

OTM'10 Keynote

Beng Chin Ooi

Beng Chin Ooi, National University of Singapore (NUS)

Short Bio

Beng Chin is Professor of Computer Science at School of Computing, at the National University of Singapore (NUS). He obtained his BSc (1st Class Honors) and PhD from Monash University, Australia, in 1985 and 1989 respectively. His research interests include database performance issues, indexing techniques, multimedia and spatio-temporal databases, P2P systems and advanced applications, and cloud computing. His current system projects include BestPeer, P2P based data management system, and epiC, a data-intensive cloud computing platform.

He has served as a PC member for international conferences including ACM SIGMOD, VLDB, IEEE ICDE, WWW, SIGKDD and Vice PC Chair for ICDE'00, 04,06, co-PC Chair for SSD'93 and DASFAA'05, PC Chair for ACM SIGMOD'07, and Core DB track PC chair for VLDB'08. He is the Editor-in-Chief of IEEE Transactions on Knowledge and Data Engineering (TKDE), and a trustee member of VLDB Endowment Board. He is the recipient of ACM SIGMOD 2009 Contributions award.

Talk

“Supporting OLTP and OLAP Queries on Cloud Platforms”

MapReduce-based systems have been widely used for large-scale data analysis. Although these systems achieve storage-system independence, high scalability, and fine-grained fault tolerance, their performance have been shown to be unsatisfactory. It has also been shown that MapReduce-based systems are significantly slower than Parallel Database systems in performing a variety of analytic tasks. Some attribute the performance gap between MapReduce-based and Parallel Database systems to architectural design. This speculation yields an interesting question: Must a system sacrifice performance to achieve flexibility and scalability? Consequently, we conducted an in-depth performance study of MapReduce in its open source implementation, Hadoop. We identified various factors that have significant performance effect on the system. Subsequently, based on what we have learned, we propose a new architectural design as an attempt to support both OLTP and OLAP queries on Cloud platforms. I shall describe some of our ongoing work in this talk.