.
- [176] Jeong-Rae Kim, Yeoin Yoon, and Kwang-Hyun Cho. Coupled feedback loops form dynamic motifs of cellular networks. *Biophysical journal*, 94(2):359–365, 2008.
- [177] Valerie King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 81–89. IEEE, 1999.
- [178] Jon M Kleinberg, Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew S Tomkins. The web as a graph: Measurements, models, and methods. In *Computing and combinatorics*, pages 1–17. Springer, 1999.
- [179] Gerard J Kleywegt. Recognition of spatial motifs in protein structures. *Journal of molecular biology*, 285(4):1887–1897, 1999.

- [180] Lauri Kovanen, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(11):P11005, 2011.
- [181] Paul L Krapivsky and Sidney Redner. Organization of growing random networks. *Physical Review E*, 63(6):066123, 2001.
- [182] PL Krapivsky and S Redner. Network growth by copying. *Physical Review E*, 71(3):036118, 2005.
- [183] Lachezar Krumov, Immanuel Schweizer, Dirk Bradler, and Thorsten Strufe. Leveraging network motifs for the adaptation of structured peer-to-peer-networks. In *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*, pages 1–5. IEEE, 2010.
- [184] Sanjeev Kulkarni, Nikunj Bhagat, Maosong Fu, Vikas Kedigehalli, Christopher Kellogg, Sailesh Mittal, Jignesh M Patel, Karthik Ramasamy, and Siddarth Taneja. Twitter heron: Stream processing at scale. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 239–250. ACM, 2015.
- [185] Jérôme Kunegis. Konect: the koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 1343–1350. ACM, 2013.
- [186] Aapo Kyrola, Guy Blelloch, and Carlos Guestrin. Graphchi: large-scale graph computation on just a pc. In *Presented as part of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, pages 31–46, 2012.
- [187] Kisung Lee and Ling Liu. Scaling queries over big rdf graphs with semantic hash partitioning. *Proceedings of the VLDB Endowment*, 6(14):1894–1905, 2013.
- [188] Min-Joong Lee, Jungmin Lee, Jaimie Yejean Park, Ryan Hyun Choi, and Chin-Wan Chung. Qube: a quick algorithm for updating betweenness centrality. In *Proceedings of the 21st international conference on World Wide Web*, pages 351–360. ACM, 2012.
- [189] Jure Leskovec. Social media analytics: tracking, modeling and predicting the flow of information through networks. In *Proceedings of the 20th international conference companion on World wide web*, pages 277–278. ACM, 2011.
- [190] Jure Leskovec, Lars Backstrom, Ravi Kumar, and Andrew Tomkins. Microscopic evolution of social networks. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 462–470. ACM, 2008.
- [191] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):2, 2007.
- [192] Jure Leskovec, Kevin J Lang, Anirban Dasgupta, and Michael W Mahoney. Statistical properties of community structure in large social and information networks. In *Proceedings of the 17th international conference on World Wide Web*, pages 695–704. ACM, 2008.
- [193] Ian XY Leung, Pan Hui, Pietro Lio, and Jon Crowcroft. Towards real-time community detection in large networks. *Physical Review E*, 79(6):066107, 2009.
- [194] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944.
- [195] Richard Lippmann, Joshua W Haines, David J Fried, Jonathan Korba, and Kumar Das. The 1999 darpa off-line intrusion detection evaluation. *Computer networks*, 34(4):579–595, 2000.
- [196] Richard P Lippmann, David J Fried, Isaac Graf, Joshua W Haines, Kristopher R Kendall, David McClung, Dan Weber, Seth E Webster, Dan Wyszogrod, Robert K Cunningham, et al. Evaluating intrusion detection systems: The 1998 darpa off-line intrusion detection evaluation. In *DARPA Information Survivability Conference and Exposition, 2000. DISCEX'00. Proceedings*, volume 2, pages 12–26. IEEE, 2000.
- [197] Richard P Lippmann, Isaac Graf, Dan Wyszogrod, Seth E Webster, Dan J Weber, and Sam Gorton. The 1998 darpa/afri off-line intrusion detection evaluation. In *Proceedings of the First International Workshop on Recent Advances in Intrusion Detection (RAID). Louvain-la-Neuve, Belgium*, 1998.

- [198] Chengwen Liu and Hao Chen. A hash partition strategy for distributed query processing. In *International Conference on Extending Database Technology*, pages 371–387. Springer, 1996.
- [199] Peter S Loubai. A network evaluation procedure. *Highway Research Record*, (205), 1967.
- [200] Yucheng Low, Danny Bickson, Joseph Gonzalez, Carlos Guestrin, Aapo Kyrola, and Joseph M Hellerstein. Distributed graphlab: a framework for machine learning and data mining in the cloud. *Proceedings of the VLDB Endowment*, 5(8):716–727, 2012.
- [201] 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.
- [202] Yi Lu, James Cheng, Da Yan, and Huanhuan Wu. Large-scale distributed graph computing systems: An experimental evaluation. *Proceedings of the VLDB Endowment*, 8(3):281–292, 2014.
- [203] Chi-Keung Luk, Robert Cohn, Robert Muth, Harish Patil, Artur Klauser, Geoff Lowney, Steven Wallace, Vijay Janapa Reddi, and Kim Hazelwood. Pin: building customized program analysis tools with dynamic instrumentation. In *ACM Sigplan Notices*, volume 40, pages 190–200. ACM, 2005.
- [204] Andrew Lumsdaine, Douglas Gregor, Bruce Hendrickson, and Jonathan Berry. Challenges in parallel graph processing. *Parallel Processing Letters*, 17(01):5–20, 2007.
- [205] Xiaoke Ma and Lin Gao. Biological network analysis: insights into structure and functions. *Briefings in functional genomics*, 11(6):434–442, 2012.
- [206] D Margo P Macko, V Marathe, and M Seltzer. *LLAMA: Efficient Graph Analytics Using Large Multiversioned Arrays*. PhD thesis, Ph. D. Dissertation. Harvard University, 2014.
- [207] Kamesh Madduri and David A Bader. Compact graph representations and parallel connectivity algorithms for massive dynamic network analysis. In *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–11. IEEE, 2009.
- [208] Kamesh Madduri, David Ediger, Karl Jiang, David A Bader, and Daniel Chavarria-Miranda. A faster parallel algorithm and efficient multithreaded implementations for evaluating betweenness centrality on massive datasets. In *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–8. IEEE, 2009.
- [209] 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.
- [210] Gurmeet Singh Manku and Rajeev Motwani. Approximate frequency counts over data streams. In *Proceedings of the 28th international conference on Very Large Data Bases*, pages 346–357. VLDB Endowment, 2002.
- [211] Pekka Mark and Lennart Nilsson. Structure and dynamics of the tip3p, spc, and spc/e water models at 298 k. *The Journal of Physical Chemistry A*, 105(43):9954–9960, 2001.
- [212] Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [213] Claudio Martella, Roman Shaposhnik, and Dionysios Logothetis. *Graph in Action*. Manning, 2014.
- [214] Sergio Marti, Prasanna Ganesan, and Hector Garcia-Molina. Sprout: P2p routing with social networks. In *Current Trends in Database Technology-EDBT 2004 Workshops*, pages 425–435. Springer, 2005.
- [215] Sergei Maslov and Kim Sneppen. Specificity and stability in topology of protein networks. *Science*, 296(5569):910–913, 2002.
- [216] Oliver Mason and Mark Verwoerd. Graph theory and networks in biology. *IET systems biology*, 1(2):89–119, 2007.
- [217] Audrey Maudoux, Ph Lefebvre, J-E Cabay, Athina Demertzi, Audrey Vanhauzenhuyse, Steven Laureys, and Andrea Soddu. Connectivity graph analysis of the auditory resting state network in tinnitus. *Brain research*, 1485:10–21, 2012.

- [218] Julian J McAuley, Luciano da Fontoura Costa, and Tibério S Caetano. Rich-club phenomenon across complex network hierarchies. *Applied Physics Letters*, 91(8):084103, 2007.
- [219] Robert Ryan McCune, Tim Weninger, and Greg Madey. Thinking like a vertex: a survey of vertex-centric frameworks for large-scale distributed graph processing. *ACM Computing Surveys (CSUR)*, 48(2):25, 2015.
- [220] Andrew McGregor. Finding graph matchings in data streams. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 170–181. Springer, 2005.
- [221] Luis AA Meira, Vinicius R Máximo, Álvaro L Fazenda, and Arlindo F da Conceição. acc-motif detection tool. *arXiv preprint arXiv:1203.3415*, 2012.
- [222] Luis AA Meira, Vinicius R Maximo, Alvaro L Fazenda, and Arlindo F da Conceicao. Accelerated motif detection using combinatorial techniques. In *Signal Image Technology and Internet Based Systems (SITIS), 2012 Eighth International Conference on*, pages 744–753. IEEE, 2012.
- [223] A. Michels, A. PÜTZ, K.H. Maurer, T. Eggert, and K.E. JÄGER. Verwendung von esterassen zur spaltung von kunststoffen, July 20 2011. EP Patent 1,910,532.
- [224] Tijana Milenkoviæ and Nataša Pržulj. Uncovering biological network function via graphlet degree signatures. *Cancer informatics*, 6:257, 2008.
- [225] Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.
- [226] Alan Mislove, Massimiliano Marcon, Krishna P Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42. ACM, 2007.
- [227] Rolf H Möhring, Heiko Schilling, Birk Schütz, Dorothea Wagner, and Thomas Willhalm. Partitioning graphs to speedup dijkstra’s algorithm. *Journal of Experimental Algorithmics (JEA)*, 11:2–8, 2007.
- [228] Shanmugavelayutham Muthukrishnan. *Data streams: Algorithms and applications*. Now Publishers Inc, 2005.
- [229] David Myr. Real time vehicle guidance and forecasting system under traffic jam conditions, November 12 2002. US Patent 6,480,783.
- [230] Walid Najjar and J-L Gaudiot. Network resilience: A measure of network fault tolerance. *IEEE Transactions on Computers*, 39(2):174–181, 1990.
- [231] Mark EJ Newman. Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. *Physical review E*, 64(1):016132, 2001.
- [232] Mark EJ Newman. Assortative mixing in networks. *Physical review letters*, 89(20):208701, 2002.
- [233] Mark EJ Newman. Mixing patterns in networks. *Physical Review E*, 67(2):026126, 2003.
- [234] Mark EJ Newman. Fast algorithm for detecting community structure in networks. *Physical review E*, 69(6):066133, 2004.
- [235] Mark EJ Newman. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582, 2006.
- [236] Mark EJ Newman. The mathematics of networks. *The new palgrave encyclopedia of economics*, 2(2008):1–12, 2008.
- [237] Mark EJ Newman, Duncan J Watts, and Steven H Strogatz. Random graph models of social networks. *Proceedings of the National Academy of Sciences*, 99(suppl 1):2566–2572, 2002.
- [238] Nam P Nguyen, Thang N Dinh, Ying Xuan, and My T Thai. Adaptive algorithms for detecting community structure in dynamic social networks. In *INFOCOM, 2011 Proceedings IEEE*, pages 2282–2290. IEEE, 2011.
- [239] Frank H Niesen, Helena Berglund, and Masoud Vedadi. The use of differential scanning fluorimetry to detect ligand interactions that promote protein stability. *Nature protocols*, 2(9):2212–2221, 2007.

- [240] Joel Nishimura and Johan Ugander. Restreaming graph partitioning: simple versatile algorithms for advanced balancing. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1106–1114. ACM, 2013.
- [241] Andre Oboler. The rise and fall of a facebook hate group. *First Monday*, 13(11), 2008.
- [242] Saeed Omid, Falk Schreiber, and Ali Masoudi-Nejad. Moda: an efficient algorithm for network motif discovery in biological networks. *Genes & genetic systems*, 84(5):385–395, 2009.
- [243] Angela Orebaugh, Gilbert Ramirez, and Jay Beale. *Wireshark & Ethereal network protocol analyzer toolkit*. Syngress, 2006.
- [244] Edigerand others. Massive streaming data analytics: A case study with clustering coefficients. In *Parallel & Distributed Processing*, 2010.
- [245] Joshua OMadadhain, Danyel Fisher, Padhraic Smyth, Scott White, and Yan-Biao Boey. Analysis and visualization of network data using jung. *Journal of Statistical Software*, 10(2):1–35, 2005.
- [246] Matthew Felice Pace. Bsp vs mapreduce. *Procedia Computer Science*, 9:246–255, 2012.
- [247] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: bringing order to the web. 1999.
- [248] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005.
- [249] Christopher R Palmer and J Gregory Steffan. Generating network topologies that obey power laws. In *Global Telecommunications Conference, 2000. GLOBECOM'00. IEEE*, volume 1, pages 434–438. IEEE, 2000.
- [250] Simona Panni and Simona E Rombo. Searching for repetitions in biological networks: methods, resources and tools. *Briefings in bioinformatics*, 16(1):118–136, 2015.
- [251] Fragkiskos Papadopoulos, Maksim Kitsak, M Ángeles Serrano, Marián Boguná, and Dmitri Krioukov. Popularity versus similarity in growing networks. *Nature*, 489(7417):537–540, 2012.
- [252] Tiago P Peixoto. The graph-tool python library. *figshare*, 2014.
- [253] Asad Amir Pirzada and Chris McDonald. Establishing trust in pure ad-hoc networks. In *Proceedings of the 27th Australasian conference on Computer science-Volume 26*, pages 47–54. Australian Computer Society, Inc., 2004.
- [254] SC Ponten, Andreas Daffertshofer, Arjan Hillebrand, and Cornelis J Stam. The relationship between structural and functional connectivity: graph theoretical analysis of an eeg neural mass model. *Neuroimage*, 52(3):985–994, 2010.
- [255] Robert J Prill, Pablo A Iglesias, and Andre Levchenko. Dynamic properties of network motifs contribute to biological network organization. *PLoS biology*, 3(11):e343, 2005.
- [256] Michael J Quinn and Narsingh Deo. Parallel graph algorithms. *ACM Computing Surveys (CSUR)*, 16(3):319–348, 1984.
- [257] Raghavan et al. Computing on data streams. In *External Memory Algorithms*, 1999.
- [258] Usha Nandini Raghavan, Réka Albert, and Soundar Kumara. Near linear time algorithm to detect community structures in large-scale networks. *Physical review E*, 76(3):036106, 2007.
- [259] Supranamaya Ranjan, Joshua Robinson, and Feilong Chen. Machine learning based botnet detection using real-time connectivity graph based traffic features, June 24 2014. US Patent 8,762,298.
- [260] Vibhor Rastogi, Ashwin Machanavajjhala, Laukik Chitnis, and Anish Das Sarma. Finding connected components in map-reduce in logarithmic rounds. In *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*, pages 50–61. IEEE, 2013.
- [261] Mark Redekopp, Yogesh Simmhan, and Viktor K Prasanna. Optimizations and analysis of bsp graph processing models on public clouds. In *Parallel & Distributed Processing (IPDPS), 2013 IEEE 27th International Symposium on*, pages 203–214. IEEE, 2013.

- [262] Pedro Ribeiro and Fernando Silva. G-tries: an efficient data structure for discovering network motifs. In *Proceedings of the 2010 ACM Symposium on Applied Computing*, pages 1559–1566. ACM, 2010.
- [263] Jason Riedy, Henning Meyerhenke, David A Bader, David Ediger, and Timothy G Mattson. Analysis of streaming social networks and graphs on multicore architectures. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5337–5340. IEEE, 2012.
- [264] Matteo Riondato and Evgenios M Kornaropoulos. Fast approximation of betweenness centrality through sampling. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 413–422. ACM, 2014.
- [265] Liam Roditty and Uri Zwick. A fully dynamic reachability algorithm for directed graphs with an almost linear update time. *SIAM Journal on Computing*, 45(3):712–733, 2016.
- [266] Martin Roesch et al. Snort: Lightweight intrusion detection for networks. In *LISA*, volume 99, pages 229–238, 1999.
- [267] Hans Rohnert. A dynamization of the all pairs least cost path problem. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 279–286. Springer, 1985.
- [268] Amitabha Roy, Laurent Bindschaedler, Jasmina Malicevic, and Willy Zwaenepoel. Chaos: Scale-out graph processing from secondary storage. In *Proceedings of the 25th Symposium on Operating Systems Principles*, pages 410–424. ACM, 2015.
- [269] Amitabha Roy, Ivo Mihailovic, and Willy Zwaenepoel. X-stream: edge-centric graph processing using streaming partitions. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pages 472–488. ACM, 2013.
- [270] Loic Royer, Matthias Reimann, Bill Andreopoulos, and Michael Schroeder. Unraveling protein networks with power graph analysis. *PLoS computational biology*, 4(7):e1000108, 2008.
- [271] Mikail Rubinov and Olaf Sporns. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52(3):1059–1069, 2010.
- [272] Semih Salihoglu and Jennifer Widom. Gps: a graph processing system. In *Proceedings of the 25th International Conference on Scientific and Statistical Database Management*, page 22. ACM, 2013.
- [273] Chris Sanders and Jason Smith. *Applied network security monitoring: collection, detection, and analysis*. Elsevier, 2013.
- [274] Atish Das Sarma, Sreenivas Gollapudi, and Rina Panigrahy. Estimating pagerank on graph streams. *Journal of the ACM (JACM)*, 58(3):13, 2011.
- [275] Thomas Schank. Algorithmic aspects of triangle-based network analysis. *Phd in computer science, University Karlsruhe*, 3, 2007.
- [276] Thomas Schank and Dorothea Wagner. *Approximating clustering-coefficient and transitivity*. Universität Karlsruhe, Fakultät für Informatik, 2004.
- [277] Michael Schatz, Elliott Cooper-Balis, and Adam Bazinet. Parallel network motif finding. Technical report, Tech. rep., University of Maryland Insitute for Advanced Computer Studies, 2008.
- [278] Falk Schreiber and Henning Schwöbbermeyer. Mavisto: a tool for the exploration of network motifs. *Bioinformatics*, 21(17):3572–3574, 2005.
- [279] Daniel A Schult and P Swart. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conferences (SciPy 2008)*, volume 2008, pages 11–16, 2008.
- [280] John Scott. *Social network analysis*. Sage, 2012.
- [281] Jan Seedorf, Dirk Kutscher, and Fabian Schneider. Decentralised binding of self-certifying names to real-world identities for assessment of third-party messages in fragmented mobile networks. In *Computer Communications Workshops (INFOCOM WKSHPS), 2014 IEEE Conference on*, pages 416–421. IEEE, 2014.

- [282] Ohad Shacham, Martin Vechev, and Eran Yahav. Chameleon: adaptive selection of collections. In *ACM Sigplan Notices*, volume 44, pages 408–418. ACM, 2009.
- [283] Micha Sharir. A strong-connectivity algorithm and its applications in data flow analysis. *Computers & Mathematics with Applications*, 7(1):67–72, 1981.
- [284] Aneesh Sharma, Jerry Jiang, Praveen Bommannavar, Brian Larson, and Jimmy Lin. Graphjet: Real-time content recommendations at twitter. *Proceedings of the VLDB Endowment*, 9(13), 2016.
- [285] Shai S Shen-Orr, Ron Milo, Shmoolik Mangan, and Uri Alon. Network motifs in the transcriptional regulation network of escherichia coli. *Nature genetics*, 31(1):64–68, 2002.
- [286] Avi Shoshan and Uri Zwick. All pairs shortest paths in undirected graphs with integer weights. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 605–614. IEEE, 1999.
- [287] Julian Shun and Guy E Blelloch. Ligra: a lightweight graph processing framework for shared memory. In *ACM SIGPLAN Notices*, volume 48, pages 135–146. ACM, 2013.
- [288] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. The hadoop distributed file system. In *2010 IEEE 26th symposium on mass storage systems and technologies (MSST)*, pages 1–10. IEEE, 2010.
- [289] Luis Solá, Miguel Romance, Regino Criado, Julio Flores, Alejandro Garcia del Amo, and Stefano Boccaletti. Eigenvector centrality of nodes in multiplex networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 23(3):033131, 2013.
- [290] Robin Sommer and Vern Paxson. Outside the closed world: On using machine learning for network intrusion detection. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 305–316. IEEE, 2010.
- [291] Robin Sommer and Vern Paxson. Outside the closed world: On using machine learning for network intrusion detection. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 305–316. IEEE, 2010.
- [292] Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2004.
- [293] Stuart Staniford-Chen, Steven Cheung, Richard Crawford, Mark Dilger, Jeremy Frank, James Hoagland, Karl Levitt, Christopher Wee, Raymond Yip, and Dan Zerkle. Grids-a graph based intrusion detection system for large networks. In *Proceedings of the 19th national information systems security conference*, volume 1, pages 361–370. Baltimore, 1996.
- [294] Isabelle Stanton and Gabriel Kliot. Streaming graph partitioning for large distributed graphs. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1222–1230. ACM, 2012.
- [295] Michael Stonebraker, Uur Çetintemel, and Stan Zdonik. The 8 requirements of real-time stream processing. *ACM SIGMOD Record*, 34(4):42–47, 2005.
- [296] Philip Stutz, Abraham Bernstein, and William Cohen. Signal/collect: graph algorithms for the (semantic) web. In *International Semantic Web Conference*, pages 764–780. Springer, 2010.
- [297] Jimeng Sun, Christos Faloutsos, Spiros Papadimitriou, and Philip S Yu. Graphscope: parameter-free mining of large time-evolving graphs. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 687–696. ACM, 2007.
- [298] Jimeng Sun, Yinglian Xie, Hui Zhang, and Christos Faloutsos. Less is more: Compact matrix decomposition for large sparse graphs. In *SDM*, pages 366–377. SIAM, 2007.
- [299] Yan Lindsay Sun, Wei Yu, Zhu Han, and KJ Ray Liu. Information theoretic framework of trust modeling and evaluation for ad hoc networks. *Selected Areas in Communications, IEEE Journal on*, 24(2):305–317, 2006.
- [300] Siddharth Suri and Sergei Vassilvitskii. Counting triangles and the curse of the last reducer. In *Proceedings of the 20th international conference on World wide web*, pages 607–614. ACM, 2011.

- [301] Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2):146–160, 1972.
- [302] Ramona S Taylor, Liem X Dang, and Bruce C Garrett. Molecular dynamics simulations of the liquid/vapor interface of spc/e water. *The Journal of Physical Chemistry*, 100(28):11720–11725, 1996.
- [303] 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.
- [304] Claire Tomlin, George J Pappas, and Shankar Sastry. Conflict resolution for air traffic management: A study in multiagent hybrid systems. *IEEE Transactions on automatic control*, 43(4):509–521, 1998.
- [305] Ngoc Hieu Tran, Kwok Pui Choi, and Louxin Zhang. Counting motifs in the human interactome. *Nature communications*, 4, 2013.
- [306] Sacha Trifunovic, Franck Legendre, and Carlos Anastasiades. Social trust in opportunistic networks. In *INFOCOM IEEE Conference on Computer Communications Workshops, 2010*, pages 1–6. IEEE, 2010.
- [307] Charalampos Tsourakakis, Christos Gkantsidis, Bozidar Radunovic, and Milan Vojnovic. Fennel: Streaming graph partitioning for massive scale graphs. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 333–342. ACM, 2014.
- [308] Charalampos E Tsourakakis. Fast counting of triangles in large real networks without counting: Algorithms and laws. In *2008 Eighth IEEE International Conference on Data Mining*, pages 608–617. IEEE, 2008.
- [309] Charalampos E Tsourakakis, U Kang, Gary L Miller, and Christos Faloutsos. Doulion: counting triangles in massive graphs with a coin. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 837–846. ACM, 2009.
- [310] Alexander Ulrich, Ralph Holz, Peter Hauck, and Georg Carle. Investigating the openpgp web of trust. In *Computer Security–ESORICS 2011*, pages 489–507. Springer, 2011.
- [311] Leslie G Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.
- [312] Luis M Vaquero and Manuel Cebrian. The rich club phenomenon in the classroom. *Scientific reports*, 3, 2013.
- [313] Vinod Kumar Vavilapalli, Arun C Murthy, Chris Douglas, Sharad Agarwal, Mahadev Konar, Robert Evans, Thomas Graves, Jason Lowe, Hitesh Shah, Siddharth Seth, et al. Apache hadoop yarn: Yet another resource negotiator. In *Proceedings of the 4th annual Symposium on Cloud Computing*, page 5. ACM, 2013.
- [314] Alexei Vázquez. Growing network with local rules: Preferential attachment, clustering hierarchy, and degree correlations. *Physical Review E*, 67(5):056104, 2003.
- [315] Srdjan Čapkun, Levente Buttyán, and Jean-Pierre Hubaux. Small worlds in security systems: an analysis of the PGP certificate graph. In *Proceedings of the 2002 workshop on New security paradigms*, pages 28–35. ACM, 2002.
- [316] Chad Vicknair, Michael Macias, Zhendong Zhao, Xiaofei Nan, Yixin Chen, and Dawn Wilkins. A comparison of a graph database and a relational database: a data provenance perspective. In *Proceedings of the 48th annual Southeast regional conference*, page 42. ACM, 2010.
- [317] Chris Walshaw, Mark Cross, and Martin G Everett. Parallel dynamic graph partitioning for adaptive unstructured meshes. *Journal of Parallel and Distributed Computing*, 47(2):102–108, 1997.
- [318] Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge university press, 1994.
- [319] Duncan J Watts and Steven H Strogatz. Collective dynamics of small-world networks. *nature*, 393(6684):440–442, 1998.
- [320] Duncan J Watts and Steven H Strogatz. Collective dynamics of small-world networks. *nature*, 393(6684):440–442, 1998.

- [321] Jim Webber. A programmatic introduction to neo4j. In *Proceedings of the 3rd annual conference on Systems, programming, and applications: software for humanity*, pages 217–218. ACM, 2012.
- [322] Sebastian Wernicke. A faster algorithm for detecting network motifs. In *International Workshop on Algorithms in Bioinformatics*, pages 165–177. Springer, 2005.
- [323] Sebastian Wernicke. Efficient detection of network motifs. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 3(4):347–359, 2006.
- [324] Sebastian Wernicke and Florian Rasche. Fanmod: a tool for fast network motif detection. *Bioinformatics*, 22(9):1152–1153, 2006.
- [325] Jason Weston, Sayan Mukherjee, Olivier Chapelle, Massimiliano Pontil, Tomaso Poggio, and Vladimir Vapnik. Feature selection for svms. 2000.
- [326] Tom White. *Hadoop: The definitive guide*. ” O’Reilly Media, Inc.”, 2012.
- [327] GJ Van Wylen and RE Sonntag. Fundamentals of classical thermodynamics. 2013.
- [328] Li Xiao, Lixue Chen, and Jingzhong Xiao. A new algorithm for shortest path problem in large-scale graph. *Appl. Math*, 6(3):657–663, 2012.
- [329] Jierui Xie, Stephen Kelley, and Boleslaw K Szymanski. Overlapping community detection in networks: The state-of-the-art and comparative study. *ACM Computing Surveys (csur)*, 45(4):43, 2013.
- [330] Reynold S Xin, Daniel Crankshaw, Ankur Dave, Joseph E Gonzalez, Michael J Franklin, and Ion Stoica. Graphx: Unifying data-parallel and graph-parallel analytics. *arXiv preprint arXiv:1402.2394*, 2014.
- [331] Reynold S Xin, Joseph E Gonzalez, Michael J Franklin, and Ion Stoica. Graphx: A resilient distributed graph system on spark. In *First International Workshop on Graph Data Management Experiences and Systems*, page 2. ACM, 2013.
- [332] Reynold S Xin, Josh Rosen, Matei Zaharia, Michael J Franklin, Scott Shenker, and Ion Stoica. Shark: Sql and rich analytics at scale. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of data*, pages 13–24. ACM, 2013.
- [333] Guoqing Xu. Coco: sound and adaptive replacement of java collections. In *European Conference on Object-Oriented Programming*, pages 1–26. Springer, 2013.
- [334] Da Yan, James Cheng, Yi Lu, and Wilfred Ng. Blogel: A block-centric framework for distributed computation on real-world graphs. *Proceedings of the VLDB Endowment*, 7(14):1981–1992, 2014.
- [335] Da Yan, James Cheng, Kai Xing, Yi Lu, Wilfred Ng, and Yingyi Bu. Pregel algorithms for graph connectivity problems with performance guarantees. *Proceedings of the VLDB Endowment*, 7(14):1821–1832, 2014.
- [336] Wan-Shiou Yang, Jia-Ben Dia, Hung-Chi Cheng, and Hsing-Tzu Lin. Mining social networks for targeted advertising. In *Proceedings of the 39th Annual Hawaii International Conference on System Sciences (HICSS’06)*, volume 6, pages 137a–137a. IEEE, 2006.
- [337] Hilmi Yildirim, Vineet Chaoji, and Mohammed J Zaki. Dagger: A scalable index for reachability queries in large dynamic graphs. *arXiv preprint arXiv:1301.0977*, 2013.
- [338] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J Franklin, Scott Shenker, and Ion Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 2–2. USENIX Association, 2012.
- [339] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: cluster computing with working sets. *HotCloud*, 10:10–10, 2010.
- [340] Matei Zaharia, Tathagata Das, Haoyuan Li, Timothy Hunter, Scott Shenker, and Ion Stoica. Discretized streams: Fault-tolerant streaming computation at scale. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pages 423–438. ACM, 2013.
- [341] Ellen W Zegura, Kenneth L Calvert, and Michael J Donahoo. A quantitative comparison of graph-based models for internet topology. *IEEE/ACM Transactions on Networking (TON)*, 5(6):770–783, 1997.

- [342] Jiong Zhang, Mohammad Zulkernine, and Anwar Haque. Random-forests-based network intrusion detection systems. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 38(5):649–659, 2008.
- [343] Zhao Zhao, Maleq Khan, VS Anil Kumar, and Madhav V Marathe. Subgraph enumeration in large social contact networks using parallel color coding and streaming. In *Parallel Processing (ICPP), 2010 39th International Conference on*, pages 594–603. IEEE, 2010.
- [344] Da Zheng, Disa Mhembere, Randal Burns, Joshua Vogelstein, Carey E Priebe, and Alexander S Szalay. Flashgraph: Processing billion-node graphs on an array of commodity ssds. In *13th USENIX Conference on File and Storage Technologies (FAST 15)*, pages 45–58, 2015.
- [345] Jianlong Zhong and Bingsheng He. An overview of medusa: simplified graph processing on gpus. In *ACM SIGPLAN Notices*, volume 47, pages 283–284. ACM, 2012.
- [346] Jianlong Zhong and Bingsheng He. Medusa: Simplified graph processing on gpus. *IEEE Transactions on Parallel and Distributed Systems*, 25(6):1543–1552, 2014.
- [347] Shi Zhou and Raúl J Mondragón. Accurately modeling the internet topology. *Physical Review E*, 70(6):066108, 2004.
- [348] Shi Zhou and Raúl J Mondragón. The rich-club phenomenon in the internet topology. *IEEE Communications Letters*, 8(3):180–182, 2004.
- [349] Shi Zhou and Raúl J Mondragón. The positive-feedback preference model of the as-level internet topology. In *IEEE International Conference on Communications, 2005. ICC 2005. 2005*, volume 1, pages 163–167. IEEE, 2005.
- [350] Philip R Zimmermann. *The official PGP user's guide*. MIT press, 1995.
- [351] Uri Zwick. All pairs shortest paths in weighted directed graphs-exact and almost exact algorithms. In *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on*, pages 310–319. IEEE, 1998.

Own Publications

- [s1] Dirk Bradler, Benjamin Schiller, Erwin Aitenbichler, and Nicolas Liebau. Towards a distributed crisis response communication system. volume 9, pages 10–13, 2009.
- [s2] Tim Grube, Benjamin Schiller, and Thorsten Strufe. Monotone sampling of networks. In *Proceedings of the 2nd International Workshop on Dynamic Networks and Knowledge Discovery*, page 37. Citeseer, 2014.
- [s3] Irina Heimbach, Benjamin Schiller, Thorsten Strufe, and Oliver Hinz. Content virality on online social networks: Empirical evidence from twitter, facebook, and google+ on german news websites. In *Proceedings of the 26th ACM Conference on Hypertext & Social Media*, pages 39–47. ACM, 2015.
- [s4] Sven Jager, Benjamin Schiller, Thorsten Strufe, and Kay Hamacher. Stream-t.g: Algorithms for analyzing coarse grained rna dynamics based on markov models of connectivity-graphs. In *International Workshop on Algorithms in Bioinformatics*, pages 197–209. Springer, 2016.
- [s5] Giang Nguyen, Stefanie Roos, Benjamin Schiller, and Thorsten Strufe. Swap: Protecting pull-based p2p video streaming systems from inference attacks. In *2016 IEEE 17th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pages 1–9. IEEE, 2016.
- [s6] Stefanie Roos, Benjamin Schiller, Stefan Hacker, and Thorsten Strufe. Measuring freenet in the wild: Censorship-resilience under observation. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 263–282. Springer, 2014.
- [s7] Hani Salah, Benjamin Schiller, and Thorsten Strufe. Comon: A system architecture for improving caching in ccn. In *Computer Communications Workshops (INFOCOM WKSHPS), 2014 IEEE Conference on*, pages 213–214. IEEE, 2014.
- [s8] Benjamin Schiller, Dirk Bradler, Immanuel Schweizer, Max Mühlhäuser, and Thorsten Strufe. Gtna: a framework for the graph-theoretic network analysis. In *Proceedings of the 2010 Spring Simulation Multiconference*, page 111. Society for Computer Simulation International, 2010.
- [s9] Benjamin Schiller, Jeronimo Castrillon, and Thorsten Strufe. Efficient data structures for dynamic graph analysis. In *2015 11th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*, pages 497–504. IEEE, 2015.
- [s10] Benjamin Schiller, Clemens Deusser, Jeronimo Castrillon, and Thorsten Strufe. Compile- and runtime approaches for the selection of efficient data structures for dynamic graph analysis. *Journal of Applied Network Science*, 1(9), 2016.
- [s11] Benjamin Schiller, Sven Jager, Kay Hamacher, and Thorsten Strufe. Stream-a stream-based algorithm for counting motifs in dynamic graphs. In *International Conference on Algorithms for Computational Biology*, pages 53–67. Springer, 2015.
- [s12] Benjamin Schiller, Giang Nguyen, and Thorsten Strufe. Resilient tree-based live streaming in reality. In *IEEE P2P 2013 Proceedings*, pages 1–2. IEEE, 2013.
- [s13] Benjamin Schiller, Giang Nguyen, and Thorsten Strufe. Resilient tree-based live streaming for mobile scenarios. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2014 IEEE International Conference on*, pages 160–162. IEEE, 2014.
- [s14] Benjamin Schiller, Stefanie Roos, Andreas Hofer, and Thorsten Strufe. Attack resistant network embeddings for darknets. In *Reliable Distributed Systems Workshops (SRDSW), 2011 30th IEEE Symposium on*, pages 90–95. IEEE, 2011.

- [s15] Benjamin Schiller and Thorsten Strufe. Dynamic network analyzer building a framework for the graph-theoretic analysis of dynamic networks. In *Proceedings of the 2013 Summer Computer Simulation Conference*, page 49. Society for Modeling & Simulation International, 2013.
- [s16] Benjamin Schiller and Thorsten Strufe. Gtna 2.0-a framework for rapid prototyping and evaluation of routing algorithms. In *Proceedings of the 2013 Summer Computer Simulation Conference*, page 23. Society for Modeling & Simulation International, 2013.
- [s17] Benjamin Schiller, Thorsten Strufe, Dirk Kohlweyer, and Jan Seedorf. Growing a web of trust. In *Local Computer Networks (LCN), 2015 IEEE 40th Conference on*, pages 100–108. IEEE, 2015.
- [s18] Jan Seedorf, Bilal Gill, Dirk Kutscher, Benjamin Schiller, and Dirk Kohlweyer. Demo overview: fully decentralised authentication scheme for icn in disaster scenarios. In *Proceedings of the 1st international conference on Information-centric networking*, pages 191–192. ACM, 2014.

Online Sources

(last visited: 27 October 2016)

- [o1] ACC. <http://www.ft.unicamp.br/docentes/meira/accmotifs/>.
- [o2] Amazon Kinesis. <https://aws.amazon.com/kinesis/>.
- [o3] Amazon S3. <https://aws.amazon.com/s3/>.
- [o4] Analysis of the PGO WoT. <http://pgp.cs.uu.nl/plot>.
- [o5] Argus (netflow generator). <http://qosient.com/argus>.
- [o6] AspectJ. <https://eclipse.org/aspectj/>.
- [o7] Blogel. <http://www.cse.cuhk.edu.hk/blogel/>.
- [o8] Boost Graph Library. http://www.boost.org/doc/libs/1_61_0/libs/graph/doc/.
- [o9] Cassandra. <http://cassandra.apache.org>.
- [o10] Cisco routers for recording packet traces. <http://www.cisco.com/go/epc>.
- [o11] DARPA'98 dataset. <https://www.ll.mit.edu/ideval/data/1998data.html>.
- [o12] DB-Engines Ranking of Graph DBMS. <http://db-engines.com/en/ranking/graph+dbms>.
- [o13] Docker. <https://www.docker.com>.
- [o14] Facebook dataset. <http://konect.uni-koblenz.de/networks/facebook-wosn-links>.
- [o15] Fanmod. <http://theinf1.informatik.uni-jena.de/~wernicke/motifs/>.
- [o16] FlashGraph. <http://flashx.io>.
- [o17] Flink. <http://flink.apache.org>.
- [o18] Flink: Gelly. <https://flink.apache.org/news/2015/08/24/introducing-flink-gelly.html>.
- [o19] Flink: Spargel. https://ci.apache.org/projects/flink/flink-docs-release-0.8/spargel_guide.html.
- [o20] Flow Exporter. <http://www.flowtraq.com/corporate/product/flow-exporter>.
- [o21] Flume. <https://flume.apache.org>.
- [o22] G-Tries. <http://www.dcc.fc.up.pt/gtries/>.
- [o23] Gephi. <https://gephi.org>.
- [o24] Giraph. <http://giraph.apache.org>.
- [o25] Gnuplot. <http://gnuplot.sourceforge.net>.
- [o26] Gnuplot: fit command. http://gnuplot.sourceforge.net/docs_4.2/node82.html.
- [o27] Graph Processing System. <http://infolab.stanford.edu/gps/>.
- [o28] Graph-Tool. <https://graph-tool.skewed.de>.
- [o29] GraphChi. <https://github.com/GraphChi>.
- [o30] GraphJet. <https://github.com/twitter/GraphJet>.

- [o31] GraphLab. https://turi.com/products/create/open_source.html.
- [o32] GraphStream. <http://graphstream-project.org>.
- [o33] Graphviz. <http://www.graphviz.org>.
- [o34] Gremlin. <http://gremlindocs.spmallette.documentup.com>.
- [o35] Guava. <https://github.com/google/guava>.
- [o36] Hadoop. <http://hadoop.apache.org>.
- [o37] Hadoop: YARN. <https://hadoop.apache.org/docs/r2.7.0/hadoop-yarn/hadoop-yarn-site/>.
- [o38] Hama. <http://hama.apache.org>.
- [o39] HBase. <http://hbase.apache.org>.
- [o40] Heron. <http://twitter.github.io/heron/>.
- [o41] Hypertext 2009 dataset. <http://konect.uni-koblenz.de/networks/sociopatterns-hypertext>.
- [o42] JChart2D. <http://jchart2d.sourceforge.net>.
- [o43] JGraphT Library. <http://jgrapht.org>.
- [o44] jpathwatch. <https://jpathwatch.wordpress.com>.
- [o45] JUNG. <http://jung.sourceforge.net>.
- [o46] Kafka. <http://kafka.apache.org>.
- [o47] Kavosh (*offline*). <http://lbb.ut.ac.ir/Download/LBBsoft/Kavosh/>.
- [o48] KONECT: The Koblenz Network Collection. <http://konect.uni-koblenz.de>.
- [o49] Kudu. <http://kudu.apache.org>.
- [o50] Ligra. <http://jshun.github.io/ligra/>.
- [o51] LinkedIn. <https://www.linkedin.com>.
- [o52] MapR FS. <https://www.mapr.com/products/mapr-fs>.
- [o53] Medusa. <https://github.com/JianlongZhong/Medusa>.
- [o54] Mesos. <http://mesos.apache.org>.
- [o55] Mizan. <https://github.com/khayyatzy/Mizan>.
- [o56] Monte Media Library. <http://www.randelshofer.ch/monte/>.
- [o57] Neo4J. <https://neo4j.com>.
- [o58] NetworkX. <http://networkx.github.io>.
- [o59] Okapi. <http://grafos.ml/index.html#Okapi>.
- [o60] openCypher. <http://www.opencypher.org>.
- [o61] OpenMP. <http://openmp.org>.
- [o62] OrientDB. <http://orientdb.com>.
- [o63] Pajek. <http://mrvar.fdv.uni-lj.si/pajek/>.
- [o64] PGP WoT graphs. <http://www.lysator.liu.se/~jc/wotsap/wots2>.
- [o65] PGP WoT graphs. <https://wot.siccegge.de>.
- [o66] PowerGraph. <https://github.com/jegonzal/PowerGraph>.
- [o67] Pregelix. <http://pregelix.ics.uci.edu>.

- [o68] Protein Data Bank: 1IL8. <http://www.rcsb.org/pdb/explore.do?structureId=1il8>.
- [o69] Protein Data Bank: 4AOA. <http://www.rcsb.org/pdb/explore.do?structureId=4AOA>.
- [o70] PyMOL. <https://www.pymol.org>.
- [o71] RabbitMQ. <https://www.rabbitmq.com>.
- [o72] Samza. <http://samza.apache.org>.
- [o73] Scikit-Learn. <http://scikit-learn.org>.
- [o74] Scikit-Learn: Lasso Scores. http://scikit-learn.org/stable/modules/generated/sklearn.linear_model.Lasso.html.
- [o75] Scikit-Learn: Random Forest classifier. <http://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html>.
- [o76] Scikit-Learn: Random Forest Regressor. <http://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestRegressor.html>.
- [o77] Scikit-Learn: Support Vector Machine classifier. <http://scikit-learn.org/stable/modules/generated/sklearn.svm.LinearSVC.html>.
- [o78] Signal/Collect. <http://uzh.github.io/signal-collect/>.
- [o79] SNAP: Small-world Network Analysis and Partitioning. <http://snap-graph.sourceforge.net>.
- [o80] Spark. <http://spark.apache.org>.
- [o81] Spark: GraphX. <http://spark.apache.org/graphx/>.
- [o82] Spark: SQL. <http://spark.apache.org/sql/>.
- [o83] Spark: Streaming. <http://spark.apache.org/streaming/>.
- [o84] Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data/>.
- [o85] Stanford Network Analysis Platform (SNAP for C++). <http://snap.stanford.edu/snap/>.
- [o86] Stanford Network Analysis Platform (SNAP for Python). <http://snap.stanford.edu/snappy/>.
- [o87] Stanford Network Analysis Project. <http://snap.stanford.edu>.
- [o88] STINGER. <http://www.stingergraph.com>.
- [o89] Storm. <http://storm.apache.org>.
- [o90] Storm: YARN. <https://github.com/yahoo/storm-yarn>.
- [o91] Tinkerpop. <http://tinkerpop.apache.org>.
- [o92] Twitter Cassovary. <https://github.com/twitter/cassovary>.
- [o93] Twitter Streaming API. <https://twitter.com/streaming/overview/>.
- [o94] X-Stream. <http://labos.epfl.ch/x-stream>.
- [o95] ZeroMQ. <http://zeromq.org>.

Repos and Online Sources (last visited: 27 October 2016)

- [c1] DNA. <http://dynamic-networks.org/dna/>.
- [c2] DNA Doc: Gvis. <https://github.com/BenjaminSchiller/DNA.doc/tree/master/gvis>.
- [c3] DNA Doc: Plotting. <https://github.com/BenjaminSchiller/DNA.doc/tree/master/plotting>.
- [c4] DNA Formats. <https://github.com/BenjaminSchiller/DNA.doc/blob/master/doc/FORMATS.md>.
- [c5] DNA Repo. <https://github.com/BenjaminSchiller/DNA/>.
- [c6] DNA.diss Repo. <https://github.com/BenjaminSchiller/DNA.diss/>.
- [c7] DNA.gdsAnalysis Repo. <https://github.com/BenjaminSchiller/DNA.gdsAnalysis/>.
- [c8] DNA.gdsMeasurements Repo. <https://github.com/BenjaminSchiller/DNA.gdsMeasurements/>.
- [c9] DnaWrapper Repo. <https://github.com/BenjaminSchiller/DnaWrapper>.
- [c10] GTNA. <http://dynamic-networks.org/gtna/>.
- [c11] Stream Repo. <https://github.com/BenjaminSchiller/Stream>.
- [c12] WoT models. <https://github.com/BenjaminSchiller/GTNA/tree/master/src/gtna/networks/model/wot>.

