IBM 3340 Direct Access Storage Facility, 1973

The “Winchester” program pioneered the use of low cost, low-mass, low-load, landing heads with lubricated disks, becoming the dominant technology for at least the next twenty years. In this particular product, the heads remained with the disk in a removable data module, a packaging concept that marks the beginning of the Original Equipment Manufacturer (OEM) disk drive industry's departure from following IBM.

Why it’s important

Al Shugart in a 2000 interview stated “the low mass lightly loaded head or, as some people call it, the Winchester head," was one of the four most significant events in the history of mass storage.

In combination these technologies provided:

- lower head flying height, thereby increasing capacity by enabling bits on a track to be closer together;
- lower head cost thru simplified design and process manufacturing;
- higher yields and reliability by having the heads and disks permanently associated.

These technologies and variations thereof were ultimately adopted by all hard disk drive (HDD) manufacturers and dominated the industry into the 1990s.

The particular embodiment, heads, disks and other components packaged in a removable data module was not followed by the OEM disk drive manufacturers, who instead chose to continue with the then-conventional disk pack and heads. This marks the beginning of the OEM industry deviating from IBM standards for the of disk drive technology.

Discussion

The IBM 3340 Direct Access Storage Facility was conceived to respond to IBM’s requirements to provide a lower cost storage subsystem for its upcoming low end System/370 mainframe computer models and to respond to the ever increasing competition for such storage facilities from the newly emergent and aggressive plug compatible disk subsystem manufacturers. IBM San Jose management decided to respond with technology rather than a stripped down IBM Model 3330 and initiated the project under Ken Haughton in 1969. The product was announced in May 1973 and began shipping in November 1973.

The concept was to have a removable Data Module containing both the magnetic storage media (disks) mounted on a spindle, and a carriage on which the magnetic heads were mounted. The heads were to start and stop in contact with the disks on a dedicated landing zone but fly over the disk on an air bearing generated between the magnetic head and spinning disk while reading and writing. This was in sharp contrast to the conventional disk pack technology where only the media were removable while the heads, spindle and carriage were all a part of the disk drive.

The Winchester data module is frequently referred to as "sealed;" however, it was not in a conventional sense "sealed" since the module had a roll-top type of door that opened (it actually rolled down) to allow the actuator in the drive to connect to the carriage in the module and to connect to the drive's clean air system. The module then sealed against the drive so that the exposure to the ambient environment was minimized.

Haughton and the San Jose laboratory management were attracted to the concept of a fixed media drive because the cost aspects of always having the same head reading the data that did the writing was compelling, that is, tolerance and alignment requirements are reduced enormously. The strategy was to go to fixed disk drives, and the original plan did include a proposal for a drive called “Weatherbee” that later became “Madrid” and later the IBM Model 3350 (though design did not begin for several years). Haughton made several trips to the Data Processing Division (Sales) to understand the application of the removability capability of the disk pack in small systems but the sales message was “they gotta be removable because they always have been.” Even though the number of pack changes was small, resistance to fixed packs was very high and the DPD viewpoint won out.

Because the heads would be shipped in the Data Module, a great deal of emphasis was put on coming up with a very low-cost head. The start/stop in contact requirement meant the heads needed to be very low mass and with a very low load force to minimize friction and wear. For reference the magnetic head target cost (<$1.00), mass (0.25 gram) and load requirements (10 grams) were all an order of magnitude less than that of any existing production head.

Feasibility of start/stop in contact with a lubricated disk and Data Disc licensed heads was demonstrated by Joe Ma in two projects, first in a single disk buffer for the Rand Corporation and then as the IBM Aries. The 3340 started with the Data Disc head [2007Haughton, p. 11] but when it turned out to be too expensive and potentially unreliable, the team, principally Mike Warner, invented what is today know as the “Winchester” technology, see US Patent 3,823,416 -- "Flying Magnetic Transducer Assembly Having Three Rails," and head development details discussion below.

The initial requirement was for two drives in one box, each having a capacity of 30 MB per Data Module. This 30/30 configuration led to the code name “Winchester.” Contrary to urban legend, the modules were always removable [2007Haughton, p. 15]. Subsequent market analysis led to requirements for modules of 35 and 70 megabytes and these capacities were the ones announced.

While the Winchester head’s configuration enabled substantially lower flying height (18 vs. 50 micro inches [2007Haughton, p. 11]) which in turn could have allowed much higher areal density – the actual product’s specifications were only moderately higher than of competing products such as OEM versions of the 3330 model 11 (12.2 vs. 10.5 megabytes/surface). In part this is due to conservative management, Haughton said he “picked out of the air” the 3340 track density of 300 tpi and has “kicked myself ever since for not saying 500.” [2004Panel p. 9-10]

Another key innovation in the 3340, by Donald Frush, was its implementation of defect skipping. It is impossible to manufacture an error-free disk so in prior art the minute imperfections in a disk drive’s disk led to loss of large regions (one or more tracks or sectors) of data or possibly the rejection of the entire disk or pack. In the 3340, and ultimately in all subsequent disk drives, the recorded data are split into pieces that avoid the defects, see US Patent 3,997,876

Richard B. Mulvany and Rudolf W. Lissner of the 3340 team invented the Data Module itself [1974Mulvany]. Placing the heads in the Module improved manufacturing costs since the head always worked with the same disk as opposed to disk pack drives which had to deal with differences amongst all disk packs. However, the low-end market required the Data Modules to be removable, which in turn required a complex module load/unload mechanism in the drive which was far more complex and expensive than the cam ramp load mechanisms used in other contemporaneous disk pack disk drives.

The significant manufacturing advances were the automated processing of multiple heads with embedded sensors and programmable machines. This and other aspects of manufacturing are discussed in some detail in “Innovations in Disk File Manufacturing,” [1981Mulvaney].

The combination of low capacity point, high data module cost and complex Module load mechanism created the opportunity for the OEMs to respond with products based upon conventional and much lower cost disk pack technology – ultimately Control Data Corporation succeeded in the market with its SMD line of disk drives and disk packs, which dominated the non-IBM market until the early 1980s. SMD marks the first major departure from IBM storage media standards. Only Control Data and Nippon Peripherals, Inc., a Japanese government sponsored consortium, produced a media compatible 3340 drive. Several media vendors, e.g., Memorex, produced compatible Data Modules. There were no future Data Module products after the 3340; the industry continued with disk packs and then transited into fixed media.

In May 1973, the IBM 3340 was announced concurrently with and as the only disk drive for the IBM System/370 model 115 (a low-end system). It was also made available on higher-end systems but not exclusively - on such systems it competed with the much higher capacity IBM 3330-11 and PCM offerings. Due to the relatively high cost of the medium and its low capacity, it failed to gain any significant volume in higher-end IBM computing systems. Neither the 3340 nor the IBM low-end systems were big sellers, whether there is a causal relationship, either way, is a subject of speculation.

About the time the 3340 was introduced, IBM decided to move sales of disk packs and Data Modules from its Data Products Division (mainframes) to its Information Records Division (punch cards, media). Haughton was concerned, “what a mess that’s gonna be when we’ve just moved all the technology into the data module.” and set out on a "million mile" journey to find a solution to that problem. The solution was to have either organization's salespersons sell the products, but even then it remained a problem since the commissions on the higher priced 3330 were more attractive to the DPD salespersons. [2004Panel p. 20-21].

Jack Harker notes the low volume may have been a blessing,"… if we thought of it in advance, it was pretty good strategy that we got a low production product comparatively, into production with the 3340, and all its new technology, and ramped up, getting good yields, and then we moved on to the 3350, which was a high-volume product, using the same technology." [2004Panel p. 21]. The IBM Model 3350 turned the Data Module into a non-removable head disk assembly, which has remained the fundamental packaging concept of hard disk drives to this date. Some observers have called the 3350, 317 MB per Module, the real Winchester, since as a product it was a huge success for IBM and perhaps even more so for IBM's PCM competitors in their double density versions, 635 MB per module. It should be noted that the rest of the computer industry did not immediately follow IBM with fixed media, but instead continued to provide removable disk packs as the 1975 CDC 9676 (300 MB) and the 1983 DEC RA60 (equivalent to 400 MB)

The Winchester head concepts of low mass, low load, low flying height and an inductive transducer continued thru several generations, well into the 1990s, ultimately being replaced by much smaller heads having a different air bearing and transducer, see e.g. IBM Sawmill for the first of the next generation.

Head development details

Because of the sensitivity to head cost, a joint program was formed between the development team and the manufacturing team. Management and Engineers from both areas were moved to the IBM Menlo Park laboratory and the "We Team" was formed. Ken Machado headed up the team with Eric Solyst from Development and Bob Howard from Manufacturing as the key managers.

The IBM 3340 project was begun with the Data Disc tri-pad head design that was used in the initial work. The tri-pad head design consisted of a Y shaped barium titanate head with two taper flat air bearings at the leading edge of the slider and a taper flat bearing at the rear in which the ferrite recording element was glass bonded in place. A stainless steel flexure was bonded to the back of the head as a suspension to allow the head to gimbal while flying over the rotating disk.

Early on, problems surfaced with the tri-pad head. The cost of manufacturing, with the separate magnetic element and the challenge of generating the taper for the trailing air bearing pad, was insurmountable. The "We Team" concept ferreted out these problems early on. In addition, the inherently non-symmetrical shape of the tri-pad head created problems during high speed accessing by the servo.

At this point, work was started in coming up with an all magnetic ferrite and glass head. Several members of the "We Team" had been involved in the IBM Model 2305 program in which an eight element ferrite and glass fixed head was manufactured in a batch process conceived by Eric Solyst. The head was made of two pieces of magnetic ferrite machined and lapped and glass bonded together to form a magnetic core. It resulted in a head with a back gap with a large cross section and a front gap with a small cross section. This provided the necessary magnetic efficiency for reading and writing on a disk. The process for making these heads involved slicing, dicing, grinding, polishing and lapping. This was done with low cost equipment fitted out with very accurate computer control and measuring systems (glass scales or laser interferometers). It was the so-called "stiff finger" process because hundreds of parts could be loaded on the machine table and the button was pushed to start the automatic process cycle. After completing the cycle, the operator would unload the batch and reload for the next cycle. An operator could run up to four machines at a time. Yields were good and the process and equipment existed. The problem was that no suitable head design existed to use this capability for the 3340.

The assignment to come up with a new design was given to Mike Warner, an engineer who had worked on the IBM Model 2305. The initial design work focused on miniaturizing the head for low mass and to allow more heads per batch for low cost. The design breakthrough came during the air bearing simulation work. It was discovered that long narrow taper flat bearings when properly configured and loaded could be made to quickly fly off the disk during start up, thus reducing wear. It could also be designed to make the head pivot about the trailing edge of the taper flat bearing. The long narrow taper flat bearings provided a pressure profile with two peaks, one at the intersection of the taper and the flat and the second at the trailing edge. The pressure would bleed off due to side flow between the peaks. These peaks would respond to the disk vertical motion and the disk circumferential curvature to keep the trailing edge of the head at a very constant spacing.

Placing a taper flat bearing rail on either side of the head provided both lift and roll stability. By adding a center rail and placing the read/write flux gap at the trialing edge of this center rail, very good read/write magnetic gap to disk spacing could be maintained. The center rail was machined to the width of the desired recording track. The machining of the widths of the three rails was done on one of the stiff finger machines with a diamond cup wheel. This eliminated many costly steps required of a discrete core type head that had been used in previous disk drives. The design allowed a husky core to reside at the rear of the head for easy winding of a fine copper wire coil. This configuration provided a large back gap cross section and a small focused front gap for good magnetic efficiency. With minor modifications, the 2305 batch process could be used for the 3340 head production. The head thus conceived met the 3340 requirements.

A stainless steel suspension with integral load spring and a spot welded load beam was clipped into a notch machined in the back of the head. A dimple was formed in the suspension and a load beam pushed against the dimple to provide a controlled load point location at the center of mass of the head. The suspension load beam structure was compliant in the Z, pitch and roll directions yet stiff in the X, Y and yaw directions. The load was adjusted by laser heat to eliminate creep from mechanical bending. This process and equipment was developed by Richard Kurth. A channel down the center of the load beam was formed for the coil wires encased in a tiny plastic tube. Tabs were also provided to capture the tube and wires. This channel also stiffened the load bean allowing a thin stainless material to be used. The suspension components were all made of photo-etched stainless steel provided in strips to allow the batch process philosophy to be applied to the suspension as well as to the head fabrication. Small holes in the etched material were used for tooling and alignment. The suspension was designed by Dick Wilkenson and Mike Warner. After some revisions, this symmetrical head suspension assembly proved satisfactory in the high-speed servo system.

The stainless steel suspension assembly was spot welded to a small stainless steel mounting plate that had a spud in the center that fit in a hole in the arm. The head suspension assembly could be individually tested prior to mounting on the four-headed arm assembly. The attachment to the aluminum arm was via a swedging process. A tool was pushed through the spud expanding the stainless steel beyond its yield point and pushing into the aluminum arm. The aluminum, with a lower modulus of elasticity, would not yield but would tightly grip the spud. This made for a low cost attachment of heads to arms and eliminating screws or glue. This also allowed head suspension assemblies to be secured to both sides of the arm in a single stroke of the tool through the back-to-back mounting plate spuds. George Pal did the design for the attachment process.

A cam surface was formed on part of the load beam to aid in inserting the arm on the actuator between the disks. A “pickle fork” tool was inserted in slots on either side of the arm and contacted the cam surface holding the heads in a retracted position in the arm. After inserting the arms between the disks in a Data Module, the pickle fork was gently withdrawn, lowering the heads onto the surface of the disk.

A thin circuit board on which a diode selection matrix was mounted was attached to the aluminum arm. Mounting electronics on the arm and the development of the Module was managed by Jack Swartz. A cut out was provided in the arm to receive the electronic module and keep its mass at the arm center and to reduce the arm profile. The diode selection module allowed the number of wires in the cable to be reduced thus reducing drag during carriage accessing. The head leads were soldered to the circuit board out on the arm and cable was attached to the circuit board near the base of the arm. Stainless steel springs surrounded the cable from the arm to the connectors in the Data Module. These springs supported the plastic cables when mounted in the Data Module. The cables were grounded to bleed off any static charge generated by the plastic cable during accessing.

One of the concerns of the program management was the vulnerability of the center (narrow) rail to damage by the disk due to curvature along the radius of the disk. This curvature would result in the center rail touching the disk when the disk was convex during starting and stopping. When the curvature was concave it would increase the spacing between the head and the disk when reading and writing while flying. After many measurements and testing it was determined to be a second order effect.

Quotes

- Haughton: "Yeah. They were all removable, and they were right from the start" 2007Haughton, p15
- Haughton: "…And the cost was a very serious issue, probably the largest jeopardy, because, again, we started out for low-end systems. And they were the ones that were going to use it, and they couldn't afford to have too much invested in data storage. So the cost was the biggest point from that point on, …" 2007Haughton, p15
- Haughton: "IBM did license the Data Disc head, and Joe Ma built a video file in Advanced System Development division using that head and 24-inch disks. Did this under contract with the Rand Corporation, and it used those heads. That's where we started with the Winchester as well …" 2007Haughton p11
- Haughton: "we were defining the product, trying to end up with what would best fit the market requirements for the low end systems …" 2004Panel p3
- Warner: "Our target was to make a head suspension assembly for less than a dollar. And at that time, a head suspension assembly was then more like $15 to $18. So we had a dramatic cost reduction that had to take place." 2004Panel p5
- Warner: "Well, anyway, we were making these tri-pad heads, not having very good success with them, and so we started working on using some of the ideas we had for the 2305 or Zeus head." 2004Panel p8
- Solyst: "…when we first started on the Winchester program it was supposed to be a recording head in contact with the disk. And some of our engineers had designed a head that was supposed to stay in contact with the disk and could change a single element. And this head was loaded with just two grams and that very low load was necessary to eliminate wear on the disk. Well when we got started and made, you know, a couple of dozen of these heads, it turned out that they operated very erratically and we, you know, nailed it down immediately that it was because it did indeed not stay in contact. Once the disk got up to speed, some head would be flying as high as thirty-five micro inches and some would be hopping and skipping on the disk, which was not a very good situation. Well, it turned out that these in contact had some very small flats running on the disk, and once in awhile the edges would be worn down a little bit or rounded a little bit during the lapping process and then would actually result in a sort of a taper flat air bearing and therefore some heads were flying up to thirty-five micro inches. Once we discovered that, we said well, let's not fool with that. Let's design a flying head to fly in between these two values. And that's how we came up with eighteen micro inches. And I believe that the load proper was what, sixteen, fifteen grams?
  …
  Warner: Well remember, Erik, when we first started out, we had a tripod head?
  Solyst: Yeah, well that's the one I was talking about. [2005Panel p.12-3]
- Warner: "…These early guys were -- you notice are four heads per IBM arm and this -- what happened outside of IBM very quickly was this technology was adopted for eight inch drives and five and a quarter inch drives at one head per arm.
  Frank: And three and a half inch drives.
  Warner: And using the same technology also with oxide media, but also with thin

film media.
…
Solyst: Except IBM.
Warner: Except -- yes, IBM was slow to pick up on that. A long story, but the interesting thing is these worked extremely reliably. The smaller the disk size goes, the reliability goes up and the sensitivity to contamination actually drops down. …" [2005Panel p.16-7]

- Frank: "Now beyond this, I mean obviously ferrite heads lasted a lot longer. This came out in 1976, as I mentioned, and in 1979 IBM came out with the first thin film head drive, the 3370. And of course there was a lot of consternation that it had taken us almost until 1976 to be in production with these things and start to make any money and all of a sudden here was a major change in technology and of course we and the other independent suppliers went chasing that. But it turns out that that was a difficult move to follow, particularly all the semiconductor style processing technology. And in reality, most of the rest of the industry continued to use ferrite heads for the bulk of their requirements for many years and, in fact, Applied Magnetics and I think the other independent suppliers were still making ferrite heads well into the 1990s. I can remember it must have been about 1993 we had this -- finally made the call to stop making ferrite heads and put all the effort into continuing to build thin-film head volumes. And I think for several years after that that was widely regarded as a mistake. But between 1976 and 1992, the reason we were able to keep ferrite heads going is that a ton of innovations were made." [2005Panel p. 18]
- Croll: "And as a little aside, as the manager of that program two element thin film head, I kept getting visitations from our corporate headquarters in Armonk saying, "Why are you fooling around with film heads when they have all these good things?" But the thing that really I think did the ferrite heads in was the magnetoresistive ones which definitely got the advantages of the two." [2005Panel p. 21]
- Frank: "Oh easily, easily, yeah. Yeah, you know, the volumes we got up to certainly weren't like thin film head volumes today which are, you know, millions a day, but, you know, I think there were probably years when we made, well several hundred million heads I would guess probably.
  Solyst: Yeah, it's got to be at least that.
  Warner: And I remember that, Erik, when you and I don't know who else, Ken Mishado and others had to go -- I think the projected number for Winchester heads was gonna be like seven hundred thousand heads.
  Solyst: I don't recall that number.
  Warner: And they were thinking about should we proceed on with this technology? Can we afford to do that? And so the projected numbers were so low at the time. And obviously, you know, those were way, way understated values. It would have been a shame if it had all been killed because of that. [2005Panel p. 23-4]

Additional information

[1974Mulvaney] R. B. Mulvany, "Engineering Design of a Disk Storage Facility with Data Modules," IBM J. Res. & Dev., vol. 18 (1974).

[1977Kean] Kean, D. W., "IBM San Jose, A Quarter Century Of Innovation," 1977, Computer History Museum # 102687875.

[1981Mulvaney] R. B. Mulvany and L.H. Thompson, "Innovations in Disk File Manufacturing," IBM J. Res. & Dev., vol. 25 (1981).

[2004Panel] IBM 3340 and 3350 Disk Drives Oral History Panel, Computer History Museum, April 6, 2004. Computer History Museum # 102657933.

[2005Panel] Oral History Panel on Advanced Ferrite Disk Heads, Computer History Museum, November 3, 2005. Computer History Museum # 102657927.

[2007Haughton] K. E. Haughton Oral History, Computer History Museum, February 14, 2007. Computer History Museum # 102658029.

IBM Archives 3340 Site

Moderator: T. Gardner

Contributors: M. Warner