Department of Computer Science, The University of Hong Kong

CS Seminar

Low-Rank Mechanism: Optimizing Batch Queries under Differential Privacy

Dr. Zhenjie Zhang
Advanced Digital Sciences Center
University of Illinois at Urbana Champaign

Date:
June 28, 2012
Thursday
10:30 am

Venue:
Room 328
Chow Yei Ching Building
The University of Hong Kong

Abstract:

Differential privacy is a promising privacy-preserving paradigm for statistical query processing over sensitive data. It works by injecting random noise into each query result, such that it is provably hard for the adversary to infer the presence or absence of any individual record from the published noisy results. The main objective in differentially private query processing is to maximize the accuracy of the query results, while satisfying the privacy guarantees. Previous work, notably the matrix mechanism, has suggested that processing a batch of correlated queries as a whole can potentially achieve considerable accuracy gains, compared to answering them individually. However, as we point out in this paper, the matrix mechanism is mainly of theoretical interest; in particular, several inherent problems in its design limit its accuracy in practice, which almost never exceeds that of naive methods. In fact, we are not aware of any existing solution that can effectively optimize a query batch under differential privacy. Motivated by this, we propose the Low-Rank Mechanism (LRM), the first practical differentially private technique for answering batch queries with high accuracy, based on a low rank approximation of the workload matrix. We prove that the accuracy provided by LRM is close to the theoretical lower bound for any mechanism to answer a batch of queries under differential privacy. Extensive experiments using real data demonstrate that LRM consistently outperforms state-of-the-art query processing solutions under differential privacy, by large margins.

About the Speaker:

Zhenjie Zhang is currently research scientist in Advanced Digital Sciences Center, University of Illinois at Urbana Champaign. He received his Ph.D. in computer science from the School of Computing, National University of Singapore, in 2010. Before that, he graduated with a B.S. degree from the Department of Computer Science and Engineering, Fudan University, in 2004. He was visiting student at the Hong Kong University of Science and Technology in 2008 and a visiting student at AT&T Shannon Lab in 2009. Before joining the Advanced Digital Sciences Center in October 2010, he worked as a Research Assistant and Research Fellow at the National University of Singapore from 2008 to 2010. His research interests cover a wide spectrum of computer science, including real-time analytics, non-metric indexing, game theory and data privacy. He has served as a Program Committee member for VLDB 2012, ICDE 2012, WWW 2010, VLDB 2010, KDD 2010 and other conferences. He was the recipient of a NUS President's Graduate Fellowship of National University of Singapore in 2007.

All are welcome!
For enquiries, please call 2859 2180 or email enquiry@cs.hku.hk
Department of Computer Science
The University of Hong Kong