CSCI 599: NewSQL Database Management Systems

Prerequisites: Either CSCI 485, 585, 685 or consent of the instructor
Time: Spring 2012, Monday and Wednesdays 10 to 11:20 pm
Instructor: Shahram Ghandeharizadeh, shahram@usc.edu, 213-740-4781
Office: SAL 208

Introduction and Purposes

NoSQL Database Management Systems (DBMSs) are maturing to offer some form of a query language. For example, Cassandra Query Language, CQL, resembles SQL sprinkled with key-value functionality. Another simultaneous trend extends SQL DBMSs with key-value stores that act as cache managers. Examples include USC's COSAR and MIT's CacheGenie. The query interface of these systems is either an extended SQL or an object relational mapping framework such as Django or Rails. These trends are introducing hybrid systems that offer both the flexibility of a key-value store and expressiveness of SQL. The purpose of this seminar is to provide an overview of these NewSQL DBMSs. More specifically, we present and discuss:

1. Key-value stores such as Berkeley DB, LinkedIn's Voldermort, Riak, Redis, and Amazon's Dynamo,
2. Key-value SQL DBMSs such as USC's COSAR and MIT's CacheGenie,
3. Document stores such as CouchDB, MongoDB, and SimpleDB,
4. Extensible record stores such as HBase, Facebook's Cassandra, and Google's BigTable,
5. Scalable RDMBS such as VoltDB, Clustrix, and ScaleDB.

In addition, we explore applications of these systems and their definition of consistency (strong, weak, and eventual).

Course Requirements and Grades

- There are no required text books. The reading material is based on recently published technical papers available via the ACM/IEEE/Springer digital libraries. All USC students have automatic access to these digital archives.
- Grading breakdown
  - Exam 1: 30%
  - Exam 2: 30%
  - Class project: 20%
  - Class attendance: 10%
  - Homework assignment: 10%
- In class attendance is mandatory due to the nature of the technical papers discussed and presented. Absence from lecture will almost certainly impact student performance for both midterms.

Spring 2012 At a Glance

Here is an overview of the Spring 2012 semester calendar and its reading material. Papers covered by each topic are detailed below.

Jan 9, 11: Introduction
Jan 16: Martin Luther King Day
Jan 18: Introduction Continued
Jan 23, 25: Web, Databases and Caches

Jan 30, Feb 1: Web, Databases and Caches
Feb 6: Key-Value Stores
Feb 8: Projects start
Feb 13, 15: Key-Value Stores
Feb 20, 22: Key-Value SQL DBMSs
Feb 27, 29: Scalable RDBMS
March 5: Review for Exam 1
March 7: Exam 1 covers material presented from Jan 9 to Feb 29
March 12, 14: Spring Break
March 19, 21: Scalable RDBMS
March 26: In-Class presentation on projects
March 28, April 2: Cloud and Benchmarks
April 4, 9: Extensible Stores
April 11, 16: Application Transparent Mid-Tier Database Caching
April 18: Project reviews
April 23: Review for Exam 2
April 25: Exam 2 covers material presented from March 19 to April 16

Course Readings/Class Sessions
The reading material for the course is organized chronologically and based on a specific theme. This material is tentative and based on recent publications. We will make adjustments to the list as new manuscripts (currently under review) are accepted for publication.

Introduction
- R. Cattell. Scalable SQL and NoSQL Data Stores. White Paper, June 12, 2011.
- M. Stonebraker and R. Cattell. 10 Rules for Scalable Performance in Simple Operation Datastores. Communications of the ACM, June 2011, Vol. 54, No. 6.
- A. Labrindis and N. Roussopoulos. Exploring the Trade-off Between Performance and Data Freshness in Database-Driven Web Servers. VLDB Journal, Volume 13, 2004.

Web, Databases, and Caches
- K. S. Candan, W.-S Li, Q. Lou, W.P.Hsiung, D. Agrawal. Enabling Dynamic Content Caching for Database-Driven Web Sites. SIGMOD 2001.
- J. Challenger, A. Iyengar, P. Dantzig. A Scalable System for Consistently Caching Dynamic Web Data. IEEE INFOCOM, 1999.

Key-Value Stores
- B. Debnath, S. Sengupta, Jin Li. SkimpyStash: RAM Space Skimpy Key-Value Store on Flash-based Storage. SIGMOD 2011.
- M. Seltzer. Beyond Relational Databases. Communications of the ACM, July 2008, Vol. 51, No. 7.
- G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: Amazon's Highly Available Key-Value Store. In SOSP 2007.

Suggested Reading:

- J. Rao, E. J. Shekita, S. Tata: Using Paxos to Build a Scalable, Consistent, and Highly Available Datastore. PVLDB 4(4): 243-254 (2011)

Key-Value SQL DBMSs

- S. Ghandeharizadeh and J. Yap. SQL To Trigger Translation: An Application Transparent Approach to Cache Consistency. Submitted for publication.
- P. Gupta, N. Zeldovich and S. Madden. CacheGenie: A Trigger-Based Middleware Cache for ORMs. Middleware 2011, December 2011.

Scalable RDBMS

- S. Kadambi, J. Chen, B. F. Cooper, D. Lomax, R. Ramakrishnan, A. Silberstein, H. Garcia-Molina. *Where in the World is My Data?* VLDB Conference, August 2011.
- R. Kallman, J. Natkins, H. Kimura, A. Pavlo, A. Rasin, S. Zdonik, E. Jone, S. Madden, M. Stonebraker, D. Abadi. H-Store: A High-Performance, Distributed Main Memory Transaction Processing System. VLDB 2008.

Suggested reading:

- B. F. Cooper, R. Ramakrishnan, U. Srivastava, A. Silberstein, P. Bohannon, H. Jacobsen, N. Puz, D. Weaver and R. Yerneni. *PNUTS: Yahoo!'s Hosted Data Serving Platform*. VLDB Conference (industry track), 2008.
- M. Stonebraker, S. Madden, D. Abadi, S. Harizopoulos, N. Hachem and P. Helland. The End of an Architectural Era (It's Time for a Complete Rewrite). In VLDB 2007.
- P. Helland. Life Beyond Distributed Transactions: An Apostate's Opinion. In CIDR 2007.

Extensible Stores

- F. Chang et al. Bigtable: A Distributed Storage System for Structured Data. In OSDI 2006.
- O. Kennedy, Y. Ahmad, C. Koch. DBToaster: Agile Views for a Dynamic Data Management System. CIDR 2011.

Cloud and Benchmarks

- S. Das, D. Agrawal, and A. El Abbadi. G-Store: A Scalable Data Store for Transactional Multi key Access in the Cloud. SoCC 2010.
- Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, R. Sears. Benchmarking Cloud Serving Systems with YCSB. SoCC 2010.
- S. Patil, M. Polte, K. Ren, W. Tantisiriroj, L. Xiao, J. Lopez, G. Gibson, A. Fuchs, B. Rinaldi. YCSB++: Benchmarking and Performance Debugging Advanced Features in Scalable Table Stores, SoCC 2011.

Application Transparent Mid-Tier Database Caching

- P. Larson, J. Goldstein, and J. Zhou. MTCache: Transparent Mid-Tier Database Caching in SQL Server. In Proceedings of the 20th International Conference on Data Engineering. Page 177-189, April 2004.
- C. Bornhovd, M. Altinel, S. Krinshnamurthy, C. Mohan, H. Pirahesh, B. Reinwald. DBCache: Middle-Tier database Caching for Highly Scalable e-Business Architectures. SIGMOD 2003.

- C. Garrod et. al. Scalable Query Result Caching for Web Applications. VLDB 1998.

Suggested Reading:

- The Times Ten Team. High Performance and Scalability Through Application-Tier In-Memory Data Management. VLDB 2000.
- Lou and J. Naughton. Form-Based Proxy Caching for Database Backed Web Sites. VLDB 2001.

Statement for Students with Disabilities

Any student requesting academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to me (or to TA) as early in the semester as possible. DSP is located in STU 301 and is open 8:30 a.m.–5:00 p.m., Monday through Friday. The phone number for DSP is (213) 740-0776.

Statement on Academic Integrity

USC seeks to maintain an optimal learning environment. General principles of academic honesty include the concept of respect for the intellectual property of others, the expectation that individual work will be submitted unless otherwise allowed by an instructor, and the obligations both to protect one's own academic work from misuse by others as well as to avoid using another's work as one's own. All students are expected to understand and abide by these principles. *Scampus,* the Student Guidebook, contains the Student Conduct Code in Section 11.00, while the recommended sanctions are located in Appendix A: http://www.usc.edu/dept/publications/SCAMPUS/gov/. Students will be referred to the Office of Student Judicial Affairs and Community Standards for further review, should there be any suspicion of academic dishonesty. The Review process can be found at: http://www.usc.edu/student-affairs/SJACS/.