



## Welcome

*The start of school leads to much activity in research and outreach*

The Digital Technology Center Intelligent Storage Consortium (DISC) began the 2008-09 school year by welcoming new students and returning faculty, and supporting a full slate of new and ongoing research projects.

### Faculty return

This year, DISC will benefit from the return of David Du, computer science and engineering professor, who comes back to Minnesota from his two-year assignment at the National Science Foundation (NSF) in Washington, D.C. While Professor Du remained involved with DISC during his NSF tenure, we look forward to seeing more of him, and, as always, value his continued leadership and new research contributions.

Jon Weissman, associate professor of computer science and engineering, also has returned: He spent a year abroad at the University of Edinburgh in Scotland.

In addition, David Lilja, professor and head of electrical and computer engineering; Ahmed Twefik, E.F. Johnson Professor of Electronic Communications and electrical and computer engineering professor; Yongdae Kim and Mohamed Mokbel professors of computer science and engineering will all return to lead DISC research projects this year.

### New: Storage Systems Cooperative Research plans

DISC is involved in an exciting new possible collaboration: Forming an NSF-sponsored Industry/University Cooperative Research Center (I/UCRC) with the University of California Santa Cruz.

Such a cooperative research center offers many benefits that include attracting new industry members and more university partners, and supporting new collaborative research projects within DISC.

DISC members will see very few operational differences between DISC and IUCRC. The DISC Membership Agreement is based on the NSF's. One notable benefit will be shared access to IP from all IUCRC university members. DISC members will have the option of converting their DISC Membership to an IUCRC Membership, if the program is approved by NSF.

The NFS describes the I/UCRC as follows:

*“... The Industry/University Cooperative Research Centers (I/UCRCs) program develops long-term partnerships among industry, academe, and government. ...I/UCRCs stimulate highly leveraged industry/university cooperation by focusing on fundamental research recommended by Industrial Advisory Boards.”*  
(<http://www.nsf.gov/eng/iip/iucrc>)

It will take the better part of a year to organize a new I/UCRC, so changes would not be visible at DISC until next year. We'll continue to update you on the progress of this initiative.



## Members Day

**Oct. 30, 2008**

DISC members will attend the annual DISC Members Day meeting Oct. 30. The event begins the evening of Oct. 29 with a social gathering and continues on Oct. 30 with a full day of research project reviews and discussions.

## Student Expertise

*Student work includes summer research internships and research papers*

Many DISC students spent the summer in very productive ways, such as working at off-campus internships for technology firms and writing papers for conferences this fall.

The following abstracts describe the current research papers of several students, as well as some intern research project summaries. The students plan to complete the papers this fall. For more information on these projects and for the final papers and author information, visit the DISC website at [www.dtc.umn.edu/disc/](http://www.dtc.umn.edu/disc/).

### A Novel Deduplication Approach with Global Knowledge

Guanlin Lu  
([lv@cs.umn.edu](mailto:lv@cs.umn.edu))  
Computer Science & Engineering

We are trying to improve the current chunking based approach with added global knowledge to wisely determine the chunk boundary. As the preliminary experiment indicates, our algorithm outperforms current ones significantly.

### Write-Back Caching Algorithm for SSD

Biplob K. Debnath  
([debnath@ece.umn.edu](mailto:debnath@ece.umn.edu))  
Electrical & Computer Engineering  
IBM Research

The NAND based flash systems have a fast read performance and a slower write performance. This slower write performance can be as much as an order of magnitude lesser than read in certain cases. Every overwrite without any buffering and mapping scheme would require an ‘erase’ and a ‘program’ operation of the related sector. The problem is worsened due to the fact that an ‘erase’ operation only can be performed in a higher granularity (a block level operation) than a ‘program’ (a sector level operation). Therefore, a good buffering scheme can give a remarkable performance gain for any workload that involves write operations. The recent work

related to this area called BPLRU (Block Padding LRU) uses a LRU-based buffering scheme and Log-based FTL algorithm underneath (published in *FAST’08*). But after a thorough analysis, we have come up with an all new buffering algorithm that considers two important aspects affecting the write performance on NAND based SSD—“recent-cy” and “size.” The evaluation methodology of this scheme is based on a simulator implemented for this purpose. The results gathered after running various workloads with different characteristics such as I/O patterns, which are random, sequential, semi-sequential, etc., show good benefit on all real life workload.

### Lazy Updates for the OLTP in Flash-Memory Based DBMS

Biplob K. Debnath  
([debnath@ece.umn.edu](mailto:debnath@ece.umn.edu))  
Electrical and Computer Engineering

Flash memory has rapidly increased in popularity as the primary non-volatile data storage medium in mobile devices because of its small size, low weight, low power consumption, shock resistance, and fast read performance. Moreover, because of these attractive features, flash memory is quickly being adopted for use in enterprise-class servers as onboard cache and solid-state disk drives (SSD). However, the high cost of flash memory and its relatively low write performance require new and innovative solutions to fully incorporate flash-based SSD into high-performance servers. Flash memory is particularly good for the read intensive workload, for example, decision support system (DSS). However, it can produce poor performance when used for Online Transaction Processing (OLTP) workloads because of the excessive number of data updates that these applications require. This paper aims to overcome this drawback of flash memory for OLTP systems by designing an efficient scheme for handling data updates in the flash memory. In particular, we will

explore a lazy update methodology that significantly improves the update processing overhead for OLTP. The cost of achieving such improvements is only few flash memory blocks. Our lazy update methodology not only enhances the write performance, but also increases the reliability of the flash memory by reducing the number of block erases. Furthermore, we propose a new query processing technique that complements the lazy update methodology to guarantee the correctness of the query result.

### Computing Multi-Objective Queries

Mohamed Khalefa ([khalefa@cs.umn.edu](mailto:khalefa@cs.umn.edu))  
Computer Science & Engineering  
Microsoft Research

Preference-aware queries return objects that are most interesting of the user. We aim to provide a scalable personalized system to users, that is aware of users’ preferences. Modern applications (e.g., biological data, GPS application, and sensor monitoring applications) deal with real-world data, therefore they should capture, model and query uncertainty associated with real-world data, to produce “accurate” and “representative” results. Various factors (including temporal and spatial) affect the span of uncertainty region for a specific object. Therefore current work addresses efficiently computing the skyline and top-k queries, which are special type of preference-aware queries, over uncertain data. The primary results show significant improvement over current approaches.

### D3: Demystifying Data Deduplication

Pramod Mandagere  
([npramod@cs.umn.edu](mailto:npramod@cs.umn.edu))  
Computer Science & Engineering  
IBM research

The effectiveness and tradeoffs of deduplication technologies are not well understood—vendors tout deduplication as a “silver bullet” that can help any

enterprise optimize its deployed storage capacity. This paper aims to provide a comprehensive taxonomy and experimental evaluation using real-world data to evaluate the origins of duplication in data, and to correlate the space savings with data types, along with the overheads involved in CPU resource, ingestion rate, and reconstruction time. Our experimental results show that for different data, inherent deduplication rates can vary by as much as 300 percent.

## Data Center Power Management

**Nagapramod Mandagere**  
 (npramod@cs.umn.edu)  
 Computer Science & Engineering  
 IBM Research

Enterprise and personal computing have undergone a rapid shift toward internet/web-based service oriented architectures. Consolidation of resources into large unified data centers has been a key enabler of this trend. A key technology, such as virtualization of resources, has increased the pace of consolidation with obvious benefits of lower management costs. Data centers have not only reduced overall cost of management, but also have led to ease of deployment of complex software solutions. Data centers can be seen as concentrated clusters of computing and data storage resources. Various surveys and reports commissioned by both the government and industry estimate that the total cost of powering the nation's data centers in 2007 was about \$5 billion, a figure that is expected to rise to about \$7.5 billion by 2011. Though this power budget for data centers accounts for only 1.5 percent of overall electricity consumption in the United States, it is the concentrated nature of the power demand that is a cause for major concern. Consequently, data centers can be viewed as stress points on the national power grid. Understanding the nature of

power usage in a data center is the key first step in any approach to improve power efficiency.

## Reliable Data Deduplication Performance

**Nohhyun Park**  
 (parkx408@umn.edu)  
 Electrical & Computer Engineering  
 LSI Advanced Development

Data deduplication is often defined by well-known design choices, such as indexing scheme, chunking method, etc. However, data that is being deduplicated, as well as HW that is running the deduplication process, also play a vital role. This research hopes to find correlation between these parameters and deduplication performance both in speed and compression ratio. A simulator is built to test various design param-

eters on different set of data distribution and dynamics. We hope that this research will enable data deduplication to deliver predictable and reliable performance over wider range of data environment.

The generic x86 processor that is used today for data deduplication consumes a lot of power and is relatively expensive. We believe using HW support of data deduplication could lower this cost extensively while offering equal performance. The difficulty is that some processes within the data path can be very slow, and a complex control mechanism is needed to handle different cases. We hope to use deep packet inspection hardware to handle these different cases while optimizing the data path of most frequent cases.

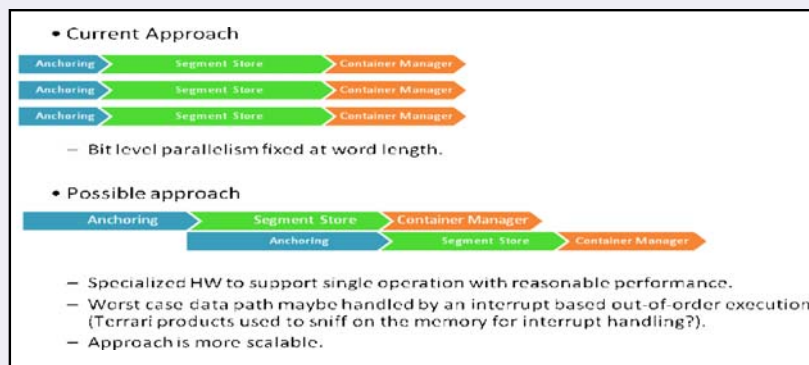


Fig. 1 Current approach versus specialized HW approach ideas

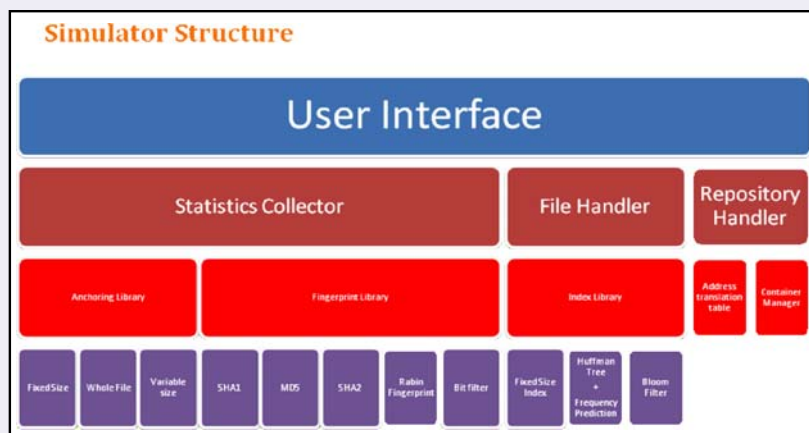


Fig. 2 Simulator to test design parameters

# Security Improvement

*Researchers explore integrated infrastructure for secure and efficient long-term data management*

**D**avid Lilja, professor and head of electrical and computer engineering, and Yongdae Kim, associate professor of computer science and engineering, are leading a research project on integrated infrastructure for secure and efficient long-term data management. The National Science Foundation provides funding support for the research.

## Theoretical foundation of authenticated file encryption

In a cryptographic file system, a user who has correct attributes for a file will eventually obtain the file encryption key for the file. With the file encryption key, a file can be encrypted or decrypted using techniques of symmetric cryptography.

As part of this project, researchers are investigating new techniques for encrypting files to provide both integrity and privacy. Although previous literature on this topic exists, the current state-of-the-art is still far from complete. In this area, the research goal is to provide a rigorous security definition and a provably secure scheme, based on symmetric key cryptography primitives like universal hash functions and authenticated encryption schemes. Researchers completed formalizing the notion of security for such schemes,



David Lilja



Yongdae Kim

produced a framework of a scheme with security proof, and are currently working on the details of the scheme.

## Hierarchical Attribute Based Encryption (HABE)

Identity Based Encryption (IBE), first realized by Boneh and Franklin, inspired further research and extension of the notion.

One such extension is Hierarchical Identity Based Encryption (HIBE), which reduces the workload of the central server, and also reflects the fact that most organizations are hierarchically structured. In other direction, Attribute Based Encryption (ABE) lets the user to encrypt with respect to a set of attributes, instead of an identity, thus enabling more fine-grained access control than IBE.

In any public-key scheme, revocation is a big challenge. This is especially true for IBE and related schemes. One solution involves including epoch or time information as a part of identity (or attribute), and encrypt with respect also to the epoch. Epoch-based revocation

was suggested for IBE, but when applied to storage security, a user has to keep all the older private keys to maintain access to older files. Also it is not clear how this can be done for HIBE. Meanwhile, due to the expressibility of ABE, it is possible to form inequalities about epochs, therefore a single policy can match many epochs.

To provide efficient revocation capability to HIBE, researchers are working on an efficient HABE scheme, a natural extension of both HIBE and ABE, which would enable researchers to incorporate the epoch information into HIBE, by combining a hierarchical identity and an epoch using logical connectives.



## For more information

Questions about a project at DISC?  
Interested in exploring a topic?  
Wondering about working with a DISC  
faculty member?

For information about DISC projects or  
membership, please contact:

Cory Devor, Ph.D.  
Director of Business and Research  
Development  
Digital Technology Center  
499 Walter Library  
612-625-1716 (o)  
952-484-8550 (c)  
devor@dtc.umn.edu

[www.dtc.umn.edu/disc/](http://www.dtc.umn.edu/disc/)

*The University of Minnesota is an equal  
opportunity educator and employer.*