

DIGITAL HARDWARE: MASS STORAGE

DR. BOYLE: I think that we are now ready to start on the other most exciting area that goes with mass digitizing. As you saw from some of the figures of the amount of magnetic tape necessary to store a map, you realize that you have to get into other means of storage. I am not a great believer in high density tapes for the future, but one of the great breakthroughs that has occurred, and I do say occurred, on an experimental and a very advanced experimental stage, is the method of drawing dots by laser beams on to disks from which you can then extract the data and use it in the ordinary way. The density of packing that you can have will no doubt come out in the discussions. If it needs any prodding I will drop in a few figures myself. We should have had John Davis from the National Security Agency, but maybe he has been tied down somewhere and muzzled, so we do not have him today. He was to have been the introducer. He has very wide knowledge of the work that different people are doing, and I am sorry he is not here. Maybe he will turn up this afternoon, in which case he will have the opportunity to speak.

I am going to ask Mr. Jamberdino from Rome Air Development Center, who has a fairly general knowledge of the different people again in the area, to speak a little longer than the others and speak first to give you more of the general background of this type of storage.

To me, it looks as if it is a vital necessity for cartography; it is a very good thing that it came along. In cartography we are riding on the backs of people who have developed these things for other aspects, partly for home video television work, partly for people like Xerox who want to store their data in digital form. These things have not been primarily developed for cartography. I think that we have to realize that in many of these very expensive developments, cartography just does not have the funds. We are very lucky that these developments have occurred. What we want to do is learn from them now, learn the capabilities, learn how we can use them, what are the difficulties, how we can interface to the sort of thing we want, how we can perhaps twist their arms a little to change the design to be a little more suitable for our applications. So, Mr. Jamberdino, would you now come up, please, and introduce this work.

HIGH DENSITY RECORDING FOR MASS STORAGE

MR. A . JAMBERDINO: Dr. Boyle, thank you for that introduction. As you have just mentioned this is the first time that I have heard that I was to present an overview of MASS STORAGE TECHNOLOGY DEVELOPMENTS, so if I stumble a little there will be two reasons for it. First of all, I'll be recovering from the shock; secondly I may seem to dwell an inordinate amount of time on this subject because last night I called my wife and she said we are going to get another 12 to 18 inches of snow today. Topped with the three feet of snow that is already on the ground right now, I feel pretty happy to be here. So, if I linger awhile, I'm sure you will know the reason why.

To begin with, Dr. Boyle, I would like to congratulate you on assembling such a distinguished collection of speakers. I am sure that with the people represented here, the audience will receive a complete, comprehensive update on where today's technology is for Mass Storage. So rather than pre-empt their presentations, I will concentrate on our (RADC's) involvement in the mass storage data handling business.

At Rome Air Development Center, we are somewhat proud of the fact that we have what we call a full spectrum laboratory. Under this charter we exploit technology developments from initial concepts, through exploratory development, advanced development, engineering development, and even to a pre-production type of a model. With that, we have access to all technology bases -- at least we feel we do -- with all the technology bases pertaining to a particular subject. Therefore the numbers you see on my viewgraphs will reflect our interests in not only today's problems but those we postulate as future requirements. So when I start talking about one gigabit recorders, it obviously does not reflect today's commercial data handling needs but what we in the Military feel are going to be the data rates of the future.

In addressing the problems of super wide band recording technologies we were forced to address meaningful record times and therefore extremely high packing densities. By addressing these extremes, I believe that we have generated some alternative approaches that will be applicable to your particular needs, that is the mass storage needs unique to the catographic community.

May I have the first Viewgraph please. (Fig 1) We sort of walked into the mass digitization or mass storage arena through the backdoor. The scenario represented here typifies a classic C³ (Command, Control, Communications) operational situation, where information is transmitted either by land lines, airborne, or satellite sensors

C³ TACTICAL EXPLOITATION




FIG 1

to a central exploitation facility. The information there will be collated, analyzed and assembled into a single product for the Tactical Commander. Our mission then is to provide the necessary technology base to accommodate the data handling or data management tools to accomplish this horrendous task. To accomplish this task we at RADC have generated three basic system configurations. The terminologies used are our own and do not necessarily reflect the Air Force's nomenclature. Consider first what we call the insurance recorder; that is a device where, all of the information that is transmitted has to be recorded or somehow stored in the event, and, hopefully this will not occur, that the system, the major exploitation system, becomes inoperative. And, as we know in computer technology today, that certainly doesn't happen. In the event it does occur however, we had to consider the necessary technology to provide an insurance recorder or insurance devices. In doing so, with the crazy numbers, I mentioned -- multi-megabits per second, gigabits per second, etc., data rates coupled with a scenerio of a 24-hour workday, where information is being transmitted in either a continuous mode or at discrete intervals at very high data rates, totally different approaches to conventional data storage had to be addressed. Typical daily volumes of 10^{12} bits demanded alternative recording technologies namely the holographic recording approach, and the high density magnetic recording devices that I will discuss later.

The next class of devices we call the "queuing" devices. Here the purpose of this device is to get the appropriate amount of information to the appropriate analyst, to the appropriate display system or what have you, so he can do the exploitation that is necessary for a finalized product. To accommodate this scenario we at RADC generated technology investigations into basically emulating the IBM computer transport. Typically, investigations towards start-stop times on the order of milliseconds, shuttling speeds on the order of a thousand of inches per second, on 2 inch instrumentation recorders, and duty cycles commensurate with 24-hour workdays, so as to allow a computer programmer the environment he is used to when handling a volume of data at this magnitude were initiated. The emphasis here then is to provide the data handling tools necessary to accommodate a number of analysts, or photo interpreters in a timely manner. The requirement here calls for medium access times (seconds), instant replay, and the capacity to accommodate very large volumes of data. Again, I will defer our solutions to this problem to later.

Finally, the third category of development, we have named the "Buffer Device". For lack of a better term the "Frame Buffer". In this scenerio consider a frame of say the "LANDSAT" sensor. Each frame contains 10^9 - 10^{10} bits. A PI must interrogate each frame,

isolate the area of interest, perform any editing, if required, then produce the processed output for further exploitation. This then calls for a much faster access, primarily on-line type of access with the flexibility of access to any data bit in milliseconds, and ideally on an erasable, reusable medium. Here we are addressing the alternatives of laser and electron beam disks. Both offer the data rate and volumes that we have postulated.

Next Viewgraph please. (Fig 2) That sort of sets the stage on the crazy numbers that I have postulated earlier. This viewgraph summarizes the various approaches that we have exploited. I would like to direct your attention to the last column of this viewgraph. Here I have tabulated on the top line what is available today in the state of the art. The bottom line depicts future predictions for that technology. To relate these numbers to what I think you are more familiar with, take, for example, the chart that Dick Clark mentioned this morning where he needs eight to ten IBM compatible 1600 bits per inch tapes to accommodate, say, one chart, or one map. The technologies represented here can accommodate on the order of 3700 maps or charts on just one single roll of tape, film or disk.

As you go down the column in the magnetic longitudinal technology, we have just delivered a 270 megabit per second recorder, which relates to a packing density yielding 2×10^{11} total bits of storage on one reel of tape. With this development we also had to address bit error rates. For the kind of readout fidelity required, error correction and detection schemes were exploited, and typical error rates of 1×10^{-6} were achieved. This machine however, is consistently achieving a BER in excess of 1×10^{-7} , and depending on the sophistication of the error correction scheme 1×10^{10} BER is achievable. Going down the chart, magnetic rotary technology has successfully demonstrated storage capacities on one roll of 3×10^{11} , and BER's of 1×10^{-6} , on similar 14-inch diameter, two-inch wide tape.

Going further down to the disk technologies, both electron beam and laser beam disk approaches are being exploited. I won't dwell on these capabilities because I'm sure that both Dr Laub and Mr Zernike will explain them in more detail. Typical storage capabilities of 5×10^{10} and 10^{11} bits have been predicted with single channel rates approaching 300 megabits per second. Again, the people on the panel, I am sure will be more specific and give more detail. Finally still another approach we are pursuing for gigabit type applications is the holographic recording approach on moving film. Here again, using this as a potential storage device, we postulate these kinds of numbers: 10^{11} - 10^{12} bits total stor-

SUMMARY OF POTENTIAL CAPABILITIES

	BANDWIDTH INPUT/OUTPUT	PACKING DENSITY BITS/CM ²	TOTAL BIT CAPACITY
MAGNETIC LONGITUDINAL	270 MBPS 600 MBPS	1.4 x 10 ⁸ 4.3 x 10 ⁸	2 x 10 ¹¹ 7 x 10 ¹¹
MAGNETIC ROTARY	20 MBPS 60 MBPS	6.4 x 10 ⁶ 6.4 x 10 ⁶	1.2 x 10 ¹¹ 9.5 x 10 ¹¹
DISK (ELECTRON BEAM)	200 MBPS	1.3 x 10 ⁷	1 x 10 ¹⁰ Per Disk
Disk (LASER)	200 MBPS	6.8 x 10 ⁷	5 x 10 ¹⁰
LASER HOLOGRAPHIC	900 MBPS 3 GBPS	1.5 x 10 ⁶ 3.0 x 10 ⁶	1.9 x 10 ¹⁰ 1.8 x 10 ¹²

FIG 2

age on a 14 inch reel. We have also demonstrated a consistent 1×10^{-6} - 1×10^{-8} bit error rates using this approach. In terms of ranking the magnetic longitudinal and rotary head are available now; electron beam and laser disk, and laser holographic approaches are in the experimental model stages of development. We are sponsoring these technologies and all have potential of meeting tomorrows data handling needs.

Let us go on to the next Viewgraph. (Fig 3) These are photographs of the equipments that I have discussed. The center photograph as I mentioned before we have just delivered, the 270 MBPS longitudinal recorder is currently operational and we are accommodating 2×10^{11} bits storage capacity per reel with a bit error rate of 1×10^{-7} .

Again, the data rate represented here does not apply to today's commercial rates but if one thinks for a minute the time it would take to present say the "LANDSAT" image to a $PI - 10^9$ bits would require approximately 17 minutes at todays computer I/O rates. I believe that data transfer rate then will become a very important parameter in the very near future.


Figure 4 (upper left). This a is a picture of the rotary head device we developed. Here again, this approach has been exploited in the early '70's. Today we are achieving something like 30 and 60 megabit data rates, but, more important, the packing density is equivalent to, and slightly exceeds the longitudinal approaches.

Figure 4 (lower left). This is a photograph of the holographic recorder currently under development. We are using high resolution silver halide as the recording material (35MM and 70MM). We have demonstrated input output rates in excess of 700 megabits per second and bit error rates of 1×10^{-6} and 1×10^{-8} . Total capacities for this device are 10^{11} - 10^{12} bits on a 14 inch reel.


Figure 5. This is a photograph of the RADC developed HRMR, Human Readable Machine Readable mass storage device. We are presently evaluating its ultimate capability. Here we have a flexibility of recording on microfiche, three different formats. One format contains 3×10^7 bits storage on a single three by five microfiche. Another alternative would be textual data in terms of pages of written texts, and then finally a third, graphic data, all three computer generated. The total storage capacity on this device right now is about 10^{11} bits, and access time, I think, is on the order of 15 seconds.

Next Viewgraph. (Fig 6) This is an artists conception of an optical video disk. Again, the panel members will go into much more

RADC WIDEBAND DIGITAL RECORDING



1979
600 Mbps
COMPUTER
COMPATIBLE




1977
300 Mbps




1976
240 Mbps
2 INCH

WIDEBAND RECORDING


**GROUND BASED
15 MHz RECORDER
30-60 MBPS**




**240 MBPS
MAGNETIC
RECORDERS**




**1GB/S
EXPERIMENTAL
HOLOGRAPHIC
RECORDER**



**AIRBORNE
DUAL CHANNEL
15 MHz RECORDER**







OPTICAL VIDEO DISC RECORDER

FIG 6

detail, so I will just brush over it lightly. We have become very excited about the potential that this device offers. As you will see this approach offers the potential of 10^{10} - 10^{11} bits of storage on a 1/4" x 12" disk, access times of milliseconds and cost predictions of \$10 per disk. Literally a computer programmer's dream. Again, I won't dwell on this subject but as you can see the disk resembles a home phonograph record.

When considering an electron beam disk approach data rates, spot sizes and access times are similar to those of the laser disk approach. We are addressing a disk format on silicon dioxide target, and we are looking at the following characteristics. 100-400 MBPS, 10^{11} bits storage, 50 msec access times. You should be aware that when using an electron beam approach, all components would have to be bottled into a vacuum type environment.

Now let me try to summarize what I have just said. In our opinion, high density recording technology offers a cost effective solution to the mass data storage problem. Several approaches are available: magnetic (longitudinal and rotary head); electron beam (silver halide and SiO_2), laser beam conventional and holographic (silver halide and thermoplastic). Each technique has its advantages and its shortcomings. Magnetic recording technology is rapidly approaching its packing density limitation and has very questionable archival characteristics. It is however, reusable and is the most utilized, understood and versatile approach available today. Electron beam technology offers orders of magnitude improvements in packing density as well as bandwidth, plus the advantages of reusability (SiO_2) and archivability (silver halide). It is, however, component size limited and requires a vacuum environment. Laser beam technology, in both conventional and holographic formats offers orders of magnitude in storage benefits and is archival in nature, however, it requires post-processing support. The fundamental technology having been developed, it now becomes a matter of time and dollars to continue the engineering development to fulfill specific requirements.

Thank you, Dr Boyle.

DR. BOYLE: We will probably ask you to come back. We now have about 25 minutes, and I am going to ask a pair of manufacturers now to speak, but try not to duplicate things. They are working in unison on this, Leonard Laub from Xerox and Frits Zernike from Philips, will be talking to you about the techniques of producing these discs and the type of hardware operations. First of all, Leonard Laub will tell you the first part of the story, and then Frits will tell you the second part. Would you like to start off?

DR. LEONARD LAUB: Here again, as with scanning, the use of disc storage for data is a piggyback phenomenon. This is perhaps an even more elaborate one than the business of scanning, because the origins of this technology in its present form do not even start with a company like Xerox, which has fairly clean and objective needs for data storage and can quantify them in terms that amount to stacks of paper and number of dollars saved. They go back to the much more subjective and, at least in some people's thoughts, larger type of market having to do with home entertainment. I used to work with Zenith Radio Corporation, and, of course, Frits works for part of the great Philips organization. Those two companies as well as several others around the world had decided that it would be a good idea to provide people with players for television records, something very much analogous to a phonograph. You put in a record, push a button and the recorded program becomes visible on your television set. You hook it up to the antenna terminals, the nice idea being in the United States that this would be something new to sell for the manufacturers such as Zenith who had already put a color television set in every conceivable home.

In Europe the justification was a little bit different, in that the government-operated broadcasting is typically so "good for you" and so restricted in its accessibility that people do not get as much television as presumably they want. Video players would enable them to buy their television programming to see when they want, even programs which are not "good for you". These are all possible market-building aspects and may someday be explored. I

regret to say that at the moment, even though the technology of the home video disc player is very well in hand in a number of companies, up to the point of nice product engineering, the market perception on the part of the manufacturers is a little bit hazy. They are not even sure whether, even at the impressively low manufacturing cost numbers they can achieve, this will be something that the world really wants. So, as a really low cost consumer device the home video disc player may be a bit in coming. However, parts of the technology which went into the home video player are forthcoming. They are forthcoming either as specialized types of video player or, in the context of today's concerns, data players and, finally, data recorder-player devices. I would like to introduce to you some of the origins of this to give you a feel for what it is that is common, and then we can get to the more specific aspects of what is happening, and Frits will tell you about the more recent work.

What I am going to describe to you is a hopeful glimpse of the future. Here is a view of what could be a typical scene in a school, and, if you change the roles of the people somewhat, a typical scene in a home sometime in the not too distant future: A young lady is being instructed by a combination of live instructor, and material, curriculum, guidance, testing or other material which is on a video disc. The video disc rereads its material through a television set to which it is connected.

Now, this typed player, known generically as an institutional player, has the property that it can be operated just as a program source, which means that you can put in a disc, push a button and sit back for the next half-hour while the program unreels, but is also capable of being controlled entirely by a computer, either a micro-processor or, by tie-line, some gigantic computer. In that role it can serve as a local mass storage device, local resource or library for an educational operation which is controlled by the computer, an interactive education approach in which students are shown material and then are tested, and on the basis of the testing and on questions that the students may ask. The computer decides where to go next on the disc. Instead of going to the very next section which is, say, English II, it will instead go to the remedial section on how to spell a word or what are these funny squiggles called letters, whatever type of remediation is called for. Or, if no remediation is called for and the second section is not needed because the student already knows it, the computer will zip over to the third section. This exploits the idea of random access on the disc and is, of course, one of the properties that everyone who uses disc-type storage likes. Now, here is a disc: This particular example is probably the simplest sort of optical disk. It is a sheet of vinyl plastic about six mils or 150 microns thick

which has been embossed with a pattern of circular tracks. That whole pattern is pressed on to the disc, not in a line-at-a-time or spot-at-a-time fashion, but simultaneously over the whole area by shoving that piece of plastic into a steam-heated press and letting it sit there for about 15 seconds.

Now, that is a very rapid type of information transfer, when you consider that the disc contains a half-hour of TV program. In more objective terms, assuming that the content of the TV program can be described in objective terms, this is a transfer rate of something on the order of 10^7 to 10^{10} bits per second. It takes 15 seconds to duplicate this whole thing. I might point out that this ultra-simple flexible transparent disc costs only about a quarter to make. Even the much more elaborate types of discs which protect the data much better and work in more sophisticated players, some of which Frits Zernike will mention, are still likely not to cost their manufacturer more than a dollar in volumes of maybe a few thousand. Now, this is just as true whether what is recorded on these discs is picture or data. Now, what do we mean recording? What is actually on the disc. The disc contains, in the video case, a spiral track. The spiral track is probed by a light spot which comes in from a microscope objective suspended over the disc, and then typically the disc spins and moves back and forth under the objective, or the objective moves back and forth over the disc.

That radial motion provides, if you wish it, a rapid access. On the disc, there is a time scale that goes round in a circle, so that the supposed one dimension of time is actually made to cover a two-dimensional area. And, as is the case with all other discs, radial motion enables you to get coarse time positioning, while the tangential motion obtained just by waiting for the disc to come around gives you fine time positioning. So, if you do not want to start at the beginning, if you want to get right to, say, one specific frame of video, or, if this is data, one bit of data, you do not have to, say, unreel a thousand feet of tape to get to it, you just go to the right radial location on the disc and wait some milliseconds for the right azimuthal location on the disc to come around. This is the sort of thing which comes out of the video disc when you send its signal into the right sort of decoder, and that into a TV set. But there are no little pictures on the video disc; all there is is an array of little marks. The tracks that I was mentioning contain marks which are about six-tenths of a micron across, and about one or 1.2 micron from one to the next.

At this, they are not as densely packed as they might be. In fact, this video disc technology is based on an arbitrary attempt to place a half-hour of video onto a disk; the density potential of this

type of storage is much higher.

Now, where is the video? For one thing, there is no area marked out -- I am talking about tracks. In fact, one turn of the spiral, one rotation of the disc is what encodes one frame of video. Remember that a television image is actually made up of a bunch of scan lines. Anybody with access to a television set -- and that is a fairly large sample of the population -- can see these scan lines anytime he likes. The TV process works by taking a two-dimensional image and scanning it and making a one-dimensional signal out of it. So, in the case of the video disc, that one-dimensional signal is encoded along a track. The method of encoding is something similar to what is used in video tape. It is a frequency modulation type of encoding in which no attempt is made to preserve the gray scale by levels, by depths of recording. Instead, what you are doing is making a mark and then making a space and then another mark. The duration of this mark and the duration of this space both contain information.

I want to trot out a pet peeve of mine, and that is the jumbling of the words "digital" and "binary". Here is a wonderful example of a system which is binary: Either you are in a mark or you are not in a mark, two states in terms of depth of depression or of presence or absence of mark. But the duration of this mark is continuously variable. An analog signal in this case is being recorded; it has no discrete levels. Mark boundaries are not located only at fixed clock positions. In the case of video recording the differences are subtle. But, let me assure you, each of these marks is of a different length, and the center-to-center distance from mark to mark is also different from one to the next. This is a binary system, it is not digital. However, it is quite possible to make the marks unitary, to make each mark like the next. That leads to a very conservative form of data recording, conservative in the sense that each of these well-defined marks has the job of carrying only one bit. But, it so happens that you can make about 10^{11} marks on the surface of one disc of about 30 or 35 centimeters diameter, about 12 to 14 inches. That means that even if you use this very conservative type of data recording scheme you can pack about 10^{11} bits on a disc. One you have done it, either by some master recording or by some local recording -- these techniques I will mention later, and Frits will have much more to say -- you can then replicate it.

The economic basis of these video discs for home presentation is the idea that once you have made a recording on the surface, that you can copy it with surface-pressing methods, such as embossing molding or casting, similar to the method of production of phonograph records. That is why the cost of replication is so

low, and that is why it is so fast to replicate: it is because you are copying a surface all at once. It is different, for instance, from copying a film or a tape, in which you must run through the entire length of the film or tape each time you want to make a copy.

Once you have the surface with the marks on it, in order to get the information out, the principle that is employed is to make, let us call it here, an "optical stylus", a small spot of light. "Small" depends on who is doing it and for what purpose, but typically the diameter of the spot of light in the plane where it is well focused will be somewhere between three-quarters of a micron and one-quarter of a micron, quite a compact spot. Because of the nice property of photons that you can put a lot of them in a place (This is different, by the way, from electrons; just a practical point, that you can concentrate light almost arbitrarily, whereas there are strict limits to the concentration of an electron beam), you can put a lot of pep into this small spot. Reading is not a particularly hard job; even a fairly faint light, such as a one or two milliwatt laser, suffices to make a bright enough spot to read the signal. In recording, depending on the materials and rates, lasers of five or ten or twenty milliwatts frequently are sufficient to enable you to make an "optical stylus" capable of making readable marks at a rate of five or ten million per second. But, at any rate, we are using light spots as though they were styli.

Now, this gives rise immediately to several nice things. One of them is that the only thing which is touching the information-bearing surface is a spot of light. When we are reading, all we have to do is make sure that we do not turn up the reading laser so high as to start melting or otherwise deforming the surface. All of these thermal effects have a very well-defined threshold. Staying below it means that you can read the same area again and again and again and nothing restrictive happens. The information is quite unaffected by reading. The other is that there is no problem in filling in most of the volume between the lens and the information-bearing surface with some transparent but solid stuff, such as a plastic or glass cover which can be as much as a millimeter thick. What this means is that the surface actually exposed to the world -- to the fingerprints, the dust, whatever -- is way out of focus relative to the surface on which the information is being recorded. If you use, say, a millimeter thick cover, the cross-section of the light at the entrance surface is a couple of millimeters across. That means that, except for really big globs of dirt, the light beam doing the reading does not know there is dirt or even fine scratches on that external surface. The information is thus quite nicely protected; you can now seal

this inside layer away from atmospheric pollution and other nasty things. So, those are some of the nice things about using light.

Now, in order to get the densities that are called for, in order, for instance, to put tracks one micron or one and a half microns apart, it is not necessary to build a super-precise machine. I have, and many other people have gone through the exercise of buying the air bearings, the granite tables and the various other devices which are necessary to make a rotating spindle which wobbles by less than one micron -- it is a very hard job. It is the kind of thing that gives a good old-time engineer a lot of pleasure. I am not old enough to be old-time and am not an engineer, but rather a physicist. I have, however, learned from some very potent engineers that a practical technique which does a better job than even the most solid and most expensive type of stable table is to take any old spindle, maybe a \$40 synchronous motor with Class 3 ball bearings that wobbles a great deal, and chase after the wobbling that spindle introduces by use of servos. Now, place a disc on this spindle. Every time it goes around, it has probably moved in and out, in the case of video discs, probably 100 or 150 microns, which is 100 or 150 times the track-to-track spacing. It is doing this perhaps thirty times each second.

Luckily, another thing we can do with light is to steer it. In the optical system is a mirror against which the light beam bounces. This mirror is tipped according to a servo loop which decides whether the spot has gone off to one side or the other of the track. The result is that even though the disc is moving like so, with the servo loop closed, the spot stays in the track -- not off as is shown here -- and moves exactly the same way the disc does. So, we make up for the imprecision of the machine with a servo loop. We take a wobble that is 100 times or maybe even 300 or 500 times greater than our tolerance and crunch it back to essentially nothing. The same thing has to be done out of that plane, in that the discs are typically very wavy, in the sense that the plane which is to be read, whether it is on the surface or buried inside the disk structure, does not stand still vertically to the half micron or so precision that is called for by the depth of focus of this tiny spot; it moves up and down tens or even sometimes hundreds of microns.

So, we put the lens in a voice-coil type device, have it dance up and down, again in response to a servo loop, and voilà, we have what looks like an extremely precise machine, even though it is cheap and simple and light and not really very precise at all. We do this by having recorded the disc in such a way that it has guidance information on it in the form of the tracks. Once those tracks are there, it is possible for the servo loops to acquire

them, chase after them, and track them very, very nicely, no matter how rapidly the disc moves around.

Now, I have been discussing the case of light going through a transparent disk. It is equally possible, and many people find it advantageous, to have the light bounce off a disk whose information-bearing surface is reflective. One process that is used in the Philips consumer video disk is to make the disk itself a rather thick transparent object, say, one or 1.1 millimeter thick, of something like acrylic plastic, have the little marks made on the far side, cover those marks with metal and then smear a little lacquer or other goo over the back so that fingers and other scratching instruments coming in from the back cannot do any damage. Handling of the disc only makes scratches on a surface which is, again, out of focus compared to the reading layer. So, this type of disc actually produces a very pretty picture even when badly handled. I personally have made some grievous thumbprints, fingerprints, even noseprints, on discs skidded them across the floor, and so forth, and they look quite nasty afterwards, but the reproduced picture looks excellent. You can do everything short of bending the disks in half or really deeply scoring them with, say, a screwdriver, and the picture is substantially unaffected. Now, in error rate terms, this means that the raw error rate of the disks is typically better than 10^{-4} , and if care is taken it can be as good as 10^{-6} . That error rate is stable despite storage and despite usage. This is a very important matter, because all the clever error correction codes in the world may not help you if storage or usage degrades the information over a period of time.

In the consumer type of video disc, the marks are depressions, pressed into the surface in the active molding, casting or embossing process. The point here is that it is easy to make mass produced discs if they are copied by a surface replication method. Optically, these little depressions act like dark areas, and a complete optical equivalent to them is openings in a continuous metal film; I think Frits will have something to say about that. The point here is that while a lot of people, particularly the people in this room, are interested in central maintenance of some sort of archives and widespread low-cost distribution of it, and so would want to use this sort of video player technology with a mass replication just as it is, a lot of other people, and certainly Xerox is among those, want to do local recording of information with the same speed and flexibility as the playing is done. For those people, instead of surface relief, instead of dimples or bumps, it seems a better idea to make marks in the form of, say, openings in a metal film or some other thing which you can produce directly with a focused laser spot.

Now, typically, in the laboratory production of master disks for video recording, rather bulky and heavy duty equipment is used, with the idea that here, cost is no object; you do not have to perform marvels of engineering, you just build a solid machine and do the job. So this machine has a lens in a rather beefy type of focus servo -- that lens is looking down on this metalized surface -- and as the disk spins the surface is being marked very precisely by a focused beam of light which comes down through the lens and is turned on and off by a light modulator, a shutter of sorts. That is what produces the marks. The disk spins continuously but the light beam is not always on, so you get sort of a Morse code pattern on the surface of the disc. This type of recorder is the one which is used to make, as I say, master records for mass replication. But, a lot of laboratories, so many -- well, not all will stand up and be named -- but many laboratories are working on technology of using the servos in the player, using the idea that the player is a good device for following a pre-established track with a light spot and keeping it in focus, keeping on that track, and turning the players into recorders, retaining the small package. A typical machine may be about a half meter across, about 18 inches or so, and about six inches high. A disc cartridge about 13 inches wide goes in through this slot. This is an example of a machine which is essentially equally happy being a recorder or a player. The only substantial difference is that in its recording mode it has a more powerful laser and light modulator inside, compared to the player, which uses a continuous beam of rather low power.

Players of this type now exist. Thompson-CSF, for instance, will be selling machines of this sort during the early part of '78, and within a month or so some people in the U. S. will be lucky enough to take delivery of one of these. I may even take delivery of one of them if I can persuade some people in New York. As video players they will be quite useful. But machines of not significantly greater size than this will become available as recorders, recorders which offer high data rate, into the tens of megabits per second, high storage density and rapid random access (tens of milliseconds or, at worst, hundreds of milliseconds access, as people now get with magnetic discs, and, of course, for the same reasons.) So this is likely to be more than just an archiving device, more than just a giant bucket or bathtub into which to throw your gigabits and terabits and whatever -- there are some new prefixes and I have not gotten them all straight for 10^{10} and whatever. Some people, I hear, want 10^{21} bits; that is a DMA-bit. (Laughter.) However many you have, it is not going necessarily to be true that you put them all in a format which is dense and cheap, but which is awkward. That may very well be appropriate for some main archive, but I think you will find more and more that there will be the type of flexibility now associated with the 10^7 , 10^8 , 10^9 bit

level -- it is going to be associated instead with the 10^{11} , 10^{12} and maybe even to the 10^{13} and 10^{14} bit level, the sort of flexibility associated with cartridge discs which are loaded by unskilled operators, queried or written by computer and accessed very, very quickly.

The concept of packaging the disc is one which is under active exploration. Various people have various needs for protection. In the Thomson player it has begun to be approached on the basis, not of making the disc itself intrinsically rugged, but of putting it in a cartridge. I just want to show you some detail of that. The package that goes in the slot in the other machine has a tray which contains the thin flexible disc previously described. That tray locks into a box, into a sleeve, as Thomson call it. You then shove the whole works into the player. The player retains the drawer, grabs the disc, lifts it up and spins it for playing purposes. When it is done, it drops the disc back into the drawer, and then you take the sleeve and shove it back into the player; the whole locked package comes out, and so human hands never touch the disk. This is the loading and unloading sequence. This type of packaging is, of course, only one of many that have been thought of. You might say in general this area of exploitation is just beginning to be well defined. It is with the aid of some people from, for instance, RADC who are in touch with specific needs that questions of how much protection and how expensive a package are finally being addressed in some rational way.

I hope that I have been able to give you some feeling for the potential of this technology in image and data storage and retrieval; it is definitely forthcoming. Thank you. (Applause).

Dr. RAY BOYLE: Frits Zernike from Philips will be continuing on the second part of the description of hardware, another hardware manufacturer, continuing from where Leonard Laub discussed his part this morning. They have gotten together, so there hopefully is no overlap. I have great pleasure in asking Frits to come on and continue what I felt was a very interesting description of the work by Leonard Laub this morning.

MR. FRITS ZERNIKE: Thank you. We heard a lot now about analog disks. I would like to talk a little bit more about the kind of thing that I am involved in, which is strictly digital disks, and digital very much in the way in which Len Laub this morning said he would not like to consider digital disks: we have either a hole or not a hole, and every hole we have is a bit. I will go into that a little more in detail, because it is not quite that simple, a little later. But, basically, it is a hole or not a hole. The difference between our system and what I would like to call the consumer video disk system, is that in the video disk you can spend a great deal of money making a master-making machine and a process which can make thousands and millions of disks, which you can then buy for a few dollars. Of course, if you buy it for a few dollars that means that somebody makes it for a quarter. But you can do that because you spend a lot of time making the master. We are looking at a rather different application, which is the application in which you make only one recording of your data, so you have a completely different kind of machine that you have to do that with.

Let us look at those differences just a little bit. In analog TV; as I said, you spend a great deal of money on the master-making machine. The master-making machine then makes the record very much in the same way as a phonograph record is made, it is essentially an embossing method. But the more basic difference is that you are not really all that concerned with the number of errors in your disk, because the TV system always has a number of so-called fly-wheel circuits. If you look at, for example, the way you people see me standing here, obviously if there were a red dot behind me all of a sudden on this wall or a blue dot on this red wall, you would know that that blue dot should not be there, so you can do a certain amount of error correction, so to say, in the machine itself. When you do digital data recording you do not want to do that, because usually the redundancy of the data has already been eliminated, and you want every bit that you can get.

Now, the machine that we are working on, and Len Laub was working on a similar thing, and there are more people doing it, is what we like to call a DRAW machine, a direct read after write machine. That means that we have a machine on which we record -- we do not have to develop-- we record and we can read that recording immedi-

ately after. When I say immediately after, I really mean we read it a few microseconds after we have written it. What we can do is this; we can record a string of data, we can store that string of data in a buffer. We can read it back out again as soon as we have recorded it. If it is not right, we say, well, let us record it once again. The only penalty that we pay for that is that we cannot use the full raw data capacity of the disk. So, in order not to have to rewrite too very, very often, we also have a certain amount of error coding in our system. What we have done at this moment, and it is still very much an experimental type of thing, is we have decided that we have an error correction system, which can correct three errors, but we have decided that as soon as we see more than two errors on our read-back system, we say this is no good, let us write it once again. With the data we have now it looks like we would pay for that in an overhead of maybe 80 percent. In other words, we have to have 1.8 times as much raw data capacity as the actual data capacity which you as the user would see.

What we basically do to do this recording is very simple. As we call it, we burn holes in a thin metal layer, which is on top of either a plastic or glass disk, or actually I should say it is on the bottom of it because, as Len pointed out this morning, we read it through the plastic or through the glass; again so that any kind of handling that you do will not show up. You can indeed put a big thumbprint on the glass and you will never see that in the data. That also means that we can make what we call an air sandwich. We put two of these disks back to back and seal them, so the actual information is completely sealed. We can seal it in a clean room, and when it gets to you it is still clean because it is sealed. The question that many people ask me is: yes, but you burn the metal out of the hole, and doesn't that metal splatter all over the place? Well, that is not what happens. What really does happen? I have an inkling, but I am not about to tell you, because I am not that sure of it, and nobody else knows what is going on yet. But it is not really splattering the metal out. Something else is happening, and there is basically no material being evaporated from the layer. So that is a problem we do not have.

Let me tell you some of the advantages and maybe disadvantages of this system. As I said already, one of these advantages is the DRAW capability, which gives us a possibility of correcting the errors as they happen, so to say. Another advantage which we think is definitely there, but we do not know yet because we have not been in the business that long, is that we think that the thing is archival. Basically, if you compare a magnetic tape, which could lose its direction of magnetization-and does, to a hole in a metal sheet, then I think you will all agree with me that that hole in the metal sheet is more durable than the magnetic disk. But whe-

ther it is really archival for seven or ten years, well, maybe in seven or ten years I will be able to tell you that, because that is really what it does take, when you come right down to it. There is another advantage which some people may call a disadvantage, and that is that the material is not erasable. Once you burnt that hole, that hole is there. Now, to some people that is a disadvantage. To other people this is an advantage. Think, for example, what would happen if a big company were to present its books to the IRS written in pencil. It would never go. I think in recent history we have a couple of examples of erasability which we would just as soon not have had. Another thing that we think is maybe an advantage, but I can well understand that maybe to this audience it would be a disadvantage, is that we make only one at a time. We make one record because, basically, the people who are making that record want only that one. They want to have it as an archival record of their computer program or all their company records or whatever it is; they do not want any more than one. They do not want to pay for it, either. Our machine as it stands now can indeed only make one. But, as Dr. Boyle said this morning, depending on -- I am also here, of course, somewhat doing a little bit of market research -- if there are a lot people who say, "Gee, we would really like to have a machine that could make a hundred or a thousand," which, mind you in the thinking of the video disk people, is still a very small number, there is basically no reason why one couldn't develop a machine that could make a hundred or a thousand disks rather cheaply.




Fig. 1

Let me just show two pictures. Figure 1 is basically a schematic diagram of a machine which we now have a few of. As you can see, there is a turntable and an air sandwich disk. You can see there actually two disks with appropriate spacers in between. Actually, the space here is shown much larger than it is in reality. It is only about a hundred microns in reality. The objective comes up from the bottom. That just happens to be better from a practical point of view. Here again is the radial tracking mirror that Len talked about this morning. Actually, what we have found is that in our recording machines we would just as soon not have this mirror; we would just as soon go with an air bearing turntable. Air bearings turn out to be not all that expensive, so you can make a very good solid machine.

The particular machine that we are talking about uses a helium neon laser. We split off part of the beam here and put it back in at a small angle. That beam, this lower beam, is our read beam which, when it comes to the disk, is just a few hundred microns behind the writing beam, and reads the signal that we have just written. You can see the other beam, the more powerful beam here, has a modulator in it, a modulator which just essentially turns the beam on or off, and that burns the holes in the disk, which make the pit. Now, when you say you burn holes, most people sort of think of holes the size of a small drill. Of course that is not the case. The holes that we are talking about are on the order of a micron or less than a micron in size. On the particular machine we are now working on, and have been using for quite some time, we have been a little lazy; we get on the order 1.5×10^{10} bits per 30 centimeter diameter disk. We could get more if we wanted to, but we wanted to start with this. The holes that we have or actually the bit density that we have is about 30 million per square centimeter, or 200 million per square inch. That is because we still have a fairly large track pitch of two microns, and holes of about a micron.

What does a disk like that look like? Well, the only other slide I would like to show you is Figure 2. Here you can see a very much experimental machine. This is the second model that we built. As you can see, the disk has pretty colors on it, which come from the fact that every one of these tracks that you see here have about 500 lines per millimeter, which in spectroscopic technique is known as a diffraction grating, which breaks up the white light into its colors. Of course, we made this picture for the annual report for North American Philips, so we said let's put some light behind it and get all the pretty colors. This particular machine as you can see is very much experimental. It is all set up on an air bearing sled.



Fig. 2

The motor has an air bearing in it so that this spindle actually turns very smoothly, and without wobble, because even though you have two micron track spacing, you want to keep that very much the same, to within about .1 microns so that you do not get cross-talk between the different tracks, and we are able to do that.

This particular machine works with blue light from an Argon laser. An Argon laser gives you lots of power. This way we do not have to worry about all our optics getting anti-reflection coated, and what have you. We can make changes very quickly and so on. We have, indeed, other machines that look much better than this one, but I do not have that pretty a picture of them. They work with a helium neon laser. I should perhaps also say that at this point we are routinely recording at a rate of about two megabits per second. And, as I have already mentioned, we get one and a half times ten to the ten, or one and a half or two times ten to the ten raw number of bits on the disk. The only other thing I mentioned already that I really should say, as I said in the beginning, that we burn one bit per hole. That is not really quite true, because our holes do not all have the same lengths. We use the so-called Miller code, which is also used in lots of communications and in magnetic data. So what we have really is three kinds of bits. One is one unit length long. One is one and a half units long. One is two

units long. Basically, we are already edging into the kind of thing that Len Laub talked about this morning.

If you think of what one could possibly do in the future, by putting in more than one bit per hole, one could certainly go a lot further. I think that numbers such as 10" bits per disk for a 30-minute disk are certainly not out of the question. But I would rather leave you with the numbers that we do now at this moment have, because, as I say, even that is in the developmental stage. So, if you say, "When can we have this?" Even that will be a few years from now. So I would rather stick to those numbers rather than, "What might we do developmentally in a year or two years from now?" Thank you very much. (Applause.)

DR. BOYLE: Thank you very much. Maybe there is a question -- And I am asking this on behalf of the audience, which I think would come up: If you are going to copy a disk, what are the sort of costs of copying a disk to re-write it -- in other words -- in the system you are talking about. Is it somewhere between ten and a hundred dollars to do it that way? I am not talking now about pressing.

MR. ZERNIKE: If you talk about taking the record and just re-recording it, it would basically, on the records that we have at the rate that I am talking about which is four revolutions per second, 240 per minute, it takes you about four hours to fill up a record. So, all you have to do is have a player and another recorder, and you could re-record that one record as many times as you want. So, basically, it would be the cost of the man who was sitting there doing it, who could just put the record on and walk away. The machine does it itself. And the cost of the record. What does a record cost? When I got into this business years ago I was once told by my then-boss, "Don't ever name a price, because if you do you'll be held to it later." So I would rather not. But it certainly is not or would not be in the hundred-dollar variety. It would be much cheaper than that. I think in the kind of thing we are looking at it is very much the case that we are looking for a very cheap record. We would rather make the recorder a little bit more expensive and keep the record cheap. Because a lot of the customers that we have talked to now want this type of machine for an enormous amount of data recording. If you talk to the satellite people and you hear the number of data they get, it's no wonder the weather reports are so bad, because they don't know what to do with it.

DR. BOYLE: The other thing that occurred to me is the aspect of the archival characteristics. People get worried that, oh, we can't possibly edit it, we are stuck with the data that's on there.

This is not, of course, true, because you only have to run a magnetic disk in parallel with your optical disk on your computer, and you can put your updates onto magnetic, and periodically make a new optical disk.

MR. ZERNIKE: Yes, you could do that. Again, it depends on what the customer wants. We see this disk very much as what computer people call tertiary storage, the type of thing where you have done all your editing, you really do not ever want to re-write that. You record it once, and that is what you want to keep. If you are interested in erasability, then we can think of a jury-rig of how to do it. We can always leave a certain amount of space on the disk. I did not go into this, but the way we record the disk now is that we sector it, and at the end of each sector we leave a certain amount of open space in which we can write, "Disregard this sector because it was all wrong". Those are the ones where we find more than one error and we re-write that again. So, if we have a good sector that we later on want to re-write, we could do what is called posting. You could go back and say, "Don't take this one, but go and look on Sector No. So-and-So for what we wrote there later". Personally, I do not like that idea. I feel very strongly that if people want erasable material, let them go to magnetic disks or magnetic tapes. This is more the thing you want for final archival material.

DR. BOYLE: In cartography, however final your data --

MR. ZERNIKE: Sure.

DR. BOYLE: Revision of maps is one of the most important things and a continuing thing, and this is how it could be used.

MR. ZERNIKE: I agree. I agree.

DR. BOYLE: I think it would be updated or disregarded and a new section put on your magnetic tape.

MR. ZERNIKE: You would read out the main part of the data from the optical disk and then make your corrections.

DR. BOYLE: Yes. That is straightforward. The other question that I had asked, which I was hoping perhaps you would say a little bit more about, but you could later, is that some people seem a little worried how they would connect one of these disks to a computer, how they could convert these holes into meaningful data. To me, as an engineer and knowing teletype units and so on, there does not seem to be any problem in this. But I just wondered if it is a straightforward, simple interface?

MR. ZERNIKE: Yes, it is that. We have heard that question enough so that we are at the moment working at interfacing one of our machines to a PDP 11, mainly because we see that interested people come in and say, "Can you really do it?" We would like to be able to say, "Yes, here it is. We have done it." But basically it is the same kind of interfacing that you have to do with any peripheral.

DR. BOYLE: Good. Thank you very much indeed. I think now we might go on to Dr. Zech, if he would come on and tell us about the work at Harris Electronic Systems Division (Melbourne, Florida).

DR. R. G. ZECH: Thank you, Ray. I guess I have a number of confessions to make. First of all, I am a pseudocartographer (in that I know enough to be conversant about the subject but not enough to be helpful) and I am also a holographer. Does anyone remember what holography is all about? (The potential it once appeared to have for high-density data storage?) Does anyone care? (Laughter)

Actually, Al Jamberdino of Rome Air Development Center mentioned in his presentation several holographic systems which Harris is developing for Mr. Jamberdino's organization. The first is the holographic Human-read Machine-read (HRMR) System which, as Al pointed out, is a microfiche mass memory system. It is based on something called "synthetic" or computer-generated holography (actually, a sophisticated form of multi-level laser recording that produces a 68-bit micro-hologram). Data is stored archivally at a density of 1.25×10^6 bits/sq. in., and recorded and read out at 0.5×10^6 bits/sec. Al also mentioned the holographic wideband recorder/reproducer system, which is capable of recording and reproducing data at very high rates. We have achieved at this date recording and reproduction rates of 750×10^6 bits/sec. I might also add that the storage density of the data is approximately 10×10^6 bits/sq. in. Those of you familiar with magnetic technology, either video or instrumentation recorders, know that this is five to ten times better storage density and three to ten times better rates than those achieved with state-of-the-art magnetic technology. Although I did not come to talk about those two systems, they do illustrate that holography is still around. I also have a certain amount of pride in them, since I have been associated with them for going on seven years now.


What I really came to talk about was something that we call the DIGIMEM (now MASTAR - Mass Archival Storage and Retrieval) System. It is a direct digital laser recording (non-holographic) system. It is the basis for an archival mass storage and retrieval system we are developing for NASA. The purpose of this

system is to archivally store LANDSAT and other types of remote sensing data. We are not in any way associated with "Video Disk," although what has been said earlier about this technology provides a very relevant background to what I am now going to talk about. Let me just give you some numbers for your consideration, because this particular system, although not a "cartographic" storage system, certainly is a remote sensing system in that what we intend to store archivally is raw LANDSAT data. Now, "raw" begs a definition here because I am not sure yet what it may mean. It may mean as the data comes off the line -- that is, as it is telemetered down to the earth. It may be after the data is processed. A multispectral scanning scene produced by LANDSAT A and B has the digital data equivalent of 2.4×10^8 bits for every scene. As I understand it, after radiometric and geometric corrections are made, its data content increases to about 5.1×10^8 bits per scene. That is a considerable amount of data when you reflect upon it. At any rate, what exactly are we doing for NASA? We are going to build for them a 10^{15} -bit archival optical mass storage system. The plan is to have it operational sometime in late 1984.

As I pointed out, the DIGIMEM System is indeed based upon laser recording, but unlike the type of laser recording (video disk) we have talked about earlier today, in which data is serially recorded, we record using a multiplicity of parallel channels. On our experimental breadboard we use eight channels, plus several channels for synchronization and other overhead purposes. The recording format is something I call a "squisk". The term squisk means a "square disk". We like many features of a disk, yet the fact of the matter is that we must handle (transport and store) the data records, and we would like to be able to handle them without human intervention. For that purpose we have designed and are building rotatable carousels (to provide rapid random access) in which to store the squisk data records. Now, these data records will not be as big as a video disk. They will be, let us say, about half the size (6 to 8 inches diameter). We plan to store about 5×10^9 bits per data record, which will correspond to approximately ten LANDSAT multispectral scanning scenes -- in fact, we will make it exactly ten, plus overhead data. The record and readout rates have not been established, but I think we are projecting something in the range of 20 to 100×10^6 bits/sec. We are talking about bit error rates

on the order of 10^{-9} (corrected), or less. A baseline plan is to have 1×10^{14} bits on-line and 9×10^{14} bits off-line. There is no reason why all 10^{15} bits could not be on-line. The data record (squisk) access time -- and the assumption made here is that we will have multiple readers -- will be about ten seconds worst case for any of the unit records. The file access time -- that is, the amount of time it will take to go from a given LANDSAT scene to any other LANDSAT scene on the same data record -- should be no more than some small fraction of a second, say, one tenth of a second. Finally, we are developing the DIGIMEM system to have two attributes which we think are quite logical at this point. It will have the capability of writing in both machine-readable and human-readable formats. I will show some examples of that a little later. In addition, the DIGIMEM data records will be replicatable in a straightforward and inexpensive way, and this in part explains the reason why we are using smaller data records. It will probably also explain the reason why the spot sizes (information packing densities) that I am going to talk about are not as small as those being talked about by Fritz Zernike and Leonard Laub.

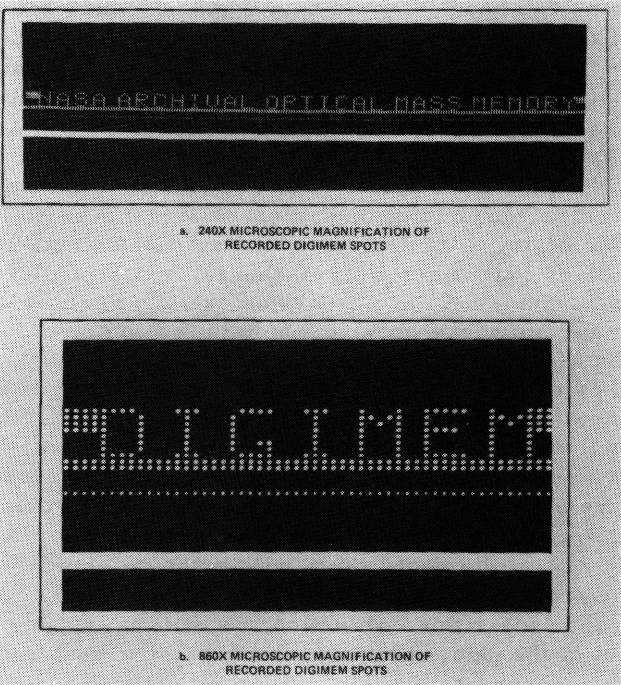
Let us go from the abstract to something a little more concrete. Would you show the first Vu-Graph, please.



VuG. 1-Feasibility breadboard Model of a Laser Optical Recorder


This is our breadboard DIGIMEM system. It appears somewhat complicated. I suppose it would not do you any justice for me to go through and explain everything that is in the system. I will, however, point out the major components. We use a helium neon laser, and, again, as Fritz said, because we are lazy we use a very large helium-neon laser, a Spectraphysics Model 125, which gives us probably 10 times more power than we need. We have an acousto-optical linear modulator array which serves as an input device (changes electrical signals into optical signals). This is the device located here. We have 32 channels available, although at this time we use only 8 of the channels for data, as will be clear when I show you a format example. We use several other channels for other purposes. The actual recording takes place here in a rotating, flexible film transport, an extremely simple device. And we do, indeed, use an air bearing platen. The air bearing has been shown to hold depth of focus -- that is, the planarity of the recording film in the region where it is being exposed, to $\pm 2 \mu\text{m}$. Further downstream in the system is something we consider relatively unique, the agile photodetector array. We call it an "electronically agile" photodetector array. An agile, or adaptive, array is required because no matter how well the servo system works -- which in this breadboard is composed of scan-descan galvonometers and an error signal derived from a separate laser beam, a tracking band and split diode detector -- we are still only "coarse servoing"-- that is, some residual positional error remains. In the present case, we are recording on a rectangular piece of film with a hole in the center of it that is mounted on a rotating spindle -- we found it didn't really matter whether the piece of film was square, circular, or rectangular -- any shape seems to work pretty well provided the air bearing is well designed and the spindle hole is centered. For recording, we put the film in, we record on it, and then we take it out. We are using a photographic emulsion at this time. We process the film, and then put it back in the system for readout. No matter how carefully done, there will be misalignment (decentering) on the order of several mils. The galvo system corrects for this, but even after this correction, we still have some slop in the readout of the data, measured to be on the order of several microns. To compensate for this, we oversample in space -- that is, if we are reading out eight data bits, we use 32 photodetectors, four detectors for each data bit, so that we can have uncorrected random spatial displacements up to ± 1 bit and still

detect the data correctly. We are also oversampling in time, a form of peak detection. I should point out that this system concept, despite the complexity conveyed by the Vu-Graph, is quite simple. We also found it to be remarkably reliable. On one occasion we selected at random a data record, put it in the reader, and allowed the system to operate unattended, with no special environmental requirements, for 72 hours. It was reading data as well after this period of time as it had been when we initially began reading data out. But, more interestingly, we found that no special requirements are required for reliable system operation. We do not use a clean room environment. We used a commercial photographic film, Eastman Kodak Type SO-173, as it comes from the box. And despite the lack of obvious precautions, the system works remarkably well. Next Vu-Graph, please.



VuG. 2-Examples of Digital Data Recorded with the DIGIMEM Breadboard that is Simultaneously Machine Readable and Human Readable

In this Vu-Graph I am showing you an example of the types of spots we record. This gives you an idea of the human-read capability of the system. Now, as it turns out, this same human-read data is also machine readable. That is, the same characters are read (decoded) by our electronics, and can then be manipulated electronically or displayed on a CRT screen. The very high magnification photomicrograph shows all the details that I think are necessary to understand the concepts. We record in a byte-oriented format -- where by "byte" I mean eight bits at a time. I should stress here that in everything we do, we try to be compatible with existing hardware. The bit sequence at the beginning of the letter "D" are signature bits. They are used for countdown -- they tell the electronics when to read user data. The two bits on the bottom of each column of spots are called data location bits. They help the electronically agile array work by providing a basis for peak detection. The single bit in between is for synchronization. Finally, at the top of this photograph, the wide band is the coarse servo correction track. Next Vu-Graph, please.



VuG. 3-Electronic Readout of Optically Recorded Digital Data
That is Also Human Readable

This is a picture taken from the face of a CRT, which is Z-modulated using the same human readable data that was shown in the previous Vu-Graph. I might point out that the spots forming the alpha-numeric are approximately $2\mu\text{m}$ in diameter. They are on $3.5\mu\text{m}$ centers. There was no special reason to choose these values. We were able to record $1\mu\text{m}$ spots on $2\mu\text{m}$ centers if we chose, but this was simply an easier way to do it. The information packing density is about 5×10^7 bits/sq. in. The user record/readout rates at this time are 1×10^6 bits/sec. Next Vu-Graph, please.

MASTAR PLAN AND GOALS

SYSTEM PARAMETER	o PAST	o PRESENT	o FUTURE	o FUTURE
	BREADBOARD (1976-1977)	ADVANCED BREADBOARD (1978-1980)	ENGINEERING MODEL (1982)	OPERATIONAL SYSTEM (1984)
STORAGE FORMAT	FLOPPY DISC*	FLOPPY DISC*	FICHE*	FICHE*
PACKING DENSITY (Mb/IN ²)	50	100	100	250
RECORD/READ RATES (Mb/s)	1	10	50	100
UNIT RECORD CAPACITY (BITS)	10^6	10^8	5×10^8	8×10^9
UNIT RECORD ACCESS TIME (SEC)	—	—	10	1
DATA FILE ACCESS TIME (SEC)	—	—	1	0.1
TOTAL MEMORY STORAGE (BITS)	—	—	10^{13}	10^{15}
RECORDING MEDIUM	AgX FILM	AgX FILM	AgX FILM?	EP OR METALLIC THIN-FILM?


*NOW 4" X 5" SHEET FILM - 9" AND 12" DISCS ARE COMMERCIALY AVAILABLE, AND MAY BE SELECTED FOR FUTURE SYSTEMS, THESE FORMATS AND "SQUISK" CONCEPT ARE CLOSELY RELATED

DATA ARE PRELIMINARY, ACTUAL SPECIFICATIONS TO BE DETERMINED BY ONE OR MORE END-USERS

VuG. 4-Plan and Goals for the Hardware Evolution of a 10^{15} -Bit Archival Optical Mass Storage System

In this Vu-Graph I have summarized what we have done, what we plan to do during the next year (1978), and then what we plan to do in the future (1979-1984). As you can see, at present the performance of our feasibility breadboard is relatively modest by comparison with video disks. This year (1978) we shall become considerably more sophisticated, and by 1982 we will be

delivering somewhere -- perhaps to Goddard Space Flight Center, perhaps to Dr. Heard's institution, the NASA Ames Institute for Advanced Computation -- this is unsettled at present -- an engineering prototype fully operational at a 10^{13} -bit storage level. If everything goes according to schedule, we will be delivering an operational 10^{15} -bit system sometime in 1984. I suppose the parameters that are of greatest interest would be those in the last column relative to 1984. "Fiche" is more or less a misnomer. I really did not know what to call it at this time. I do not believe it will be in the end a film of the type we are using now. I think a squisk is probably closer to the type of format that we will eventually use. But, as you can see, the information packing density of 2.5×10^8 bits/sq. in. will be several orders of magnitude better than what can be achieved with present day magnetic systems. The record/readout rates of 100×10^6 bits/sec., of course, can be matched by several magnetic systems. Bell & Howell, I understand, is now selling a system that is capable of recording and readout at 160×10^6 bits/sec. on 42 tracks with 33,000 bpi track density (1.2×10^6 bits/sq. in.) -- a very impressive magnetic system. You might note that sometime over the next five we will go through a transition in which we will elect to use instead of silver film either a thin metallic film type material, such as that mentioned by Fritz Zernike, or an electrophotographic film. There are several of the latter available today that are most attractive, in particular the CdS EP film being developed and now produced by Coulter Systems Corporation, which is an inorganic electrophotographic material with a speed measured in microjoules per square centimeter and resolving power measured at 1,000 cycles/millimeter. Next Vu-Graph, please.



VuG. 5-Artist's Concept of 10^{15} -Bit Archival Optical Mass Storage System in which 1×10^{14} bits are On-line and 9×10^{14} bits are Off-line

Finally, this is an artist's conception of what the 10^{15} -bit mass storage system may look like. I feel that a little artistic license was used there. I suspect the volume (size) of the actual system may be larger than is shown here for a number of practical reasons. Although, quite honestly, in principle it is possible to store 10^{15} bits in the conceptual system that we show here. The on-line data will be in the front two rows of cabinets, while the off-line storage will be in the rest of these cabinets.

I think I will end my presentation here. In summary -- optical mass storage systems with very high capacities and on-line rapid access are feasible, are being developed, and will be available for cartographic applications in late 1980's. Thank you very much. (Applause.)

WIDEBAND OPTICAL STORAGE

H. G. Heard

Institution for Advanced Computation
NASA-Ames Research Center
Moffett Field, California

ABSTRACT

This paper has five purposes. First, to focus upon the key relationships that bound the technology choices for large, archival, digital storage devices; second, to identify the motivations for selecting the optical technology for a petabit-exabit level storage system (10^{15} - 10^{18} bits); third, to present a generic example and a specific implementation of a terabit-level optical storage device; fourth, to characterize the global design space constraints that will allow one to build a technology-limited optical store; and fifth, to sketch the outline of the BYTERON concept, a wideband 10^{16} - 10^{17} bit optical store concept and contrast its performance to that of an optical store that is in operation today.

INTRODUCTION

The first purpose of this paper is to provide a perspective for assessing a given technology's value to the computational community. It is asserted that four key relationships establish the worth of a mass storage device: 1) technology maturity, 2) hardware costs versus access time to a substantial quantity of data, say a megabyte, 3) the relationship between storage capacity and access time and 4) the memory bandwidth-storage cost product. In the paragraphs that follow it will be seen that optical storage provides a technology that becomes increasingly attractive as a wideband archival store in the petabit-exabit storage volume range.

Technology maturity is most simplistically represented by the relationship between unit cost of storage (cents/bit) and time. Figure 1 illustrates this relationship for the current storage technologies. First, note that the dependent variable, hardware cost, ranges from 10^3 to 10^{-7} cents/bit over a 20-year time span. The rate of decrease of hardware cost with time is smaller for the maturing technologies; e.g., core and magnetic disks. In contrast there is rapid cost progress in the newly exploited technologies, particularly in semiconductor devices wherein gate density per chip is rapidly approaching 10^6 . It is significant that, when introduced in the late 1960's, optical storage, as represented by the UNICON, was already orders of magnitude lower in cost than other wideband technologies. Even then this storage technology competed with punched cards and

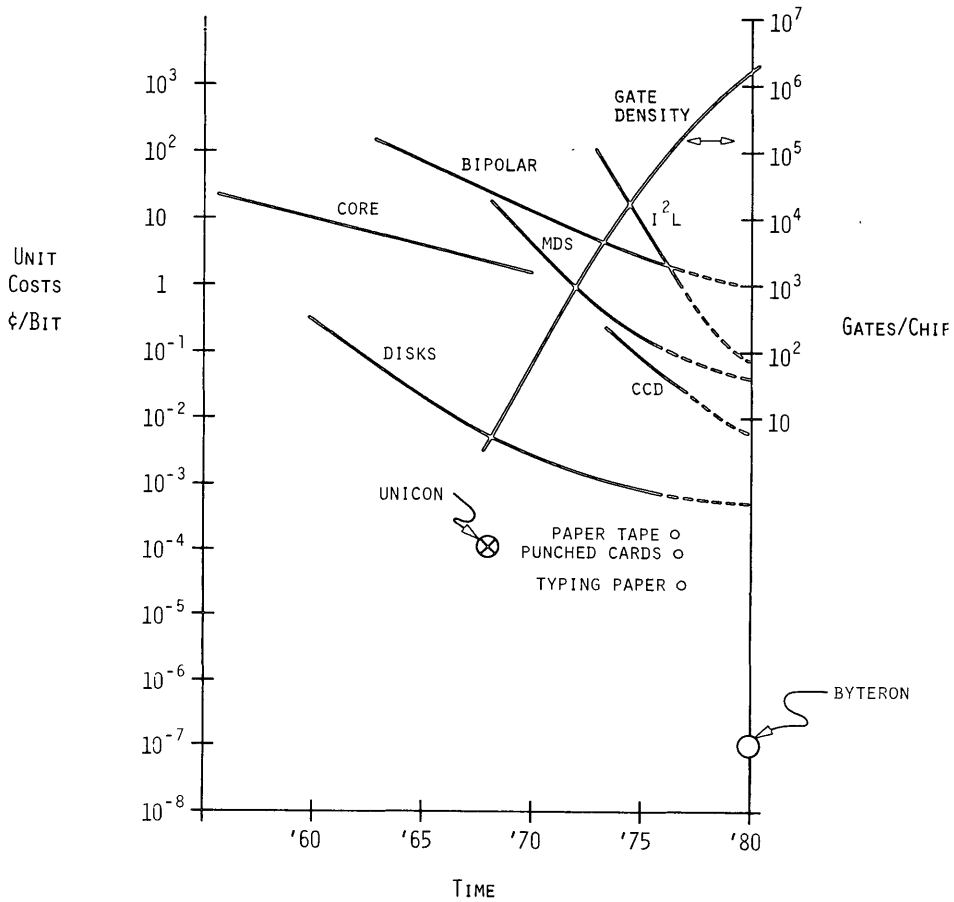


Figure 1. Read/Write Storage Cost and Complexity Versus Time

paper tape. Economies of scale are expected to enable optical storage to become even more competitive in the near future. It is projected that hardware costs for a BYTERON scale optical store, that could be available in the 1980's, will be of the order of 10^{-7} cents/bit over a 20-year time span. It is reasonable to conclude that neither semiconductor devices nor magnetic storage technologies will ever compute in the petabit-exabit storage range where optical devices are economic.

The ready availability of large scale storage, with large memory bandwidth, may be expected to have a significant impact upon the architecture of future computation systems. A cornerstone of the new design philosophy will be that wideband archival storage will be essentially free.

Figure 2 illustrates a second key relationship for storage technologies. Note that the access-time dimension of Figure 2 may be subdivided into two regions. Below approximately 10^{-3} seconds, data access is performed electrically. Storage volumes are constrained by costs to on-line capacities of the order of megabytes (8 bits/byte). Unless very significant, and as yet unknown electrical storage methods are developed in the future, charge-coupled devices and magnetic bubble shift-register-type memories notwithstanding, inevitable cost constraints will continue to force mass storage to the mechanical access domain, where large scale data availability, even with wideband channels, will be in the 10^{-3} to 10^3 seconds time range. Thus, on-line library-scale storage ($>10^{12}$ bits) will continue to be dominated by mechanical-access devices.

The on-line capacity variable in Figure 2 may be subdivided into three domains. On-line storage below 10^8 bits is generally available from a single aggregate of hardware. A few assemblies support the transition gap range, where one or a few devices with removable media, typically including tapes and disk packs, provide economic cost and performance. Magnetic-device technologies, which represent over 30 years of continuing development, currently dominate the high-storage capacity. Large magnetic tape libraries, typically containing in excess of 100,000 tapes, are found in the long-access-time multiple-device domain. Over a hundred of these libraries are now automated.

Optical storage, which very early demonstrated fast access methods and higher storage densities, has been slow to develop. A key feature of optical storage is its capability to exploit economies of scale. The total on-line storage for optical devices may grow into the 10^{14} - 10^{18} bit range, without sacrificing either memory bandwidth or

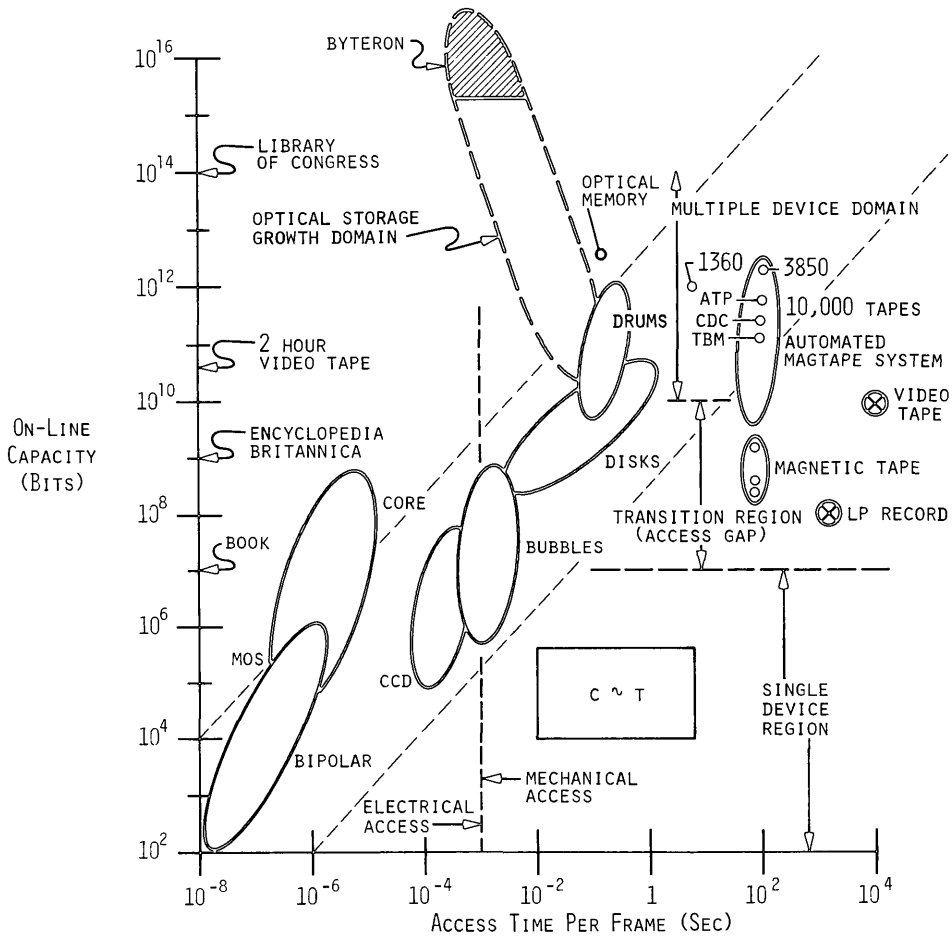


Figure 2. System Capacity Versus Access Time Write/Read Storage

access time. For example, storage strategies and hardware technologies embodied in the BYTERON concept, are projected to enable millisecond access to multiple volumes of data with gigabit/second read bandwidths.

A third key storage technology overview relationship is illustrated in Figure 3, which relates unit hardware costs (cents/bit) to the access time to one megabyte of data. Semiconductor devices, which primarily provide fast access to data, dominate the high cost area. Mass storage devices, which focus upon economic storage, are far less responsive. In a large-scale sense, the hardware cost per bit tend to vary inversely with the square root of the access time to a large block of data. Devices on the left half of Figure 3 generally tend to be more complex; devices on the right tend to be mechanical. The laser (optical) storage technology is seen to depart markedly from the general trend. Because it incorporates a relatively few complex devices and may easily exploit high storage density, a single optical mass storage system may perform the same storage functions as the 30 tape drives in an automated tape library. System reliability should accordingly be greater.

An important fourth key relationship needed for evaluation storage systems is a measure of storage quality. Strehlow¹ proposed that the product of device cost and memory bandwidth be used. He suggested that the product of an economic figure of merit (bits/cent) and a technical figure of merit (bits/second), be used to rate all storage devices. Figure 4, which plots this storage Figure of Merit 'Q' (bits²/cent second) versus on-line capacity (bits), compares the major memory technologies in this dimensionality. In this representation the electronically-agile memories dominate the high-cost low-storage volume domain. To increase the storage Figure of Merit for a given on-line capacity, it is essential that very significant improvements be achieved in both device hardware cost-performance and bandwidth. It will be seen that order of magnitude improvements in both cost and throughput are essential to improve a given technology's storage Figure of Merit.

The intermediate storage range in Figure 4 is dominated by disk technology with the (optical) video disk, assuming it can be marketed economically, yielding the highest figure of merit (in the 10¹⁰ range). The terabit storage range is populated by automated magnetic tape libraries and two types of optical storage devices: 1) the silver-halide IBM 1360 PHOTOSTORE, and 2) the metal-film storage media devices -- the UNICON 690 and 190. The domain of the ultraviolet laser




Figure 3. Cost Versus Block Access Time Write/Read Stores




Figure 4. Limiting Figures of Merit of Advanced Mass Storage Hardware Systems

store (BYTERON) is 10 petabits and beyond. The projected figure of merit of this device, 10^{16} - 10^{18} , stems from its greater inherent bandwidth, higher available storage density and overall device simplicity.

MOTIVATIONS FOR INTEREST IN METAL-FILM
PETABIT-EXABIT OPTICAL STORAGE TECHNOLOGY

Since the early 50's, large-system users have witnessed a widening gap between needed and available storage hardware for responsive, direct access, archival-quality tertiary memory. By 1974, some commercial and government organizations already had tape libraries containing in excess of 200,000 reels of magnetic tape.² System requirements analyses continue to predict yearly file growth rates of fifteen to twenty percent, with on-line 1980 data base sizes in the 10^{12} - 10^{15} size range. Studies by the U. S. National Archives indicate, for example, that U. S. government organizations yearly produce 9,000,000 reels of magnetic tape data.³ Permanent (multi-century) storage for approximately two percent, or 200,000 reels per year, of this data is needed now. Furthermore, as industry and government sources continue to convert to digital control and computer automation, to maintain competitive cost postures, the magnitude as well as the rate of file size growth will continue for the indefinite future.

Until recently, the firmly entrenched magnetic-storage technology provided sufficient growth in storage density and cost progress to discourage new technology growth and retain market dominance. There are unequivocal indications, however, that growing requirements for lower storage costs, increased storage density and archival-quality storage, cannot be met by continuing advances in magnetic tape technology.³

During the last decade, magnetic technology has been exploited to produce most of today's mass memory systems. Current magnetic-tape mass memory systems are little more than automated tape files that are cost-justified on the basis of temporarily reduced need for expensive disk storage, reduced labor costs, and increased performance derived from the now-displaced manual tape files. The cost benefits, however transient, are real and elicit user satisfaction from improved tape system reliability and reduced reruns.

Video-tape systems have been implemented by IBM, CDC, and Ampex, using either reel or cartridge magnetic tapes, to produce competitive, virtual-disk storage systems. Scores of Calcomp automated CCT libraries also number among the over fifty automated-tape storage systems that are either operational or planned to go on-stream within this year.⁴

Only two of the many optical storage technologies, that have been investigated within the last decade, have demonstrated sufficient merit, at the 10^{12} bit storage level, to challenge the well-established magnetic technology. The IBM.1360 PHOTOSTORE, which was the first introduced, implements electron-beam exposure of wet-processed silver-halid film. The 1360 ultimately evolved as a slow, low bandwidth, but capacious store.⁵ The only other practical optical storage technology to appear uses a laser to write directly upon a metal-clad polyester film. A system based upon this technology has recently been developed by the Institute for Advanced Computation. The original UNICON 690 read/writ unit, after extensive modification, was configured as a complete syste including all the software for Call-By-Name file access management, fo use as a terabit archival store.

The technology base, evolved during the successful development of this laser storage unit, has been used to investigate the viability of future, low-cost, on-line archival storage systems with capacities far greater than may be obtained by magnetic storage methods.

There are many compelling reasons for interest in very large scale laser archival storage technology. The most important technical reasons are the following:

- 1) The technology is simple and may well be the only one economically viable at storage levels of 10^{16} to 10^{18} bits.
- 2) Data are permanently recorded. No energy is required to hold information indefinitely.
- 3) Read-while-write capability enables on-line error detection and correction so that error-free data is recorded.
- 4) Mechanical contact with the record medium is eliminated and head crashes cannot occur.
- 5) Data densities in the 10^{10} bits/square inch are potentially available and a 10^{16} bit storage system may occupy 1,000 cubic feet or less.
- 6) Mechanical tolerances are tractable.
- 7) Data records may be read an indefinite number of times since read operations are nondestructive. Recorded data may be stored indefinitely.
- 8) All of the data tracks within the field of view of an objective may potentially be read simultaneously at gigabit per track memory bandwidths. Multiple read heads may be used to provide bandwidths that are limited only by electronics.

- 9) The inertialess optical beam may be scanned at gigabit rates to yield nanosecond access times to random data fields.

The initial capital investment needed to build a 10-petabit system is not small. Furthermore, considering the extremely limited market, the high risk, and the low potential return on investment, there is currently little business incentive to attract the risk capital needed to bring the technology to fruition. Nevertheless, the potential hardware and storage costs savings at the bit level are still very interesting. The potential media cost for a 10^{16} - 10^{18} bit store are so low (of the order of 10^{-10} cents/bit), that archival storage would be essentially free. Hardware costs, however, may run as high as 10^{-7} cents/bit. Sponsorship by an agency of the government will doubtless provide the incentive for bringing the optical technology to the forefront for large-scale archival storage.

THE IAC ARCHIVAL LASER STORAGE SYSTEM

The decision tree of Figure 5 gives a quick overview of the design base for an operating optical storage system, the IAC UNICON 690.

Melted metal-film optical memories are functionally similar in performance to fusible-link read-only memories; data bits, once written, may not be altered. A 500 milliwatt 514.5 nanometer wavelength argon-ion optical laser source is used to melt the metal film to post a pattern of bits on a metal-coated flexible plastic strip. The same laser is used to record and read information stored on each of the 450 on-line strips. When a strip is mounted and vacuum-retained, on either of two redundant 10-inch diameter drums, individual parallel data tracks may be recorded or read by the servo-controlled, mirror-directed laser beam.

After a strip is formatted (initialized), data are directly recorded wherever the high-intensity laser beam is focussed. A logical "1" is recorded by melting the 12.5 nanometer thick noble-metal-coated film that was previously vacuum evaporated on the upper surface of the 7-mil thick polyester strip. Once melted, the metal residue forms an annulus, the lowest potential energy form for a ruptured molten film. The exposed base film provides a 4 percent surface reflectance that yields a logical "0". The unmelted rhodium film has a 62 percent reflectance that provides a logical "1". Because of the impervious nature of the melted noble metal and the chemically inert characteristics of the heavy plastic strip, neither strip wear nor chemical oxidation are projected to impact the permanently recorded data integrity for several tens of years. These characteristics are responsible for the superior archival quality of an optical-memory data store.

Efficient discrimination between a logical "1" and "0" is provided by the large change in record surface reflectance that occurs within a

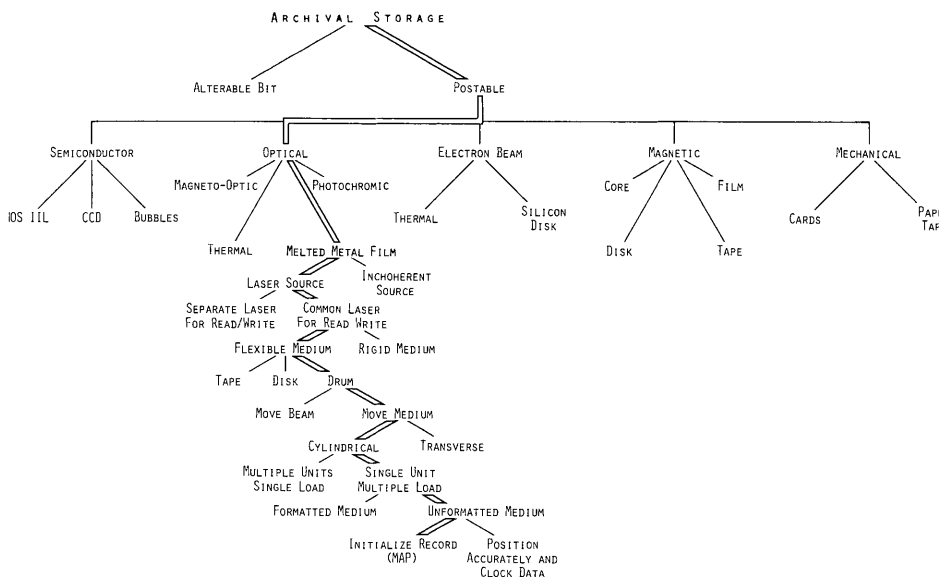


Figure 5. Optical Permanent Storage Decision Tree




Figure 6. Electron Micrograph Replica of a Typical Data Record

few nanoseconds after writing begins. The read-while-write direct-recording characteristics of the metal-film recording medium are exploited by system analog electronics, to ensure that error-free data are always recorded. Unless the surface reflectance is reduced below a well-established reference level during recording, a finite state machine (FSM) logic encodes a write error and automatically initiates rewrite of the FSM information. After several unsuccessful retrials, the FSM aborts the write process, selects a new area of the record medium on which to write, and reports the performance to the supervisory system.

Figure 6 illustrates a negative copy of an electron micrograph replica of typical data patterns produced on the plastic strip. The annular nature of the "1" bits is not revealed in this photograph. However, several important features of the data record may be observed in this figure when the average spot size of 3.5 microns is known: 1) the bit spacing is equal to the bit diameter, 2) the average track spacing is 7.5 microns, 3) there is a statistical fluctuation in recorded spot size, 4) the track spacing varies, 5) the recorded information is highly directional, and, 6) the data density could be at least doubled if the adjacent track spacing were reduced. As currently implemented, the bit cell is 3.5×7.5 microns; the data density is approximately 2.5×10^7 bits/square inch.

Four levels of error correction are currently included in lowering the raw bit error rate from 10^{-4} to 10^{-10} . During write operations, read-while-write electronics insure that error-free data are recorded, and that an on-strip record is retained containing the detailed information about each track. On-line validation of previously-written words is also encoded by a unique, highly redundant error pattern. To force independence of error statistics, all records are also redundantly written on a spatially separated area on the strip. Each record is written with its check sum for single-bit error detection. Polynomial error correction is also provided by an (80, 64) Fire Code, implemented in the hardware for on-line multiple-bit burst error correction.

While the 10^{-10} error rate of recently recorded data is marginally acceptable, dust build-up may occur over a long period of time and cause the error rate to deteriorate. Three approaches are planned to further reduce the error rate and extend the archival quality of the data: 1) add an additional on-line hardware implemented error correction code to bring the bit error rate to 10^{-13} , 2) blow a jet of dry nitrogen on the record surface when an error is detected, and, 3) under software control, completely regenerate the record from the redundant file copy whenever an uncorrected error is detected on either copy of the data. These methods are projected to give an archival error rate at the system level on the order of 10^{-12} .

SYSTEM ARCHITECTURE

The system architecture of the IAC direct access layer storage system is illustrated in Figure 7. A Digital Equipment Corporation (DEC) KI-10 (36-bit) processor and the central memory comprise the bus-organized host machine. A group of three DEC PDP-11 (16-bit) processors communicate through specially-designed high-speed data channels and memory-mapping interfaces to provide five data and control functions: 1) memory management, 2) access control, 3) master-file directory access and security, 4) data structure manipulation, and, 5) laser subsystem control. The three processors effectively interface at the central memory bus level to minimize demands upon the host and retain the needed subsystem autonomy. The laser subsystem processor (controller) manages the 18 strip-storage containers, the two read/write drums, the on-line accessed memory strips, the tracking system, the continuous bit-stream pattern recognition system (that enables use of an unformatted data medium), and all of the data verification and error correction logic necessary to achieve an acceptable system-level error rate.

The PDP-10 host processor operates under a version of Bolt Beranek and Newman (BBN) TENEX, release 133 monitor. A TENEX process, UFS, manages the laser memory file system and provides a very friendly user interface. Functioning as a high-latency (~ 10 second), high-bandwidth (5 megabaud) virtual disk, the system responds as a Call-by-Name File to enable system access to any of the 450 on-line memory strips. Currently supporting up to 512 simultaneous user directories, with 8,192 files per directory, the laser memory provides user-accessed permanent storage of up to 8 TENEX pages (512 x 32-bit) on the 11,606 tracks recorded on each plastic strip. Because the UFS software provides a highly-impenetrable isolation between users, files may be, and are, efficiently recorded sequentially. File access to a track of pages mimics a disk seek at the system level.

Separate system-support software controls the management processor. Briefly stated, a myriad of laser-system hardware, management, and control functions are handled by assembly-language-level software, that intricately interlocks the mechanical and electronic components of the system and provides a high degree of user data protection. An ultra-conservative hardware and software design policy was implemented in the system design and development as the loss of one fully recorded strip of data could represent the storage equivalent of 5 to 10 rolls of magnetic tape.

SYSTEM PERFORMANCE

The laser storage system was originally designed to provide tertiary storage to support the ILLIAC IV array processor. User data files were to be transmitted over the ARPANET for storage and staging. ILLIAC IV

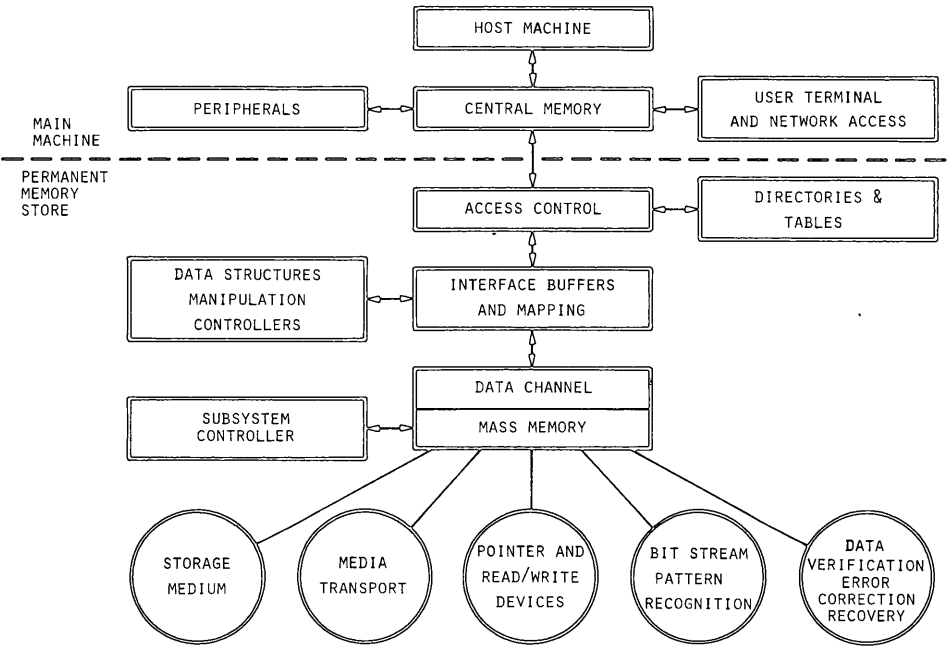


Figure 7. General Configuration of a Permanent Store

results files were to be stored on the laser memory for subsequent use perusal using TENEX processes. This heritage resulted in a TENEX-oriented record structure with up to 8 records per data track. Data words of 64-bit length are derived from ILLIAC IV requirements. One ILLIAC IV page (1024 64-bit words) is recorded as 4 records on a track

The PDP-11 controller software restructures its 16K data access from central memory as 16,784 bit pages. System software enables the laser store to handle either 32- or 36-bit words from central memory. When 36-bit words are recognized, they are handled redundantly as both left and right justified 32-bit words; 36-bit words reduce the record efficiency by a factor of 2. Reconstruction of the original word is handled, during read, by effectively or-ing the redundant middle 28 bits. Storage of 36-bit files on laser memory is only half as efficient as 32- or 64-bit files. Each record consists of a 256-bit quadredundantly recorded header, 16K bits of data, 260 bits of error-correction code, and 128 checksum bits. Each of up to 8 fixed-format records of 20,928 bits on the 11,606 data tracks, is transmitted at a peak data rate of 5 megahertz. The average data rate, when all of the record keeping and redundant operations are accounted for, is approximately 3 megahertz.

GLOBAL DESIGN CONSTRAINTS ON PETABIT-EXABIT LASER ARCHIVAL STORE

There are seven major global design limits that establish bounds on the volume of design space within which all metal-film stores may exist: 1) recorded spot size, 2) data cell size, 3) write bandwidth, 4) read bandwidth, 5) media surface speed, 6) record width, and, 7) record environment. These variables unite to form an equation of state and define a design volume, a subspace of which permits a practicable optical store to be designed.

Record spot size is limited by spherical aberration and diffraction to approximately a wavelength. Spots of the order of a wavelength have been recorded with a 514.5nm laser. It is projected that a lower limit for the spot size of 0.3 microns may be achieved with a near-ultra-violet laser.

Cell size on the UNICON 690 is limited by the tracking system. Each track is resolved by a line-following, sample-data servo system that places data in the space between clock bits previously recorded on the record surface during initialization. An integrating track system, that follows a reference established by a laser interferometer, should permit bits to be recorded in every bit cell. Accordingly, a 0.3×0.3 micrometer cell size represents an upper storage limit of nearly 10^{10} bits/square inch.

Write bandwidth is limited by the physical mechanisms of laser-induced material damage. Tests indicate that laser-induced damage may be produced reliably with picosecond length pulses. A mode-lock gas laser, operating at 87 megabits/second, or a quadrupled, mode-locked, solid-state laser, operating at 1000 megabits/second, may thus establish a reasonable upper limit to the write bandwidth.

A unique characteristic of optical stores enables multiple-track reading. To the extent that the tracks may be resolved, focussed upon a detector array, and electronically deskewed, all of the tracks within the field of view of an objective lens may be read simultaneously. Since over a thousand adjacent tracks may be easily resolved, 128 bytes of data may be read simultaneously; each of the 1,024 tracks may be simultaneously read serially at detector-response limited speeds. As solid-state detectors are available with nanosecond response times, a read bandwidth upper limit is probably 1024K megabits/second, per read head. Multiple read heads may be used, of course, to produce read bandwidth that is practically unlimited.

Media speeds must remain significantly below the velocity of sound if vacuum loading is to be used to retain traction between the storage medium and the record drum. A reasonable upper limit for linear speed is approximately 500 feet/second. Higher surface speeds could be attained in a vacuum with the forfeit of single media replacement.

Recording media width is limited by economic producibility and accurate mensuration. Since a laser interferometer may reliably measure distances up to 50 feet, record width becomes constrained by more pragmatic considerations. The length of a data sheet is governed by the mounting mechanism. Single sheets of 8-10 foot length could readily be handled.

The most serious limitation to storing information at submicron levels is dust. Methods have been developed to sandwich the record material. Data may be recorded and read within the focal depth of an appropriately designed optical system. Thus dust will no longer limit storage lifetime.

The preliminary design parameters for the BYTERON, a 10^{16} -bit laser archival store, have been selected as a subspace within the above constraints. Assuming a 0.3 micron spot, a 0.3×0.3 micron storage cell size, and a 90 percent storage efficiency, 1.74×10^6 square inches of record area will be required. If the largest conveniently manufactured record width is set at 48 inches, 36,328 lineal inches of record space will be required. A little over 350 data sheets of 96-inch nominal length will more than suffice for storage. A mechanical storage mechanism, that will provide on-line indexed access to individual data strips, could be conservatively constructed within 1,000 cubic feet if the film were loaded on a 3-foot diameter drum.

(Film sheets and drums of this size are currently used in laser recorders in map making.)

It is of value to consider briefly some of the comparative storage parameters for the BYTERON and UNICON 690. A convenient medium for comparison is the set of fully recorded 1600 BPI CCT data tapes; some relative parameters are listed below:

TABLE 1: STORAGE COMPARISON

<u>PARAMETER</u>	<u>UNICON 690</u>	<u>BYTERON</u>
Tapes/Strip	5	30,000
Tapes/Pack	125	750,000
Tapes/System	2,250	$> 10^7$

SUMMARY AND CONCLUSIONS

This paper addressed a substantial conceptual advance in laser archival storage technology. A device, called the BYTERON, focusses an ultra-violet laser beam upon a recording medium to create a permanent data record. The new technology represented by this recording methodology leads to attractive storage densities, per-bit hardware costs, multi-rack read bandwidths and on-line storage capacities in the petabit-exabit range (10^{15} - 10^{18} bits). Because each data record is permanent, truly archival storage (~ 25 years) may be obtained at data densities that are beyond the range of magnetic storage technology. Direct access may be provided, to the storage equivalent of over 30,000 1600 BPI magnetic tapes, in less than a quarter of a second. A remote user may select a file from the storage equivalent of tens of millions of magnetic tapes in seconds.

Optical storage technology, when extended to the petabit-exabit storage range, offers an exciting opportunity to obtain a truly archival storage at per-bit hardware and medium costs that are lower than that of any other known media. Inherently a simple device, an optical storage system may be constructed in a small volume because storage densities of the order of 10^{10} bits/square inch are potentially attainable. Read and record bandwidths in the gigahertz range are projected and limited only by electronic device technology.

ACKNOWLEDGMENT

The work reported herein was jointly sponsored by the Defense Advanced Research Projects Agency and the National Aeronautics and Space Administration.

REFERENCES

1. W. H. Strehlow, "Symposium on Advanced Memory Concepts," Stanford Research Institute Report, June 1976.
2. UNICON 690: Metal-film data storage medium originally manufactured by Precision Instruments, Inc., Santa Clara, California.
3. Thomas K. Gaylord, "Optical Memories," Optical Spectra, June 1974, p. 29.
4. Charles M. Dollar, "Problems of Magnetic Recording in Archival Storage," Proceeding Comcon 77 (Spring), February 1977, p. 28.
5. Robert Howie, "Summary at Third IEEE-CS Workshop on Mass Storage Systems," Palo Alto, California, April 1977.
6. Garret L. Boer, "A Look at Mass Storage," UCRL-76835, NASA #N76 2586, Lawrence Livermore Laboratory, University of California, May 1975.
7. D. Milan, "Laser-Induced Damage at 1064nm," Applied Optics, Vol. 16, 1977, pp. 1204-1213.

DR. BOYLE: Thank you very much. We are getting bigger and bigger systems, you see, in our discussions. I am now expecting Bill, warning him that he must have 10^{20} bits when he starts his talk. (Laughter.) I would just like to mention that on some calculations I did this summer, it looks as if we could put one topographic quad sheet, if we compacted it in a clever way, into 10^7 bits. That means, just to give you some of the parameters, on a 10^{10} bit optical disk we could get a thousand topographic quads, just to give you some idea of the calculations, the way they work out. The last one we were talking about was, of course, very, very much larger than that. Remember, every time the digit goes up in the power, you are ten times. Bruce

Mr. BRUCE OPITZ: I am from the Defense Mapping Agency, Hydrographic Center. My role, in the Advanced Technology Division, is more as a potential user of some of these devices that you have heard about here today than a developer. You heard during the keynote speech that DMA is looking at a future data base requirement in the mid to late '80's of 10^{15} to 10^{16} bits. To give you a feel for the problem that we are all facing, certainly DMA is facing, 10^{15} bits is the equivalent to a minimum of one million conventional magnetic tapes. That is a true minimum. That is not allowing the industry standard of 20 percent of a magnetic tape being used, and that is not allowing any backup. So, when you are talking about 10^{15} , or 10^{16} , we have tremendous problems unless some of these devices we've heard about today become a reality.

Currently there are some systems available which are billed as mass storage systems -- Calcomp, Precision Instruments, IBM, CDC, and Ampex. These systems are typified by a medium amount of storage, 10^{11} , 10^{12} , or in that range, on-line. The key to those systems is that they offer on-line storage to a large computer center; they use conventional methodology, not the kind of optical techniques that we were talking about today; and that they are typified by horrible mechanical devices; most of them look like a Rube Goldberg device.

What we are looking for is something closer to a video disk or perhaps Digimem system for high density storage. One of the things that we as users must start to do, and I think this has been indicated by some of the comments, especially from North American Philips, is to start giving some feel of just what our requirements are. DMA is currently involved in looking at the configurations that will be required in the late '80's time frame. Some of the questions that we have to ask ourselves and provide answers to the developers of these hardware devices is, do we really need on-line access? Or is an off-line access, similar to the way we operate our tape libraries today, is that sufficient? Is our prime inter-largely archival? I think I should define archival as being slowly

changing. And I think most cartographic data is slowly changing in terms of every six months or much less than that -- not that you are going to record a file and keep it there forever. But you are going to record a file and slowly change it so you can afford to look at it as an off-line archival medium. If we are going archival, are we going to have one copy of these devices or two copies hooked into our major computer center? Or are we going to look at this as an off-line device, still using some conventional medium like mag tape or magnetic disk, in an on-line mode, copying from the mass storage, perhaps optical storage, to the magnetic disk, and using that as the transfer medium to our computing system?

Another question and one that has been brought up is that is there a requirement for duplicating the file in the form that it is being stored, or will duplication continue to be on the basis of sending out mag tape? Is the user community at large going to require these massive amounts of data, or are they only going to be held in a few major repositories like NOAA, GS, DMA -- and are they the only people who are going to require the mass storage? I just want all the users here and potential users to start thinking about these questions now, not tomorrow, and to start a dialog with the hardware developers so that we can make sure that the right technology for our purposes will be pursued. This technology is off in the '85, '90 time frame, I feel, at the very least. What we have right now is, at very best, some devices which are maybe a half step beyond being laboratory curiosities. We are a long way from a true integrated system for practical use. Thank you. (Applause).

DISCUSSION
DIGITAL HARDWARE, MASS STORAGE

DR. BOYLE: Thank you very much. I think we have now completed our panel presentations. The floor is yours.

MR. BILL ROSEMAN: Bill Roseman, Iometrics Corporation. I would like to point out for those of you attending this conference that there is a small company in this area down in Santa Clara that does have a video disk, prototype; it is up and working. I am sure that if any of you get down in that area that they would be more than willing to show you the disk in operation so you can actually see what this mystifying video disk actually looks like in operation. Thank you.

DR. BOYLE: Can you give the name of the company?

MR. ROSEMAN: The company is Videonics of Hawaii, Inc., and they are located in Santa Clara at 3022 Scott Boulevard.

DR. BOYLE: I am also sorry that Leonard Laub had to go, but he would also be very pleased, I am certain, to see any of you down in Los Angeles at Xerox. There are some very interesting things to be seen.

MR. ZERNIKE: The comment I would like to make is that maybe Leonard Laub would like to have all of you come down to Los Angeles. Please do. I definitely would not like to see all of you come to Biercliff because we do want to get this thing developed at some point in time. (Laughter).

MR. ZERNIKE: All I can say in answer to the comment from the company in Santa Clara: We do have video disks, and we can show them. We make them every day. I can tell you that as far as analog video disk is concerned, some of you may remember the end of the movie "The Sting." I can tell you that Robert Redford does an absolutely perfect job of falling down, because I have seen him fall down and come back up thousands of times now by showing the video disks to interested customers. I would rather not do it another thousand times, thank you.

DR. ZECH: I would just like to extend an invitation to all of you who might want to come to Melbourne, Florida, particularly at this time of year -- I assure you, the weather is beautiful. We would be happy to not only show you some disks, but we will show you an HRMR system, a wide band recorder-reproducer system, an acoustic traveling wave system, and some of the other wonderful mysterious things that are going on in electroptics. I think the point I am

trying to make is that it is not that I am sure Fritz Zernike at Philips does not know how to make disks, it is not that Harris Electro-Optics could not come out with some sort of product. What we are trying to do is, first of all, find out what the problem is. Once we have established what the problem is, to try and find a systems oriented solution to it. Although I did not mention it, over the last year or so, part of my relationship with NASA has been to slowly accumulate information about real needs. NASA is a very big organization. As I am sure it should be clear now, Dr. Heard is pursuing a somewhat different approach than what I am suggesting for yet a different part of NASA. Thank you.

MR. JAMBERDINO: As long as everybody is inviting everybody, I would invite anyone, if you have the guts, to come to Rome, New York. (Laughter). There is about four and a half feet of snow, and I am running out of a time and place to put the snow. I would like to make some comments, however. The point was well made that one should define your requirements. We have heard inputs here from 10^{12} , 10^{13} , and 10^{16} kind of memories. You people are unique in that you have these kinds of requirements. I also mentioned the fact before that it took 30 years and three billion dollars to bring the magnetic technology to where it is today. I do not assume that it is going to take that much, but unless you people give us the requirements that you need and back it up with money, we will stay here and tell you all the beautiful things that can be done in the future. There is a very big problem in terms of defining your needs and backing them up with money. That is all I have to say.

DR. BOYLE: Yes. It needs cartographic thinking. I myself like the idea for automated cartography, I am not talking about LANDSAT imagery -- only chunks of 10^{10} bits. Could I quickly ask Harry Heard if we will be seeing anything of this down at the NASA Center when we go on Friday?

DR. HEARD: Yes. We will open the Unicon and let people actually look at the various parts of it so they get a feel for the technology.

DR. BOYLE: Thank you.

MR. HAL MOELLERING: Hal Moellering, Ohio State University. I have a general question which I would like to address to the members of the panel. There was a recent issue, special issue of the Institute of British Geographers. I think it was Volume II, No. 1, new series. It was devoted specifically to cartography. The lead article in that volume was by Robinson, Morrison, and Merky. The title of the article was "Cartography, 1950 to 2000." They were reviewing the developments in cartography, some of the innovative developments and techniques in terms of computers and so on. They were making some

suggestions and predictions about some of the products of the future. One thing they mentioned which I have not really heard here, and that is the type of product called the orthophoto map. Now, one of the problems that has been discussed here is the cost of digitizing. In orthophoto maps, you are essentially processing the information optically all the way through. Now, it is true, if you are going to produce orthophoto maps, at some point you are going to have to put annotation on the map -- names, things like that -- which cannot be gotten, say, with their photos. So, there is some digital information processing going on. However, if you do produce things like orthophoto maps, it seems that then you are essentially finessing, or perhaps might be able to finesse, a large part of the digitizing process, whereby you are taking the information you are gathering at some point, air photos or other sources; you are digitizing into the numerical domain, then processing it into the numerical domain, and then re-transforming it back into the graphic domain, and then produce the map. Is it not possible, say, using something like orthophotos to maintain your information in a two-dimensional optical state instead of some kind of linear numerical state, and thereby -- perhaps we shall say "finesse" a good part of this problem of digitizing or scanning, however you like to call it?

DR. BOYLE: Would anybody like to respond to that? I think you have beaten them.

MR. JAMBERDINO: I think we need the definition of orthophoto.

MR. MOELLERING: An orthophoto map, an orthophoto is essentially a type of air photograph which has been rectified optically to look as if it is orthogonally viewed from all points in the image. That image then could be printed on a map, and you could add other information like important highways or, say, political boundaries, names, things like that which would not appear in a photograph. But what you have done, you have gathered this information with an optical system, you have kept it in that domain. You are obviating a lot of the problems that you might get from making these transformations. And, digitizing is essentially a transformation from some kind of graphic domain into a numeric domain.

MR. PALERMO: Not addressing the problem of actually going to an orthophoto type of a storage or medium, the problem basically driving DMA today, and therefore driving governmental labs like RADC and ETL, is not the production of graphics. They have had digital systems in-house for a number of years, and in some cases, like Hydro Center just produced their first digital production graphic. It is the call for other digital cartographic products which are necessary for the systems being used within DOD, and DMA is one of those agencies that has to produce this data. It is not a question

of whether we can get around it with some other product. That question has already gone by. Most of DMA's problems now in the production of these products is how do we most economically produce digital data? It is not a question of whether we are or whether we are going to get around producing digital data. It is, we are doing it, we have to do it. What is the best way of getting it? The data has to be stored. What is the type of system for the future that is going to have to store these amounts of data? The thing that industry and the commercial field has to answer above and beyond that is, is it necessary for them to go to a digital data? The DMA has been tasked that they will, and they are. That is the basic difference driving all of these systems, many of the government DOD labs and the other areas.

MR. OPITZ: I would like to make a comment on that, too. The Hydrographic Center, in addition to supporting the fleet, also has a charter to support the Merchant Marine. And we sell our products, which are currently hard copy products, to the Merchant Marine, and the digital products are mostly going to support advanced weapon systems. But I think it is a big mistake for us to look and think that we are going to get away with producing only hard copy products in the future. There are more and more signs of the user, of the civilian user requiring a digital product. In the hydrographic area there are coming along both collision avoidance systems and navigation systems which are going to require digital charts. There have been experiments, some systems have been sold, where data was supplied by the vendor. The Navy has been looking at collision avoidance systems using digital charts. More and more, at least in the hydrographic area, the digital product is going to become a requirement. I think that maybe all the mapping agencies are going to have to look at the point where the digital product is going to become more popular and more demanded than the hard copy product.

DR. BOYLE: This means that both your ends will be digital: Your collection of data and your output will be digital.

MR. MONTERO: I would like to pose not a question, but a thought. And that is, with the capability of extremely fast mass digitization as we heard this morning, and the capability of virtually immediate processing of those data, is not perhaps it more worthwhile to consider re-scaling original source material and storing that rather than spending enormous sums of money for mass storage devices? Any comments?

DR. BOYLE: Would anybody like to comment on that?

MR. OPITZ: I think a lot of the same comments apply.

MR. ZERNIKE: The comment I would like to make, as I said before, I do not want to quote prices, but certainly the kind of mass storage that we are now talking about to customers in other fields will not go past -- and I think I am being quite generous -- will not go past a hundred thousand dollars for the recorder and \$50,000 for the player. As a matter of fact, I think they will be quite a lot cheaper than that. That includes the interface but does not include the computer, quite obviously. Again, if you compare that to what Dr. Heard was talking about, that would not be mass storage. It would be small storage.

DR. BOYLE: I think we have to look at the job we are trying to do. You must do your cartographic thinking. We are talking here about things that are possible to do at this time.

MR. GENE SLOTTOW: Gene Slottow, University of Illinois. I just wondered, in considering these various memories, you might remark on how you compare the electron beam memories, which I think did not get too much discussion this afternoon, along with these other mass memories.

MR. JAMBERDINO: There are a very few electron beam capabilities in the industry now. We have tweaked several. You know the work GE is doing on Beemos, I think Datamation is another company; CBS is no longer in the business. It went from CBS to EPSKO to CARSON-ALEXIO, and we lost the people that did it. Image Graphics is doing some work in this general area. I think the fact that it has to be enclosed in a vacuum type of a system creates a difficulty, and how many bits can you really store and how many tubes, if you will, are going to be required. This is still a question. I think we are still in the exploratory domain at this point in time.

DR. BOYLE: Gene, I wonder if you could just comment on this work at the University of Illinois where you have been doing some experimental work -- on the use of optical disks in teaching systems. I was very intrigued by the comment I had from there, that one of the important things was to speed up the access to less than a hundredth of a second, because students just could not stand to wait for a tenth of a second for their answers.

MR. SLOTTOW: I do not think that last one is quite right. When you get to a tenth of a second as far as a student is concerned, it is instantaneous. If you wait a few tenths of a second, then that becomes noticeable. Most of the design in the Plato system, for example, at the University is essentially built around the idea of a one-tenth of a second response time. There is not much work, I think, in the use of disks so far in educational systems. The disks are primarily used for just archival storage and for lesson

materials. For a person working on-line, he is really interacting with the computer with a very, very fast transfer memory, and working at something like a hundred million to six-hundred million bits per second transfer rate between the main memory and the storage memory, the time sharing memory. That is one of the reasons why you can support 500 to 1,000 people with that kind of response time.

There is some work going on in video disk investigation, more or less along the lines of the discussion this morning and early this afternoon, in which holes are basically produced in metal film. This will provide a rapid access support memory to some kind of future terminal.

DR. BOYLE: Thank you. Waldo?

DR. WALDO TOBLER: At ten meter resolution, and depending a little on how you quantize it, ten to the 13 is roughly adequate to have the whole topography of the earth stored. In the United States there are approximately 35 to 40,000 topographic maps now in existence, and I do not know the comparable number for the whole world; and, of course, there is the bathymetry and so on. Could you make an estimate how long it would take to get that information into one of these mass storages? Because, as you recall, Shimmof said that at current 6,250 bits per inch tape drives, it would take 30 years, 24 hours a day, to fill up one of these memories. Is this a reasonable problem? Is this a problem?

DR. BOYLE: All these things are problems, yes. (Laughter). I did some estimates on the 1:24,000 series, the topographic quad series for the states, and it looks as if a good operational period of time for digitization was five years. This was with acceptable cost and high data quality. I think to answer that "off the cuff" would be asking for trouble. Was there another comment?

MR. TOWNS: My name is Towns from Graphics Technology in Rockville, Maryland. We all know the RS-232 interface that is used in machines talking to each other in a standard, ASCII formats, and so forth. Is any sort of planning or thinking going on along this line to try to standardize the methods by which such obviously highly technical interfacing is going to have to be done for these new machines? Any planning being done for the future so that, if we do have a very fast recording holographic system, we can standardize for those of us who have to make and store data?

DR. BOYLE: The interface will have to be directly into the storage system. I do not think you could use one of the slower standard interfaces for this. It will be special for a particular mini-computer. I see a mini-computer being directly attached by its own interface, and then that mini-computer will talk to something else by one of the standard interfaces.