The Technology of Data Stashing

Neville Holmes, University of Tasmania

Computers only exist to process data.

For some time I have been bemused by the proliferation of names used for the flash memory devices that have consigned diskettes to history.

I have read of flash, jump, key, key chain, pen, solid state, thumb, and USB drives; of data, flash, and memory keys; of flash and USB plugs; and of data, flash, intelligent, and memory sticks. Wikipedia cites 50 different names from English usage and another 30 from other languages (en.wikipedia.org/wiki/USB_flash_drive). Use of the term "drive" here is unfortunate. Just as cars had drives, so did disk packs and diskettes. A drive is something that moves what it is attached to.

Whatever you call them, and I favor data plug, they, and their cousins hidden within devices like iPods and mobile phones, are widely used indeed.

That may be why, while recently serving as a judge for a technical award committee, the entry for flash memory jumped out at me when I read the initial list of nominations.

I voted for flash memory for two reasons. First, it provides a basic and distinct technology for storing data. Second, it exerts a great effect on the culture of data storage, a technology more important than that of computers. After all, computers only exist to process data. Data stashing—the storage of data for computers but external to them—has a long and fascinating history that brings out the significance of the data plug and flash memory.

CARDS

My first experience with data stashing involved the medium of punched cards. One card usually held the data for a single transaction, and users called this technology unit "record." Files held thousands of cards, even hundreds of thousands, and were processed through a variety of electromechanical machinery at several different rates.

These files could be speedily copied, resequenced, merged, and separated, and the machines' operators could see and control what was going on during any process. Individual cards could even be printed on and used as documents. However, the file cards had limited capacity, fewer than a hundred characters, which made processing more complicated if the unit record didn't fit on a single card. Also, the cards had a tendency to jam, particularly in high-speed sorters. Repairing a mangled card could be difficult and time-consuming.

Developers overcame the card-capacity problem in the 1960s by recording the data on a magnetic strip, like that on present-day credit cards, or on a magnetic card, although this solution applied mainly to electric typewriters. Nevertheless, the larger magnetic-stripe ledger cards that NCR and Burroughs introduced in the 1950s remained popular well into the next decade.

Mass data storage provided a rather curious use for magnetic cards on, for example, the NCR 315 CRAM and the IBM 2321. Such machines had canisters of cards mounted inside a machine that could select individual cards, wrap them around a drum so that data could be speedily transferred between card and computer, then put them back in their canister. While successful, these machines had wear problems that led to the technique being dropped. One interesting unit-record technology of the 1960s that survived, the MICR check, features check data in magnetic ink that can be read by both people and machines.

TAPES

In the early days of electronic computing, scientists preferred punched paper tape to punched cards. The advantage, they claimed, was that a dropped roll of paper tape could be rolled up again, as opposed to a dropped tray of cards, which lost their sequence. The commercial world sometimes used paper tape because it was cheaper and often faster for data capture, but usually it was simply transcribed to cards or magnetic tape for further processing.

Magnetic tape was widely used for data stashing, but inherently for files of data rather than for accessible unit records. If, say, one record in a customer master file on magnetic tape needed updating, the entire file had to be transcribed. As long as data processing was batched, this was not a great disadvantage. Resequencing a file was awkward, though, because it involved repeated splitting and

Continued on page 118

Continued from page 120

remerging of the entire file, using three tape drives at the very least.

With magnetic tape, small files were wasteful because only longish reels were available for the common and expensive half-inch tape. Although a programmer could write many separate files on a reel, updating one involved rewriting the entire reel.

The open-reel DECtape, popular with minicomputers, got around this difficulty by having the tape formatted in data blocks that could be read or written individually. The quarter-inch tape cartridge, developed by 3M in 1975 for IBM's 5100 PC, let the user choose formatting sizes for files on the tape. A belt running on its surface drove the tape itself, a feature that got around the problem of variable speed in capstan-driven cassettes. This tape only had a 204-Kbyte capacity, but when bit densities increased greatly, the design became popular for the streaming tapes used for backup. Early hobby PCs used a drive to handle data on ordinary audiocassette tapes.

Tape also had the disadvantage of being limited to serial and usually unidirectional access. The Datafile on Burroughs' Datatron computer of the late 1950s got around this by having 50 200-foot loops of tape and a moveable read/write head.

IBM at one stage planned to announce a tape loop device with its System/360. This device had chambers for the loops around the periphery of the space between two flat metal plates. A selected loop would be brought from its chamber and inflated in a waisted shape, like a figure eight, between a capstan and a read/write head. Pneumatics achieved all this and the return of the loop to its chamber. IBM dropped this device in favor of a similar one using simpler magnetic hoops, which was in turn dropped, reportedly because the tape persisted in crimping its edges.

DISKS

The first disk data-storage devices I encountered, the IBM 350 and the later 1405, were physically large, integrated with the host computer, and not really designed for stashing data. For data stashing, the company introduced the IBM Disk Pack in 1962. Weighing 10 pounds, the pack could be attached to or removed from the 1311 disk drive. It had six disks and held 2 million six-bit characters. Unlike the 350, the drive had a comblike access arm with a read/write head for each surface.

The widespread adoption of flash-memory devices brings both problems and possibilities.

A succession of disk packs followed, with ever-increasing bit and track density and decreasing head clearance, which made their reliability dependent on clean surroundings. Ultimately, a disk pack code-named Winchester introduced in 1973 for the IBM 3340 unit was a hermetically sealed unit complete with its own access arm, which made it rather heavy. After that, designers sealed high-capacity data storage disks in with their drives.

While designers used a variety of disk packs in the 1970s for various purposes, the less ponderous data stashing started with the 200-mm diskette, originally called the Minnow. Also from IBM, although originally a Japanese invention, the diskette arrived near the end of the 1960s, during the transition from System/360 computers with built-in microprograms to the 370s with microprograms loaded through a Minnow drive. International standard 200-mm diskettes were widely used with early PCs such as IBM's 5110.

Later PCs, in particular but belatedly the IBM 5150, had drives for a 130-mm diskette, originally single-sided, popularly and verbosely called a "floppy disk." The higher-capacity 90-mm diskette that followed enjoyed the protection of a rigid case with a sliding shutter. This opened when inserted into a drive to give the read/write heads access to the two recording surfaces.

Like the early digital cassettes, digital compact disks for use with computers have developed from audio CDs, and the same has more recently happened with digital videodisks. These media do not record magnetically and can usually be only written on once. An exception, the CD-RW, can only be reused by blanking the entire disk.

CHIPS

Once semiconductor data storage chips replaced ferrite cores as the basis for the main store of computers in the 1970s, analysts widely speculated that banks of such chips would eventually render magnetic disk storage obsolete. This didn't happen because the chips required electrical power to sustain their data. However, they were used as the cache medium for large disk-storage units.

The obvious next step was to mimic the ferrite core in silicon-based circuitry. In 1984, an Australian company, Ramtron, discovered a way to do this by exploiting ferroelectric crystal behavior. However, Australia is not a good place for such speculative enterprises. It took Ramtron until 1988 to get a 256-bit chip into production, so the company went into receivership, was sold, and is now based in Texas.

Curiously, another Texas company last July started production of MRAM chips based on a magnetoresistive effect, and this announcement was greeted in the press as "the most significant development in computer memory for a decade" (news.bbc.co.uk/2/hi/technology/5164110.stm). However, these chips hold only 4 million bits, whereas flash memory data plugs already hold a gigabyte or more.

A less obvious development used electrons securely trapped in a floating-gate transistor to record digits, a method that Toshiba's Fujio Matsuoka invented in 1984. This is the technology behind flash memory, which has wiped out diskettes and become standard equipment on digital cameras and other popular consumer items. Indeed, data plugs can even be worn: the last one I bought came out of its box with a necklace attachment.

IMPLICATIONS

The widespread adoption of data plugs and other flash-memory devices brings both problems and possibilities.

In the short term, the data plug's convenience and universality make data theft easy—so easy that some businesses have taken to using superglue as a defense ("Superglue Used to Stop Data Theft," 4 July 2006; theage.com.au/news). Its small size also makes it easy to lose, although this can sometimes be helpful to police when a device is handed in, as in the case of a data plug holding photos of great relevance to a recent coroner's inquiry in New South Wales ("Child's Crucial Brimble Evidence," 27 July 2006; theage.com.au/news). The high capacity now available has already made it possible to use all of a computer's software and data from a data plug, making users independent of any particular computer (David Pogue, "A Flash Drive That Holds Your Computer," 15 June 2006; nytimes.com).

In the longer term, some of flash memory's characteristics could lead to ferroelectric or magnetoresistive technology superseding it in data plugs, disk caches, and elsewhere. However, the development of the Internet and its Web could also let us run programs and store data there, freeing us from having to wear data plugs around our necks.

Flash memory wasn't chosen as the most significant technology in the contest for which I served as a judge, so this essay can only applaud Matsuoka unofficially for his so very significant invention. ■

Neville Holmes is an honorary research associate at the University of Tasmania's School of Computing. Contact him at neville.holmes@utas.edu.au. Details of citations in this essay, and links to further material, are at www.comp.utas.edu.au/users/nholmes/prfsn.