Pushing the Data Storage Envelope

With magnetic storage devices rapidly running out of room, the coming generation of computers may depend on optical and atomic force microscopes—if they come up to speed

To catch a glimpse of the future of data storage, says Dan Rugar, a physicist at IBM's Almaden Research Center in California, you will need a powerful microscope. For that future is vanishingly small, and microscopes may play a leading role in it.

The present, in contrast, is big. Too big. The magnetic hard disks and optical CD-ROM drives now humming along in your desktop or mainframe computer store a single unit of information—a bit—in a small magnetized region on thin metallic film, or as a pit molded into a plastic CD-ROM. These bits can have an area of 1 square millionth of a meter. That may not sound like much, but space-hungry computer programs and databases call for billions more of these bits each year—and no matter how hard they squeeze, engineers are unsure they'll be able to get more than 10 billion of these bits onto a square inch of disk (inches are the standard units for describing data density). "There are some pretty daunting challenges to increasing the storage density," says Gary Gibson, a physicist at Hewlett-Packard Laboratories in Palo Alto, California.

Ultrasensitive microscopes—tools that at their sharpest resolution can image the inner workings of cells and even individual atoms—may be up to that challenge. At the Optical Data Storage conference 2 weeks ago in San Diego, rival teams of researchers, one from Almaden and the other from Stanford University and Sony Corp., reported combining elements from a high-resolution optical microscope and a high-speed magnetic disc drive to make bits as small as 310 nanometers (billionths of a meter) on a side.

But most importantly, one prototype could read and write data at a rate of 8 million bits per second—faster than a conventional CD-ROM drive. That's a giant leap forward, because while microscopic techniques have excelled in data shrinkage, they've also been woefully slow. But the San Diego presentations show that the rates "are certainly picking up steam," says Ken Babcock, a physicist and microscope specialist at Digital Instruments in Santa Barbara, California. These faster speeds give microscopes "great potential" to find their way into the next generation of data storage devices, adds Horst Stormer, AT&T Bell Laboratory's director of physical research.

That potential—and the hope of carving out even a sliver of the $42 billion worldwide data storage market—has prodded commercial giants such as IBM, Sony, Matsushita, and Hewlett-Packard to attempt to speed up other microscopes. Preliminary successes with the atomic force microscope (AFM)—the atom-imager—are starting to trickle in, with ever smaller bit sizes and faster speeds. But these novel storage devices still have to surmount formidable technical obstacles, such as reducing their bulky size and reducing their cost, which in the case of the AFM currently exceeds $50,000 each.

Conventional hard disk drives write bits by placing a magnetic head over a thin film; the magnetic field from the head rapidly reverses, reorienting the field of a small region of this film, and the transition between regions is read as a bit. In making their devices, both the IBM and Sony-Stanford teams used a variation of this technique, but one involving a near-field optical microscope originally developed by Stanford University electrical engineer and team leader Gordon Kino. The microscope uses a special hemisphere-shaped lens, known as a solid immersion lens (SIL), to focus infrared light from a semiconductor laser beam to an ultrasharp point at the lens bottom. Most of this light is reflected back, but a thin sliver of photons can "tunnel" for up to 150 nanometers past the lens.

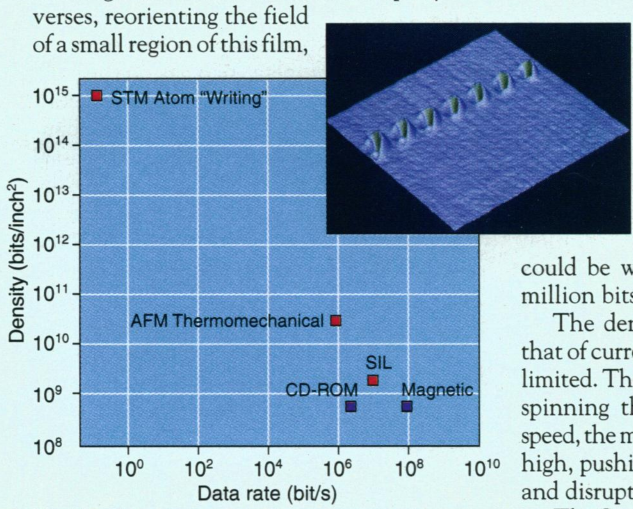

Packing it in. Computer data storage techniques using microscopes *(red)* are advancing in speed and density. Data bits created by an atomic force microscope (AFM) form craters *(inset)*, which are just 300 nanometers across.

If a magneto-optic disk is within those 150 nanometers, these tunneling photons then heat a small spot on it, causing it to lose its previous magnetic orientation. As the spot cools, it then takes on the magnetic orientation of a surrounding field generated by a nearby electromagnet, while neighboring spots remain unchanged. And that transition is seen as a distinct bit. It can be read as distinct by a laser as well, because the changed magnetic orientation affects the polarization of the light reflected back from the surface.

But all this has to be done at high speed. Rugar and his IBM colleagues Bruce Terris and John Mamin borrowed a magnetic recording technique whereby the head—connected to a mechanical arm like the arm on an old phonograph—"flies" like an airplane wing on an air cushion produced by a disk spinning at 420 revolutions per minute. That allowed the head to stay close to the disk's surface and produced a data density of 2.5 billion bits—or gigabits—per square inch, which could be written and read at a rate of 3.3 million bits per second.

The density was four times higher than that of current devices, but the speed was still limited. The air cushion prevented them from spinning the disk any faster: Beyond that speed, the moving air lifted the optical head too high, pushing the SIL away from the surface and disrupting the tunneling mechanism.

The Sony-Stanford team, led by Kino and Sony's Isao Ichimura, got around the problem by using a different strategy to create bits. They reshaped the SIL to allow more light through from the laser, so they didn't have to rely on tunneling photons. That meant they could move the head further from the disk and spin it faster. The result: a reading and writing rate of 8 million bits per second. But moving the SIL back also increased the spot size. So to compensate, they turned to a shorter wavelength green-light laser to produce their light. A shorter wavelength means a smaller spot. That gave them a data density of 4 gigabits per square inch in addition to a fast microscope.

But nobody expects the effort to speed up microscopes will stop here. The IBM researchers and other teams are working to boost the speeds of other microscopes, such as the AFM. With tips as small as 10 nanometers across and the ability to manipulate matter on that scale, AFM probes can manufacture smaller bits than SILs can. Over the last few years, for example, Mamin and Rugar have developed a scheme for writing data by rapidly heating and cooling the tip of an AFM probe, melting 100-nanometer-wide pits in a plastic disk, a size that corresponds to a data density of 30 gigabits per square inch.

The holdup here is the speed. In the past, AFM data writing and reading techniques have had a top speed of 0.1 million bits per second. But the reading speed is accelerating. To read bits, the disk is spun under the ultrasharp AFM tip, which sits on the end of a flexible cantilever arm. A laser beam is then directed onto the surface of the cantilever. And as the probe tip moves in and out of the pits, the cantilever arm vibrates, changing the direction of the reflected light, which is detected by a nearby sensor. Each change in the light pattern signifies a bit. This reading speed is limited, however, by the frequency at which the cantilever arm is able to vibrate—100,000 times per second for standard AFM cantilevers.

Last year, however, the researchers reported developing smaller cantilevers—which vibrate at higher rates—and in a prototype used them to read AFM-coded data at 1.2 million bits per second. Rugar believes they should be able to boost the reading speed further by going to yet smaller cantilevers.

Increasing the writing speed beyond the current limit "is a tougher problem," he says, because the tip must cool down before it can move to the next bit; otherwise it would melt a continuous groove in the plastic substrate. To produce a faster pit writer, "we'll probably have to make parallel arrays of tips," says Rugar. In other words: If you can't move the data faster under the tip, multiply the tips that produce the data. The basic idea is to build arrays of independently maneuverable probes that can work in parallel. While researchers are making progress in fabricating such arrays, they have yet to use such systems to demonstrate the reading and writing of data at high speeds. "We're about 90% there," says Noel MacDonald, an array builder and professor of electrical engineering at Cornell University in Ithaca, New York.

That last 10% may be tough. Writing or reading with more than one tip is complicated by the fact that recording surfaces are rarely flat. And if a series of tips is all connected to the same platform, one tip may be in contact with the surface while its neighbor is not. So the array-builders are striving to build tiny motors alongside each AFM to control its height from the surface. If they succeed, they may buy the data recording industry more speed as well as space.

–Robert F. Service

ASTRONOMY

How Quasars Make Heavy Metal

Quasars manage to stand out, even among their fellow cosmic glitterati. These distant light sources, each of which flooded the early universe with the energy of hundreds of billions of suns, have plenty of mysteries to go with their sheer brilliance. One mystery in particular has kept astrophysicists watching their journals as eagerly as show-business fans follow *Variety*: the high concentration of metals—astronomers' lingo for elements heavier than hydrogen, deuterium, and helium—evident in quasars' spectra. The big bang created only the lightest nuclei; the quasar metals, tens of times more abundant than in our own sun, must have been forged soon afterward, in the cores of stars. But astronomers could only speculate about where these early metalworks were located and how they supplied quasars with their products.

Now a group of astrophysicists has come up with a new scenario in which metal formation takes place within quasars themselves. The metals, say Pawel Artymowicz of Stockholm University in Sweden and his colleagues, come from stars trapped in the "accretion disk" of glowing gas that surrounds the center of each quasar, where a supermassive black hole is thought to lurk. At a meeting of the American Astronomical Society in Pittsburgh last month, the researchers proposed that these trapped stars devour material from the accretion disk, grow to massive proportions, and turn into prolific metal sources.

The theory, which quasar researchers like David Turnshek of the University of Pittsburgh consider "exotic" but intriguing, follows earlier theories in relying on extremely massive stars to forge metals. Such stars, perhaps 100 times heavier than the sun, live out their lives quickly, then explode as supernovas and spew out metal-rich gas. In the past, theorists puzzling over the quasar metals have assumed that the massive stars populated the extensive galaxies in which quasars were assumed to reside. Proposed in 1992 by Fred Hamann of the University of California, San Diego (UCSD), and Gary Ferland of the University of Kentucky, this scenario gained support from observations of large galaxies near Earth, whose cores are rich in metals.

But this picture got a jolt from recent

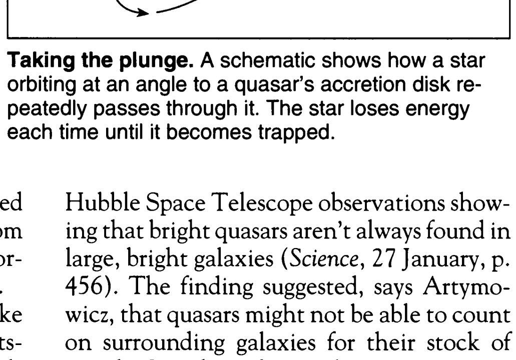

SOURCE: PAWEL ARTYMOWICZ

Taking the plunge. A schematic shows how a star orbiting at an angle to a quasar's accretion disk repeatedly passes through it. The star loses energy each time until it becomes trapped.

Hubble Space Telescope observations showing that bright quasars aren't always found in large, bright galaxies (*Science*, 27 January, p. 456). The finding suggested, says Artymowicz, that quasars might not be able to count on surrounding galaxies for their stock of metals. Searching for an alternative source, he and his colleagues Douglas Lin of the University of California, Santa Cruz, and Joseph Wampler of the European Southern Observatory in Garching, Germany, focused on the swarms of smaller stars, orbiting near the accretion disk, that astronomers have observed. Large numbers of these stars, the group realized, could easily get caught in the disk. Twice per orbit, each of these stars would punch through the disk and lose energy to it, eventually becoming trapped. The stars would then grow "very rapidly" by drawing in dense gases from the disk, says Artymowicz. Finally they would become massive enough to evolve quickly and explode as supernovas, spreading metals at roughly the right rate to match the observations.

"It's an intriguing vision," says Gregory Shields of the University of Texas, Austin. And, he says, it's one that is "likely to have a number of ramifications beyond the question of chemical abundances." Specifically, he says, the supernovas could help answer the question of why the disk gases, instead of circling the black hole indefinitely like Saturn's rings, eventually fall into the black hole, generating the quasar's brilliant light. The secret, says Shields, might be the "stirring" provided by the supernova explosions.

UCSD's Hamann, however, isn't ready to agree that the Hubble observations rule out the galactic scenario for metal formation. Even if large galaxies are missing from the quasar neighborhood, smaller galaxies, invisible to the Hubble, could still be forging at least some of the heavy elements, he says. "If [Artymowicz's] scenario is in some sense correct, I would argue that both [processes] are going on," says Hamann. Only very detailed measurements of the metal abundances could show which is the bigger contributor, he says. And that means quasars aren't likely to lose their mystique anytime soon.

–James Glanz