CMRR Report Number 21 Summer 2003

Center for Magnetic Recording Research

CMRR's Project on Intelligent Disk Drive Storage Systems ("iStor") and the "Secure Erasure" Project

Gordon Hughes, *University of California, San Diego*

CMRR's project on intelligent disk drive storage systems ("iStor") is sponsored by the Sloan Foundation Information Storage Industry Center at UCSD. Our initial project study was CMRR SMART, which developed disk drive failure warning methods allowing four times higher warning accuracy than current disk drive technology, at low false alarm rates (see "Improved disk drive failure warnings" *IEEE Trans. Reliability*, Vol. 51, p 350-7, September 2002)

A paper on SMART by ECE graduate student researcher Joe Murray has been accepted for presentation at the 2003 International Conference on Artificial Neural Networks. It compares past CMRR SMART warning methods to pattern recognition support vector methods, to unsupervised clustering, and to reverse arrangements statistical rank tests. The latter test looks for upward trends in drive internal error counts, by tracking the number of times a new error count is higher than previous counts.

Intelligent storage features offer benefits in storage system performance and to user application programs. A recent *IEEE Spectrum* article by Dr. Gordon Hughes' ("Wise Drives," August 2002) points out that work on intelligent storage dates from the 1980's (such as the Teradata data processing computer and Active Disks), pushing data-centric computing downward in computer architecture, but rarely penetrating through the computer-disk drive interface. CMRR currently has three active projects based on the potential iStor features listed in this *IEEE Spectrum* article.

Intelligent storage system ("iStor") features are enabled by changes in internal drive technology and in computer data access methodology. We see extensions of object-based storage devices (OSD in the SCSI architecture) as a basis for storage systems and user application programs to interface to iStor features. Current abstractions such as storage virtualization miss iStor opportunities because they use drives simply as block data storage

devices. OSD allows abstraction of storage away from physical device management, for virtualization and for real-time quality-of-service management. This can allow tight control of physical storage performance by data intensive user applications such as database programs.

A computer system has been configured to experimentally test the disk drive secure erase "SE" intelligent feature that CMRR put into the standard drive computer interface specs, SCSI and ATA. We are developing a CMRR SE test protocol to verify that SE deletes all user data from test drives, beyond the possibility of recovery.

A third iStor project has begun to study storage system performance improvements by using the "time-to-data" feature from the *IEEE Spectrum* article. A computer simulation will determine data access speedups made possible by letting storage systems query drives to find the quickest order to access records in a multi-task user request queue. An experimental project is in progress to study time-to-data by using existing SCSI commands which return the head and disk physical positions at any time.

CMRR is also participating in a Storage Networking Industry Association (snia.org) initiative to form a storage network users interest group at ISIC.

"Secure Erasure" Project

CMRR's storage systems project on "secure erasure" of disk drive data has just been renewed by the National Security Agency, with Gordon Hughes as Principal Investigator. This three-year $150,000 project answers a significant data storage user need to reliably eradicate data from computer hard drives for security and privacy reasons. The need for "secure erasure" arises when:

- A user releases disks or they are removed from systems for maintenance.
- Storage devices are re-configured for other uses or users, for instance in expiring leased data storage facilities at a storage service provider or data center.
- A project is completed and the data must be purged to protect "need to know" or to prepare drives for new users or applications.
- A virus has been detected and all possible traces of the offending code must be eliminated.

The "secure erase" (SE) command is now part of the standard disk drive specifications (IDE/ATA and SCSI) at CMRR's request. It is a positive, easy-to-use data destroy command, amounting to "electronic data shredding." It completely erases all possible user data areas by overwriting. SE is a simple addition to the existing "format drive" command currently present in computer operating systems and storage systems, and consequently adds little or no cost to drives.

Secure erase is required by the ATA specification, although it is optional in SCSI. The new "serial ATA" drives will be able to advertise SE as a user feature in their competition with SCSI and fiber channel drives for market share in low-end storage systems.

The new secure erase project is testing a CMRR protocol for evaluating the SE feature in individual drive models to insure that erased data is not recoverable.

Secure erasure capability will be required by the U.S. government for their disk drive purchases. Considering the security feature this capability offers to many users we expect there will be considerable commercial interest in this capability as well.

For additional information on the secure erasure initiative, contact Gordon Hughes gfhughes@ucsd.edu