

Tian Luo

luot@cse.ohio-state.edu

<http://www.cse.ohio-state.edu/~luot>

Education

2008.9 ~ present: Ph.D. student (Advisor: Professor Xiaodong Zhang)

CS&E Department, The Ohio State University, USA

2005.9 ~ 2008.7: M.S. Computer Science

Institute of Computing Technology, Chinese Academy of Sciences, China

2001.9 ~ 2005.7: B.E. Computer Science

Huazhong University of Science and Technology, China

Research Interest

I am mainly interested in research topics related to data processing and data storage.

- **Database:** Database systems are more and more being challenged by new architectures, like hybrid storage systems. Better adapting to these changes involves many exciting research topics.
- **Emerging storage technology:** Better designs of software are desired to make the most out of new storage devices. On the other hand, such devices have their own drawbacks that need to be addressed. For example, SSDs have the issues of expensive garbage collection and limited life endurance.

Skills

Languages: C, Java, Python, JavaScript

Specific Knowledge: Postgresql, DiskSim (including SSD extension), Linux I/O Module, Hadoop

Work Experiences

2011 June ~ September: Graduate intern at *Intel Labs*, Hillsboro, Oregon.

Research Experiences

2011 ~ now: DBMS and hybrid storage

- **Modify database management system (DBMS) for hybrid storage systems:** (This is an in-progress work.) A storage system is hybrid if it consists of heterogeneous storage devices, such as hard-disk drives (HDDs) and newly emerged solid-state drives (SSDs). HDDs have been the only storage choice for decades, so that DBMSs have been designed without considering hybrid storage systems. In order to fully exploit the capability of heterogeneous devices, we have designed and implemented a heterogeneity-aware software framework for DBMS storage management. In this framework, semantic information related to storage performance is identified and notified to the storage manager where database I/O requests get generated. According to the collected semantic information, requests are classified into

different types. Each type is assigned a proper QoS policy supported by the underlying hybrid storage system, so that every request will be served to its suitable storage device.

- **Differentiated Storage Services:** We propose an I/O classification architecture to close the widening semantic gap between computer systems and storage systems. By classifying I/O, a computer system can request that different classes of data be handled with different storage system policies. Specifically, when a storage system is first initialized, we assign performance policies to predefined classes, such as the filesystem journal. Then, online, we include a classifier with each I/O command (e.g., SCSI), thereby allowing the storage system to enforce the associated policy for each I/O that it receives. This work has been published on SOSP 2011.

2010 ~ 2011: Big data analysis, Storage

- **YSmart:** Database queries are the majority of workloads in data centers. The performance of query executions largely depends on effective execution plan optimizations. Complex queries often take multiple iterations to complete, involving partially duplicate operations such as multiple scans of the same table. By exploiting data and computation commonalities among these iterations, we can reduce the number of iterations and greatly reduce execution time. In this project, we showcased that this technique can effectively speedup (up to 3x) the execution of complex database queries on MapReduce platform. This work has been published on ICDCS 2011.
- **CAFTL:** Flash is an emerging storage technology. It has outstanding random I/O performance, shock resistance and energy efficiency. However, it has a vital shortcoming -- limited life endurance. Inspired by the observation of data duplication, we proposed to address the endurance issue of SSD within the device. By efficiently identifying duplicate writes among incoming requests, we can reduce the volume of write traffic to flash, and save device space. The saved space could be used for more efficient garbage collection and more effective wear leveling. This work has been published on FAST 2011.

2009 TopBT: a topology-aware and infrastructure-independent BitTorrent

- P2P systems like BitTorrent and Skype organize peers in an overlay (logical) network, which has a different topology from the underlying physical network. Such topology mismatch often results in excessive use of Internet bandwidth. We proposed a topology-aware BitTorrent client (TopBT) to find close and equally fast peers, which can save up to 30% of Internet bandwidth without compromising downloading speed. This work has been published on INFOCOM 2010. An open-source TopBT client is available for downloading.

Teaching Experiences

2009 winter Lab Instructor

- CSE 200: "Computer Assisted Problem Solving for Business", given at The Ohio State University. I supervised a lab section of 80 undergraduate students, and managed 4 lab consultants. The job also included face-to-face tutoring, and grading of quizzes and exams.

Publications

- Michael Mesnier, Jason Akers, Feng Chen, Tian Luo, "Differentiated Storage Services", Proceedings of the 23rd ACM Symposium on Operating System Principles (**SOSP'11**), Cascais, Portugal, October 23 - 26, 2011
- Rubao Li, Tian Luo, Yongqiang He, Yin Huai, Xiaodong Zhang, "YSmart: Yet Another SQL-to-MapReduce Translator", Proceedings of the 31st International Conference on Distributed Computing Systems (**ICDCS'11**), 2011. (**Best paper award**)
- Feng Chen, Tian Luo, Xiaodong Zhang, "CAFTL: A Content-Aware Flash Translation Layer Enhancing the Lifespan of Flash Memory based Solid State Drives", Proceedings of the 9th USENIX Conference on File and Storage Technologies (**FAST'11**), 2011.
- Shansi Ren, Enhua Tan, Tian Luo, Songqing Chen, Lei Guo, Xiaodong Zhang, "TopBT: A Topology-Aware and Infrastructure-Independent BitTorrent Client", Proceedings of the 29th IEEE Conference on Computer Communications (**INFOCOM'10**) 2010.