































- (2016).
- [4] Michael J. Carey and Donald Kossmann. 1997. Processing top n and bottom n queries. *IEEE Data Eng. Bull.* 20, 3 (1997), 12–19.
- [5] Michael J Carey and Donald Kossmann. 1998. Reducing the braking distance of an SQL query engine. In *VLDB*, Vol. 98. 24–27.
- [6] Josephine Cheng, Don Haderle, Richard Hedges, Balakrishna R Iyer, Ted Messinger, C Mohan, and Yun Wang. 1991. An efficient hybrid join algorithm: A DB2 prototype. In [1991] *Proceedings. Seventh International Conference on Data Engineering*. IEEE, 171–180.
- [7] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. 2009. Power-law distributions in empirical data. *SIAM review* 51, 4 (2009), 661–703.
- [8] Transaction Processing Performance Council. 2008. TPC-H benchmark specification. Published at <http://www.tpc.org/hspec.html> 21 (2008), 592–603.
- [9] Jialin Ding, Umar Farooq Minhas, Hantian Zhang, Yinan Li, Chi Wang, Badrish Chandramouli, Johannes Gehrke, Donald Kossmann, and David Lomet. 2019. ALEX: an updatable adaptive learned index. *arXiv preprint arXiv:1905.08898* (2019).
- [10] Ronald Fagin. 2016. Optimal Score Aggregation Algorithms. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. ACM, 55–55.
- [11] Christos Faloutsos and HV Jagadish. 1992. On B-tree indices for skewed distributions. (1992).
- [12] Goetz Graefe. 1993. Query evaluation techniques for large databases. *ACM Computing Surveys (CSUR)* 25, 2 (1993), 73–169.
- [13] Goetz Graefe. 2006. Implementing sorting in database systems. *ACM Computing Surveys (CSUR)* 38, 3 (2006), 10.
- [14] Goetz Graefe. 2008. A general and efficient algorithm for "top" queries. In *Data Engineering Workshop, 2008. ICDEW 2008. IEEE 24th International Conference on*. IEEE, 548–555.
- [15] Julian Huxley, Richard E Strauss, and Frederick B Churchill. 1932. Problems of relative growth. (1932).
- [16] Ihab F Ilyas, George Beskales, and Mohamed A Soliman. 2008. A survey of top-k query processing techniques in relational database systems. *ACM Computing Surveys (CSUR)* 40, 4 (2008), 11.
- [17] Business Insider. 2019. Facebook Photos Statistics. Retrieved 02/16/2019 from <https://www.businessinsider.com/facebook-350-million-photos-each-day-2013-9>
- [18] Business Insider. 2020. Amazon Prime Users Statistics. Retrieved 04/09/2020 from <https://www.businessinsider.com/amazon-more-than-100-million-prime-members-us-survey-2019-1>
- [19] Herald Killapi, Eva Sitaridi, Manolis M Tsangaris, and Yannis Ioannidis. 2011. Schedule optimization for data processing flows on the cloud. In *Proceedings of the 2011 International Conference on Management of Data*. ACM, 289–300.
- [20] Donald Ervin Knuth. 1973. *The art of computer programming: sorting and searching*. Vol. 3. Pearson Education.
- [21] D Kossmann and M Carey. 1997. On saying "enough already!". In *SQL, in Proc. of the 1997 ACM-SIGMOD Conference on Management of Data, Tucson, Arizona*.
- [22] Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. 2018. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*. 489–504.
- [23] Chengkai Li, Kevin Chen-Chuan Chang, Ihab F Ilyas, and Sumin Song. 2005. RankSQL: query algebra and optimization for relational top-k queries. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*. ACM, 131–142.
- [24] Wentian Li. 2002. Zipf's Law everywhere. *Glottometrics* 5 (2002), 14–21.
- [25] Tian Mi and Sanguthevar Rajasekaran. 2013. A two-pass exact algorithm for selection on Parallel Disk Systems. In *2013 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 000612–000617.
- [26] Chris Nyberg, Tom Barclay, Zarka Cvetanovic, Jim Gray, and Dave Lomet. 1994. AlphaSort: A RISC machine sort. In *ACM SIGMOD Record*, Vol. 23. ACM, 233–242.
- [27] Chris Nyberg, Tom Barclay, Zarka Cvetanovic, Jim Gray, and Dave Lomet. 1995. Alphasort: A cache-sensitive parallel external sort. *The VLDB Journal* 4, 4 (1995), 603–627.
- [28] David MW Powers. 1998. Applications and explanations of Zipf's law. In *Proceedings of the joint conferences on new methods in language processing and computational natural language learning*. Association for Computational Linguistics, 151–160.
- [29] Bart Samwel, John Cieslewicz, Ben Handy, Jason Govig, Petros Venetis, Chanjun Yang, Keith Peters, Jeff Shute, Daniel Tenedorio, Himani Apte, Felix Weigel, David Willhite, Jiacheng Yang, Jun Xu, Jiexing Li, Zhan Yuan, Craig Chasseur, Qiang Zeng, Ian Rae, Anurag Biyani, Andrew Harn, Yang Xia, Andrey Gubichev, Amr El-Helw, Orri Erling, Zhepeng Yan, Mohan Yang, Yiqun Wei, Thanh Do, Colin Zheng, Goetz Graefe, Somayeh Sardashti, Ahmed M. Aly, Divy Agrawal, Ashish Gupta, and Shiv Venkataraman. 2018. F1 Query: Declarative Querying at Scale. *Proc. VLDB Endow.* 11, 12 (Aug. 2018), 1835–1848. <https://doi.org/10.14778/3229863.3229871>
- [30] Anil Shanbhag, Holger Pirk, and Samuel Madden. 2018. Efficient Top-K Query Processing on Massively Parallel Hardware. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD '18)*. Association for Computing Machinery, New York, NY, USA, 1557–1570. <https://doi.org/10.1145/3183713.3183735>
- [31] David Simmen, Eugene Shekita, and Timothy Malkemus. 1996. Fundamental Techniques for Order Optimization. *SIGMOD Rec.* 25, 2 (June 1996), 57–67. <https://doi.org/10.1145/235968.233320>
- [32] Internet Live Stats. 2020. Twitter statistics. Retrieved 04/09/2020 from <https://www.internetlivestats.com/twitter-statistics/>
- [33] Peter Van Sandt, Yannis Chronis, and Jignesh M Patel. 2019. Efficiently Searching In-Memory Sorted Arrays: Revenge of the Interpolation Search?. In *Proceedings of the 2019 International Conference on Management of Data*. ACM, 36–53.
- [34] The Verge. 2019. Google's billion user services. Retrieved 04/09/2020 from <https://www.theverge.com/2019/7/24/20708328/google-photos-users-gallery-go-1-billion>
- [35] Peifeng Yin, Ping Luo, Wang-Chien Lee, and Min Wang. 2013. Silence is also evidence: interpreting dwell time for recommendation from psychological perspective. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 989–997.