



## From security to reliability

RSA labs explores new strategies for detecting disk errors.

BY SARAH JENSEN

**AS THE DIGITAL** universe continues to expand, the reliability of hard drives becomes ever more imperative. According to a 2008 EMC-sponsored IDC study, digital information is expected to double in size every 18 months in the near term, quintupling by 2012. That surge in data will necessitate increased reliance on hard drives not just for storage, but as components for large-scale archival and backup systems. RSA labs re-

searchers Alina Oprea and Ari Juels are exploring ways to ensure hard drive reliability by more quickly detecting disk errors that could compromise compliance and data availability. “Our aim is to find the best strategy for detecting latent sector errors (LSEs), errors that aren’t discovered until you attempt to read a sector,” says Oprea. After analyzing existing disk error and failure data, she and Juels have derived a new simula-

tion model for LSE incidence and an improved algorithm with the potential to optimally find those errors before they become problematic.

**TO ERR IS CYBER** Enterprise drives, bound by compliance requirements, normally develop fewer errors than home computers and PCs, but any drive can develop errors for a variety of reasons both age- and usage-related. Small particles may become lodged in the media, or fre-



*“We believe this work will lead to better-designed algorithms that will significantly improve existing technologies.”*

**ALINA OPREA**

quently accessed areas of a drive may simply wear out and become corrupted.

“If you’ve ever seen the message ‘bad sector’ when you’re trying to open a file or perform an operation on the drive, you’ve encountered a sector error,” explains Oprea. “If a redundancy mechanism isn’t employed and the errors aren’t detected by the system, you’ll find the data simply isn’t there when you try to access it.”

Enterprises typically employ RAID mechanisms (Redundant Array of Independent Disks), subsystems that provide fault tolerance; if an error is detected in time,

redundant additional drives step in to correct it. Even with a RAID-5 mechanism in place, however, an LSE coupled with a drive failure results in data loss.

#### **SCRUBBING UP**

Most current systems perform sequential disk scrubbing—checking disk sectors sequentially at a constant rate—to detect LSEs in a timely fashion. But RSA is taking a different approach, a “staggered adaptive strategy” (SAS).

“We describe our strategy as ‘staggered’ because it doesn’t scrub the disk sequentially,” Oprea explains. “Instead, our method samples a few sectors in a disk region

and then ‘staggers’ to a new disk region. LSEs tend to arise in clusters, and this repeated sampling of a region identifies them faster than the sequential process.”

“Adaptive” refers to varying the scrubbing rate—the number of sectors scrubbed over a given time period. By combining the techniques of staggering and adaptively changing scrubbing rates, Oprea and Juels developed an SAS prototype, now in use in the lab, which opens the door to a number of future investigations.

“The next steps will be determining the optimal scrubbing strategies for particular disk configurations,

including RAID configurations,” she says. She also envisions exploring ways SAS can automatically invoke recovery applications when errors are detected, and translating SAS results into the Flash realm.

SAS represents not only a step forward in the area of disk scrubbing, but a step beyond RSA’s traditional concentration on security. “This moves us into the areas of reliability and the improved availability of information,” says Oprea. “We believe this work will lead to better-designed algorithms that will significantly improve existing technologies.” ■