P2P数据管理

余 敏, 李战怀, 张龙波

余 敏, 李战怀, 张龙波

(西北工业大学 计算机学院,陕西 西安 710072)

作者简介: 余敏(1980－),女,江西波阳人,博士生,主要研究领域为数据库理论与技术,P2P数据管理.李战怀(1961－),男,博士,教授,博士生导师,CCF高级会员,主要研究领域为数据库理论与技术.张龙波(1968－),男,博士生,副教授,主要研究领域为数据流管理,查询处理.

联系人: 余 敏 Phn: +86-29-88493772, E-mail: yum@mail.nwpu.edu.cn, http://www.nwpu.edu.cn

Abstract

P2P (peer-to-peer) is the key technology of reconstructing the future distributed architecture and has a good application perspective. As the issues in P2P systems mostly come down to data placement and retrieval, P2P data management has recently become an active topic in database community. In this paper, the advantages of P2P systems are first described. Then the goals of P2P data management researches are presented. Thirdly, research of P2P data management is described from three facets, i.e. P2P information retrieval, P2P database-style queries and P2P continuous queries. Particularly, the index construction methods, semantic coordination, query semantics, query processing strategies, types of queries supported, and query optimization of P2P database-style queries are discussed in detail. Finally, the issues to be further studied are proposed.

摘要

P2P(peer-to-peer)技术是未来重构分布式体系结构的关键技术,拥有广阔的应用前景.P2P系统的大多数问题都可归结为数据放置和检索问题,因此,P2P数据管理成为数据库领域活跃的研究课题.当前,P2P数据管理主要有信息检索、数据库查询和连续查询3个子领域,取得了许多研究成果.在介绍P2P技术的优点后,指出了P2P数据管理研究的目标.然后针对上述3个方面,论述P2P数据管理研究的现状,着重讨论了P2P数据库查询的索引构造策略、语义异构的解决方法、查询语义、查询处理策略、查询类型和查询优化技术.通过比较,指出了现状与目标的差距,提出了需要进一步研究的问题.

References:

[1] Gribble SD, Halevy AY, Ives ZG, Rodrig M, Suciu D. What can databases do for peer-to-peer- In: Mecca G, Simeon J, eds. Proc. (informal) of the 4th Int'l Workshop on the Web and Databases (WebDB). Santa Barbara, 2001. 31-36.

[2] Ling B, Lu ZG, Ng WS, Qian WN, Zhou AY. PeerIS: A peer-to-peer based information retrieval system. Journal of Software, 2004,15(9):1375-1384 (in Chinese with English abstract). http://www.jos.org.cn/1000-9825/15/1375.htm

[3] Aberer K, Hauswirth M. An overview on peer-to-peer information systems. In: Litwin W, Lévy G, eds. Proc. in Informatics (14). Waterloo: Carleton Scientific, 2002. 171-188.

[4] Batini C. A survey of data quality issues in cooperative information systems. In: Tutorial of the 23rd Int'l Conf. on Conceptual Modeling (ER 2004). Shanghai: Fudan University, 2004.

[5] Bawa M, Cooper BF, Crespo A, Daswani N, Ganesan P, Garcia-Molina H, Kamvar S, Marti S, Schlosser M, Sun Q, Vinograd P, Yang B. Peer-to-Peer research at Stanford. ACM SIGMOD Record, 2003,32(3):23-28.

[6] Crainiceanu A, Linga P, Gehrke J, Shanmugasundaram J. Querying peer-to-peer networks using p-trees. In: Amer-Yahia S, Gravano L, eds. Proc. of the 7th Int'l Workshop on Web and Databases. New York: ACM Press, 2004. 25-30.

[7] Heubsch R, Hellerstein JM, Lanham N, Loo BT, Shenker S, Stocia I. Querying the Internet with PIER. In: Freytag JC, Lockemann PC, Abiteboul S, Carey MJ, Selinger PG, Heuer A, eds. Proc. of the 29th Int'l Conf. on Very Large Data Bases. San Fransisco: Morgan Kaufmann Publishers, 2003. 321-332.

[8] Tatarinov I, Halevy A. Efficient query reformulation in peer data management systems. In: Weikum G, Konig AC, DeBloch S, eds. Proc of the ACM SIGMOD Int'l Conf. on the Management of Data. Paris: ACM, 2004. 539-550.

[9] Cai M, Frank M. RDFPeers: A scalable distributed RDF repository based on a structured peer-to-peer network. In: Tolles-Efinger L, ed. Proc. of the 13th Int'l World Wide Web Conf. New York: Sheridan Printing, 2004. 650-657.

[10] Kementsietsidis A, Arenas M. Data sharing through query translation in autonomous sources. In: Nascimento MA, Ozsu MT, Kossmann D, Miller RJ, Blakeley JA, Schiefer KB, eds. Proc. of the 30th Int'l Conf. on Very Large Data Bases. San Fransisco: Morgan Kaufmann Publishers, 2004. 468-479.

[11] Nejdl W, Siberski W, Sintek M. Design issues and challenges for RDF and schema-based peer-to-peer systems. ACM SIGMOD Record, 2003,32(3):41-46.

[12] Aberer K, Cudré-Mauroux P, Hauswirth M, van Pelt T. GridVine: Building Internet-scale semantic overlay networks. Lecture Notes in Computer Science, 2004, 3298:107-121.

[13] Chirita PA, Idreos S, Koubarakis M, Nejdl W. Publish/Subscribe for RDF-based P2P networks. Lecture Notes in Computer Science, 2004,3053:182-197.

[14] Ng WS, Ooi BC, Shu YF, Tan KL, Tok WH. Efficient distributed continuous query processing using peers. Technical Report, NUS-CS01-03, Kent Ridge: National University of Singapore, 2003. 1-13.

[15] Ng WS, Ooi BC, Tan KL, Zhou AY. PeerDB: A P2P-based system for distributed data sharing. In: Dayal U, ed. Proc. of the 19th Int'l Conf. on Data Engineering (ICDE). Bangalore: IEEE Computer Society Press, 2003. 633-644.

[16] Daswani N, Garcia-Molina H, Yang B. Open problems in data-sharing peer-to-peer systems. Lecture Notes in Computer Science, 2003,2572:1-15.

[17] Harren M, Hellerstein JM, Huebsch R, Loo BT, Shenker S, Stoica I. Complex queries in DHT-based peer-to-peer networks. Lecture Notes in Computer Science, 2002,2429:242-259.

[18] Joseph S, Hoshiai T. Decentralized meta-data strategies: Effective peer-to-peer search. IEICE Trans. on Communications, 2003, E85-B(6):1740-1753.

[19] Weikum G. Towards a statistically semantic Web. Lecture Notes in Computer Science, 2004,3288:3-17.

[20] Galanis L, Wang Y, Jeffery SR, DeWitt DJ. Processing queries in a large peer-to-peer system. Lecture Notes in Computer Science, 2003,2681:273-288.

[21] Bawa M, Manku GS, Raghavan P. SETS: Search enhanced by topic segmentation. In: Callan J, Cormack G, Clarke C, Hawking D, Smeaton A, eds. Proc. of the 26th Annual Int'l ACM SIGIR Conf. on Research and Development in Information Retrieval. New York: ACM Press, 2003. 306-313.

[22] Risson J, Moors T. Survey of research towards robust peer-to-peer networks: Search methods. Technical Report, UNSWEE-P2P-1-1, Sydney: University of New South Wales, 2004. 1-36.

[23] Shen HT, Shu YF, Yu B. Efficient semantic-based content search in P2P network. IEEE Trans. on Knowledge and Data Engineering, 2004,17(7):813-826.

[24] Balke WT. Supporting information retrieval in peer-to-peer systems. Lecture Notes in Computer Science, 2005,3485:337-352.

[25] Ganesan P, Sun Q, Garcia-Molina H. Adlib: A self-tuning index for dynamic peer-to-peer systems. In: Kawada S, ed. Proc. of the 21st Int'l Conf. on Data Engineering (ICDE). Tokyo: IEEE Computer Society, 2005. 256-257.

[26] Aberer K. Guest editor's introduction. ACM SIGMOD Record, 2003,32(3):21-22.

[27] Huebsch R, Chun B, Hellerstein J, Loo BT, Maniatis P, Roscoe T, Shenker S, Stoica I, Yumerefendi AR. The architecture of PIER: An Internet-scale query processor. In: Stonebraker M, Weikum G, DeWitt D, eds. Proc. of the 2005 Conf. on Innovative Data Systems Research. Asilomar: VLDB, 2005. 28-43.

[28] Aberer K, Cudre-Maroux P, Hauswirth M, Van Pelt T. Start making sense: The chatty Web approach for global semantic agreements. Journal of Web Semantics, 2004,1(1):72-86.

[29] Ooi BC, Shu YF, Tan KL. Relational data sharing in peer-based data management systems. ACM SIGMOD Record, 2003,32(3): 59-64.

[30] Hellerstein JM. Toward network data independence. ACM SIGMOD Record, 2003,32(3):34-40.

[31] Zhou AY. The understanding and consideration on several frontiers of database. In: Proc. of the 2004 Symp. on Database Development Strategy (in Chinese). Shenyang, 2004 (in Chinese with English abstract).

[32] Crainiceanu A, Linga P, Machanavjjhala A, Gehrke J, Shanmugasundaram J. P-Ring: An index structure for peer-to-peer systems. Technical Report, TR2004-1946, New York: Cornell University, 2004. 1-19.

[33] Crainiceanu A, Linga P, Machanavajjhala A, Gehrke J, Shanmugasundaram J. An indexing framework for peer-to-peer systems. In: Weikum G, Konig AC, DeBloch S, eds. Proc. of the ACM SIGMOD Int'l Conf. on the Management of Data (Demo). Paris: ACM, 2004. 939-940.

[34] Yang B, Garcia-Molina H. Improving search in peer-to-peer networks. In: Rodrigues LET, Raynal M, Chen WSE, eds. Proc. of the 22nd Int'l Conf. on Distributed Computing Systems. Washington: IEEE Computer Society, 2002. 5-14.

[35] Lü Q, Ratnasamy S, Shenker S. Can heterogeneity make Gnutella scalable- Lecture Notes in Computer Science, 2002,2429:94-103.

[36] Lü Q, Cao P, Cohen E, Li K, Shenker S. Search and replication in unstructured peer-to-peer networks. In: Ebcioglu K, Pingali K, Nicolau A, eds. Proc. of the 16th Int'l Conf. on Supercomputing. New York: ACM Press, 2002. 84-95.

[37] Banaei-Kashani F, Shahabi C. Criticality-Based analysis and design of unstructured peer-to-peer networks as "complex systems". In: Yokokawa M, ed. Proc. of the 3rd IEEE/ACM Int'l Symp. on Cluster Computing and the Grid. Tokyo: IEEE Computer Society, 2003. 351-358.

[38] Doan A, Noy D, Halevy A. Introduction to the special issue on semantic integration. ACM SIGMOD Record, 2004,33(4):11-13.

[39] Bernstein P, Giunchigalia F, Kementsietsidis A, Mylopoulos J, Serafini L, Zaihrayeu I. Data management for peer-to-peer computing: A vision. In: Fernandez MF, Papakonstantinous Y, eds. Proc. of the 5th Int'l Workshop on the Web and Databases (WebDB 2002). Madison: Wisconsin, 2002. 89-94.

[40] Kantere V, Kiringa I, Mylopoulos J, Kementisietsidis A, Arenas M. Coordinating peer databases using ECA rules. Lecture Notes in Computer Science, 2004, 2944:108-122.

[41] Calvanese D, De Giacomo G, Lenzerini M, Rosati R. Logical foundations of peer-to-peer data integration. In: Deutsch A, ed. Proc. of the 23rd ACM SIGMOD-SIGACT-SIGART Symp. on Principles of Database Systems (PODS 2004). Paris: ACM, 2004. 241-251.

[42] Nejdl W, Siberski W. Schema-Based peer-to-peer systems. Lecture Notes in Computer Science, 2005,3485:323-336.

[43] Brunkhorst I, Dhraief H, Kemper A, Nejdl W, Wiesner C. Distributed queries and query optimization in schema-based P2P-systmes. Lecture Notes in Computer Science, 2004,2944:184-199.

[44] Papadimos V, Maier D. Distributed queries without distributed state. In: Fernandex MF, Papkonstantinou Y, eds. Proc. of the 5th Int'l Workshop on the Web and Databases (WebDB 2002). Madison: Wisconsin, 2002. 95-100.

[45] Papadimos V, Maier D, Tufte K. Distributed query processing and catalogs for peer-to-peer systems. In: Proc. (online) of the 1st Biennial Conf. on Innovative Data Systems Research (CIDR 2003). Asilornar: Wisconsin, 2003.

[46] Agrawal BAM, Seshan S. Mercury: Supporting scalable multi-attribute range queries. Computer Communication Review, 2004, 34 (4):353-366.

[47] Gupta A, Agrawal D, Abbadi AE. Approximate range selection queries in peer-to-peer systems. In: Proc. (online) of the 1st Biennial Conf. on Innovative Data Systems Research (CIDR 2003). Asilornar: Wisconsin, 2003.

[48] Sahin OD, Gupta A, Agrawal D, Abbadi AE. A peer-to-peer framework for caching range queries. In: Rundensteiner E, ed. Proc. of the 20th Int'l Conf. on Data Engineering. Boston: IEEE Computer Society, 2004. 165-176.

[49] Aspnes J, Shah G. Skip graphs. In: Proc. of the 14th Annual ACM-SIAM Symp. on Discrete Algorithms. Philadelphia: Society for Industrial and Applied Mathematics, 2003. 384-393.

[50] Ganesan P, Bawa M, Carcia-Molina H. Online balancing of range-partitioned data with applications to peer-to-peer systems. In: Nascimento MA, Ozsu MT, Kossmann D, Miller RJ, Blakeley JA, Schiefer KB, eds. Proc. of the 30th Int'l Conf. on Very Large Data Bases. San Fransisco: Morgan Kaufmann Publishers, 2004. 444-455.

[51] Daskos A, Ghandeharizadeh S, An X. PePeR: A distributed range addressing space for peer-to-peer systems. Lecture Notes in Computer Science, 2004,2944:165-176.

[52] Ramabhadran S, Ratnasamy S, Hellerstein JM, Shenker S. Brief announcement: Prefix Hash tree. In: Chaudhuri S, Kutten S, eds. Proc. of the 23rd Annual ACM Symp. on Principles of Distributed Computing. St.John's: ACM, 2004. 368.

[53] Aberer K. P-Grid: A self-organizing access structure for P2P information systems. Lecture Notes in Computer Science, 2001,2172:179-194.

[54] Freedman MJ, Vingralek R. Efficient peer-to-peer lookup based on a distributed trie. Lecture Notes in Computer Science, 2002,2429:66-75.

[55] Rowstron A, Druschel P. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. Lecture Notes in Computer Science, 2001,2218:329-350.

[56] Linga P, Crainiceanu A, Gehrke J, Shanmugasundaram J. Guaranteeing correctness and availability in P2P range indices. In: Ozcan F, ed. Proc. of the ACM SIGMOD Int'l Conf. on Management of data. Baltimore: ACM, 2005. 323-334.

[57] Triantafillow P, Pitoura T. Towards a unifying framework for complex query processing over structured peer-to-peer data networks. Lecture Notes in Computer Science, 2004,2944:169-183.

[58] Bawa M, Garcia-Molina H, Gionis A, Motwani R. Estimating aggregates on a peer-to-peer networks. Technical Report, 2003-24, Stanford: Stanford University, 2003. 1-13.

[59] Bawa M, Gionis A, Garcia-Molina H, Motwani R. The price of validity in dynamic networks. In: Weikum G, Konig AC, DeBloch S, eds. Proc. of the ACM SIGMOD Int'l Conf. on the Management of Data. Paris: ACM, 2004. 515-526.

[60] Loo BT, Huebsch R, Hellerstein JM, Roscoe T, Stoica I. Analyzing P2P overlays with recursive queries. Technical Report, UCB//CSD-04-1301, Berkeley: Computer Science Division, UC Berkeley, 2004. 1-5.

[61] Avnur R, Hellerstein JM. Eddies: Continuously adaptive query processing. In: Chen WD, Naughton JF, Bernstein PA, eds. Proc. of the 2000 ACM SIGMOD Int'l Conf. on Management of Data. Dallas: ACM, 2000. 261-272.

[62] Gedik B, Liu L. PeerCQ: A decentralized and self-configuring peer-to-peer information monitoring system. In: Titsworth FM, ed. Proc. of the 23rd IEEE Int'l Conf. on Distributed Computer Systems. Providence: IEEE Computer Society, 2003. 490-499.

[63] Zdonik S, Stonebraker M, Chemiack M, Centintemel U, Balazinska M, Balakrishna H. The aurora and medusa project. IEEE Data Engineering Bulletin, 2003,26(1):3-10.

附中文参考文献:

[2] 凌波,陆志国,黄维维,钱卫宁,周敖英.PeerIS:基于Peer-to-Peer的信息检索系统.软件学报,2004,15(9):1375-1384. http://www.jos.org.cn/1000-9825/15/1375.htm

[31] 周敖英.若干数据库前沿技术的理解与思考.见:2004年数据库发展战略研讨会.沈阳,2004.