

# Intelligent Disk Drives

---

**Gordon Hughes**

**Associate Director**

**Center for Magnetic Recording Research**

**University of California, San Diego**

**A university center for data storage technology,  
Founded 1983, sponsored by disk and tape companies.  
...Five professors, 40 students, dedicated building**

# Talk Summary

---

- **What *are* iDrives?**
- **What are iDrive tasks?**
- **How to get tasks into drives**  
Who is the customer, the technology user?
- **CMRR's iDrive Projects**  
Drive failure prediction  
Secure Erase
- **CMRR New iDrive Projects**

# Talk Summary (2)

---

- **List current and future intelligent tasks**
- **Discuss some in a little detail**
  - Large computer systems for database searching
- **Storage intelligence stops near drives**
  - Interface specs are a barrier to in-drive tasks
- **Interface specs (SCSI, FC, ATA...)**
  - How new tasks get in the specs
  - Flash memory: an example new ATA task
  - New tasks don't *have* to be put in the specs

# Intelligent Disk Drives

## Gordon Hughes, CMRR

---

- **CMRR Storage Systems projects:**
  1. Disk drive failure prediction (SMART)
  2. Secure erase (drive “data destruction”)
- **Both are examples of “intelligent drives”**
  - ...Tasks beyond data storage and retrieval
- **Other possible intelligent drive tasks?**
  - Past research work stops at drive interface
  - Computer side of SCSI, Fiber Channel, ATA,...

# What are Intelligent Disk Drives?

---

- **Disk drives have computing power**  
comparable to PCs a few years old  
32-bit  $\mu$ Ps and MBs of “free” storage  
DOS-like (proprietary) operating systems  
And with the combined power of *many* drives...
- **Computing power allows new drive tasks**
- **May allow product differentiation**
- **Interface standards and “food chain”**  
are hiding this potential

# iDrive Tasks

In drives now:now CMRR Research Project:cmrr

---

- **Data storage, retrieval and integrity** now
- **In-drive RAID-5 parity storing** now
- **In-drive failure prediction** now, cmrr
- **In-drive all-database mining**
- **Secure Erase (“data destruction”)** now? cmrr
- **In-drive file encryption or compression**
- **Support tasks for data storage networks**  
Network data error/corruption control, backup
- **Disk drive quality of service (QoS)**

# **In-Drive RAID Error Control** now

---

- **RAID-5 stripes user data over several drives**
- **Data ECC XOR stored on parity drive**
- **So writing one data sector takes 4 accesses**  
Read old data sector, old parity sector,  
XOR these with new data sector,  
Store new data sector and new parity sector.
- **SCSI spec permits in-drive XOR today**  
Reduces four accesses to two, handled by drives
- **Autonomous drive networking already here!**  
But only between drives on single SCSI/FC loops

# Drive Internal Failure Prediction cmrr

## (CMRR Research Project)

---

- **“SMART” is in ATA and SCSI specs**  
Introduced by Compaq in 1994  
SMART warns user of impending drive failure
- **Exploited existing drive error logging**  
Logging has been in drives for many decades  
Used by drive reliability engineers
- **CMRR’s SMART research program**  
Problem: high accuracy can mean high false alarms  
Simple statistical predictors work very well  
See [cmrr.ucsd.edu/smart](http://cmrr.ucsd.edu/smart)

# In-Drive File Encryption

---

- **Encrypting stored data becoming common**
- **Software encryption by computer O/S**  
Microsoft Windows 2000 and XP have it
- **Drive encryption would be far more secure**  
Hardware is more secure than software  
Encryption key stored in user inaccessible area  
Protection against brute-force cracking attacks
- **AEA  $\approx$  binary multiply: Coded=data<sup>e</sup> mod n**

# **Secure Erase** now, cmrr (CMRR Research Project)

---

- **New drive feature: “data destruction”**  
Requested by U.S. Government  
(National Security Agency)
- **It’s in the ATA and SCSI specs**  
From earlier work by CMRR
- **CMRR developed a drive erase protocol**  
Erase all user accessible areas (G&P lists too)  
Validation tests - track edge data vs. noise. (Insight)
- **In-Drive file encryption would speed SE up**

# Large Database Computer Systems

---

- **Big mainframes, tight data storage control**
- **Parallel searches of whole databases (mining)**
- **Specialized computers proposed, since 1980's**  
Teradata's parallel  $\mu$ CPU's, RAM and drives:  
Specialization *stopped* at drive interface (SMD)
- **Today, no specialized CPU's or drives**  
“Use fast-cheap commodity drives” -Jim Gray, MSFT  
Intelligent drive research also stops at interface

# RAID fiber channel Storage Networks

---

- **FC allows 127 drives per two-way loop**  
Whole-SAN/NAS mining search might get only 2/127  
 $\approx 2\%$  of total data transfer rate (all drives)?
- **Intelligent drive data mining searches**  
Send search criteria to drives,  
Drives independently search, report matches
- **Issue: non-file-aware searches, RAID striping**  
File-aware SCSI object oriented storage (OOS)

# Disk Drive Quality of Service (QoS)

---

- **Disk drive data access times are uncertain**  
Variable seek and latency delays  
Unannounced delays in drive error recovery mode
- **For video streaming, late data is useless**
- **Networks need data delay and error QoS**
- **Drive QoS for network data error control**  
Drives now provide logic-level interface CRC  
Usable by network for data corruption detection

# The Interface Specs:

## ...Both Opportunity and Barrier

---

- **Getting new SCSI, ATA, FC tasks:**

Major drive makers/users need to agree on specs.  
It works: flash memory requirements are in ATA  
Intelligent drive tasks can be added same way
- **Proprietary tasks also possible**

Specs have “Vendor reserved commands”  
A drive maker could introduce a feature,  
Gain competitive advantage.  
(Popular tasks may be put into standards)

# Implementation

- **iDrive changes cross many interfaces**  
Computer/storage network, drive, drive's CPU code
- **Technology food chain problems**  
Specialist companies influence the interfaces  
Hey! That's *my* value added!
- **Object based storage as “carrier vehicle”**  
Storage Networking Industry Association's OSD  
[http://www.snia.org/tech\\_activities/workgroups/osd/](http://www.snia.org/tech_activities/workgroups/osd/)
- **CMRR projects may help**  
RAID itself started in a university...  
Patterson, Katz & Gibson at UC Berkeley

# CMRR New iDrive Projects

---

- **Research programs**

- Drive network error correction

- Drive tests using SCSI read long

- Data delay QoS

- Storage array simulation

- Compare network access queuing to in-drive queuing

- **Publicize**

- Technology articles

- White papers

# Conclusions

---

- **Present & potential iDrive tasks listed**
- **They may allow product differentiation**  
Today, PC specs don't even *list* the drive maker!  
They *do* list SMART-3, an iDrive feature...
- **Opportunities for tape drive makers too**  
Do SMART first (tape drives have error logs too)
- **Comments and suggestions are requested**  
858-534-5317, [gfhughes@ucsd.edu](mailto:gfhughes@ucsd.edu)