History of Disk-File Development at Hursley and Millbrook

October 17, 1990

W J P Case

ID (Information Development)
MP 095 Ext 5524

Contents

1.0 Introduction . . . 3
  Additional information . . . 3

2.0 Table of Acronyms . . . 5

3.0 Product shipments . . . 7
  Patent activity since 1967 . . . 9

4.0 A small beginning . . . 11

5.0 Dolphin (5444) . . . 13
  Access . . . 16
    Friction drive with 250 ms average access time . . . 16
    Stepper motor with 125 ms average access time . . . 20
  Data channel . . . 21
  Surface analysis testing (SAT) . . . 22

6.0 Gulliver (62GV) . . . 23
  Track following . . . 23
  Access . . . 26
  Data channel . . . 27
    Basic specification . . . 27
    Encoding method . . . 27
  SAT . . . 27

7.0 Piccolo (62PC) . . . 29
  Encoding of sector and index . . . 29
  Power-on cycle . . . 30
    Requirements . . . 30
    Possible problem . . . 30
  Track following . . . 31
  Accessing . . . 32
    Calibration . . . 34
    Gain . . . 34
  Data channel . . . 34
  SAT . . . 34

8.0 Sprat (later Bluegill then Flotilla) . . . 35
  Basic specification . . . 35
Sprat at Millbrook January 1975 to December 1977 . . . 35
  Introduction . . . 35
  Outline . . . 37
  Head flying and wear . . . 40
  Backplates . . . 42
  Mathematical analysis of disk-flying . . . 43
  Head fabrication . . . 44
  Data channel . . . 46
    Introduction . . . 46

    - Concepts and terms 47
  - History 48
    - Progress during 1975 48
    - Progress During 1976 49
    - Progress during 1977 50
    - The Position in February 1978 51
  - Mathematical modelling of the record and read processes 52
  - Microprocessor control 53
    - Background 53
  - M6800 software preperation 54
  - Software and hardware for track-following 55
    - Control-system architecture and hardware 55
  - M6800 software 57
  - Devkit microprocessor application system 57
  - Servo system 58
  - Track-following 58
    - Background 59
    - Early work with analog controller 59
    - Digital servo 59
  - Track seeking 60
    - General method 60
    - Early work with analogue controllers 61
    - Digital scheme 61
  - Servowriter 62
    - Introduction 62
  - History 63
  - Burnishing 65
  - Disk surface analysis 66
    - Purpose 66
  - History 67
  - Mechanical design – machine and cartridge 67
  - Store-testing robots 68
  - Wear-test rigs 69
  - Track misregistration study 70
    - Conclusions 71
- Bluegill at Boulder and Tucson January 1978 to September 1979 72
  - EVT Experience 72
  - AVT design 73
  - Users 75
  - Staffing 75
  - Architecture 76
    - Format 76
    - Command Set 77
  - Mechanical design 78
    - Drive 78
    - Cartridge design 78
  - Microcode 79
    - Control code 79
    - Microcode tools 80
  - Control logic 80
  - Microprocessor hardware 82

Data channel 84
Motor drive 84
Read-write heads 85
Recording medium 86
Track-following and access 86
Servo-writer 87
Track misregistration (TMR) studies 88
Testers 89
Cost estimates 90
    Product cost 90
    Cartridge cost 90
    Maintenance cost 90

9.0 Swallow (62SW) 91
Track-following 91
Accessing 92
Data Channel 92
SAT 92

10.0 Woodwind (3310) 93

11.0 Flotilla 95

12.0 Harp (later Heron then Falcon) 97
Introduction to the Falcon servo 97
Accessing 97
    Access types. 100
Track-follow 101
    Compensator 101
    Low frequency integrator 102
    Filter 103
The power-on cycle and calibration 103
Servo control-loop block diagram 104

13.0 Falcon GT 105

14.0 Kestrel (9335) 109
Configuration 109

15.0 Simulator and Matchbox (1984) 113

Figure 1. Front of Hursley House

1.0 Introduction

When the department was first set up in Hursley to develop disk drives in 1966 it was called Storage Development; its name was changed to Storage Products in 1985.

Additional information

This document contains the history of the major developments in this department since 1967.

Details of all patents raised by Storage staff are included in another document. This document is called PART SCRIPT in the PART library on disk PBISIL 237 on WINVMK. PART SCRIPT contains:

- Abstracts of patents with inventors' names
- A chart of all personnel showing their time in Storage
- A description of all projects in Storage
- "Stories" that have circulated during this period, included for interest.

2.0 Table of Acronyms

The following acronyms are often used in disk-file devopment:

ADC	Analog-digital converter
BAT	Basic assurance test
BDLS	Basic design language structural
CD-ROM	Compact-disk read-only memory
CKD	Count key data
CRC	Cyclic redundancy check
DAC	Digital-analog converter
DAD	Device adapter device
DASD	Direct-access storage device
DCP	Digital control processor
DFCI	Device file control interface
DRAM	Dynamic random access memory
DSP	Digital signal processor
EC	Engineering change level
ECC	Error checking and correction
EDS	Engineering design system
EIA	Electrical industries association
EPROM	Electrically programmable read-only memory
ESD	Electrostatic discharge
FCM	File control microproccesing
FRU	Field-replaceable unit
GPD	General products division
GSD	General systems division
HDA	Head disk assembly
IDC	Insulation displacement cable
IPI	Intelligent peripheral interface
LBA	Logical block address
LED	Light emiting diode
LSI	Large scale integration
LSSD	Level-sensitive scan device

MFM	Modified frequency modulation
OPD	Office products division
PNEUC	Physical non-elementary unit code
PES	Position error signal
PD	Physical design
PLD	Power line disturbance
PWM	Pulse width modulation
RAM	Ramdom access memory
RMS	Rochester master slice
ROM	Read-only memory
SCSI	Small computer system interface
SPD	System products division
SRAM	Static random-access memory
TPI	Tracks per inch
UTS	Universal test system
VCM	Voice coil motor
VGA	Variable gain amplifier
VLSI	Very large scale integration
VPD	Vital product data

3.0 Product shipments

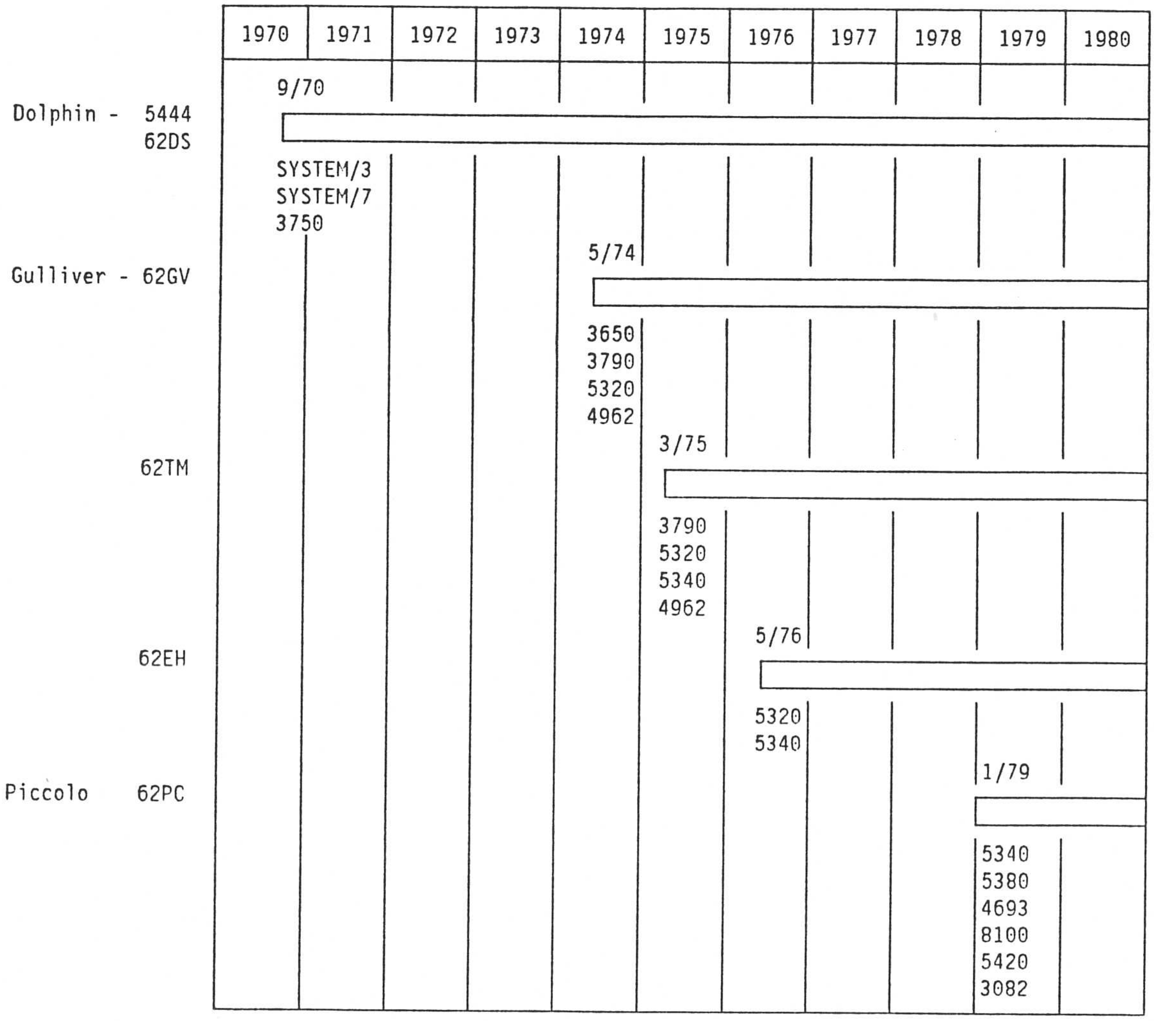

Figure 2. Shipments between 1970 and 1980

This chart has been made from a chart that John Taylor had in 1980. Information about the use of these products after 1980 is not available.

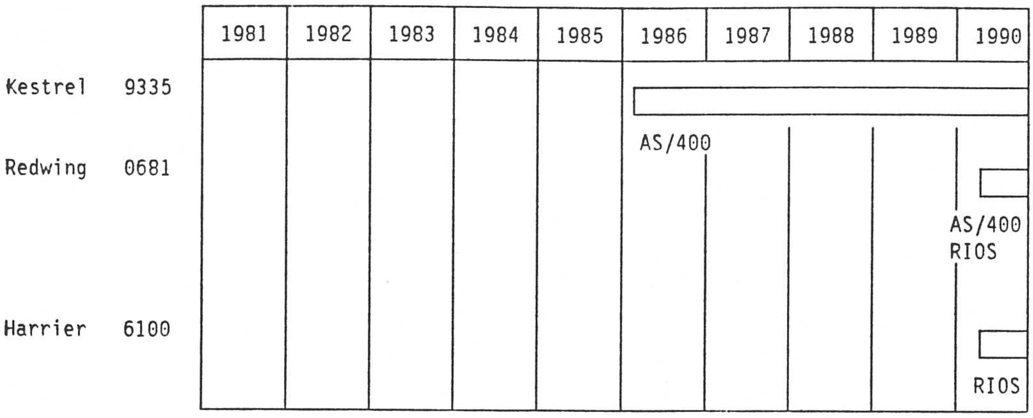

Figure 3. Shipments between 1981 and 1990

Patent activity since 1967

5444 (Dolphin) (1967 – 1969)	First fixed + removable DASD First step-motor-driven access 5440 Cartridge with recessed release handle	5 patents
62GV (Gulliver) (1970 – 1973)	Winchester technology First sealed HDA to include actuator Unique rotary VCM Subject of Corporate Outstanding Innovation Award.	5 patents Widely copied
62PC (Piccolo) (1973 – 1976)	First 8-inch DASD Volumetric efficiency from servo technology extensions	5 patents
Sprat/Bluegill (1974 – 1979)	Proved product feasibility of Bernoulli file with cartridge, contoured back plate, and sector-servo access control	11 patents
Flotilla (1982 – 1983)	Restart of Sprat/Bluegill	No patents at August 1990
Swallow (1979 – 1982)	Inrease of density both in BPI and TPI	At least 1 patent
9335 (Kestrel) (1984 – 1990)	IBM's first IPI-3 controller (A01) DASD design for robotic ASM (B01) Volumetric efficiency increased	11 patents
0681 (Redwing)	5.25-inch footprint OEM release	No patents at August 1990

4.0 A small beginning

This section was written by Chris Wallis and John Heath.

The birth of what we now know as Storage Products was in the spring of 1966, when John Herbert formed a small team of engineers to develop a "chip" file. John had recently returned from an assignment at the IBM Development Laboratory in San Jose, California. It was a time of rapid technological and business growth at that laboratory, based on the success of the disk file invented there ten years earlier. San Jose is in the Santa Clara valley, now popularly known as Silicon Valley, and the explosive economic growth based on technology caught John's imagination. He brought back a vision of seeing the same thing happen in Hampshire.

It happened that San Jose had more work than they could do themselves, and this provided an opportunity to start something at Hursley, where there were already some electro-mechanical skills. The idea of a chip file derived from a large photographic archival store that had been developed by IBM for the U.S. Government. This device, called Cypress, stored millions of photographic images the size of credit cards in small boxes, which it stowed away in miles of storage racks. The idea then was to develop a unit that could use magnetic cards (chips) that could be stored in the same boxes and be archived the same way but which could store digital data. The project was called Zenith, and Hursley gained a part of the task related to a low-cost, read and write station, Zenith LC.

The chips were pieces of Mylar 35 x 70 mm coated with iron oxide. About 32 were loaded into each box, which was about 38 mm wide, 80 mm deep, and 25 mm long. The boxes from which lids had previously been removed were lined up in a row of eight on a table that was positioned by a stepping motor.

A selected chip was positioned by the stepping motor over a "kicker" that acted through a slot in the bottom of the box, and kicked vertically upwards into the start passage leading to the read station. The passage, which was wide enough to accommodate the 35-mm width of the chip in one dimension, and its thickness in the other, had several air jets that blew the chip along the passage to the read station.

The read station had a drum about 80 mm in diameter, against which the chip was retained by air pressure. Inside the drum, a rotating head scanned the chip (positioned laterally across the chip by a stepping motor), and rotated continuously around the periphery of the drum to scan the chip along its length. The head had to be coupled via a transformer to avoid twisting its wires off. After the access, the chip was returned the way it had come by reversing the airflow.

The average access time to get a box in position was 288 milliseconds, and the time to move a chip to the read-write station and then replace it was 134 milliseconds. The average access time including write and readback check was 730 milliseconds.

The storage capacity of this "high-tech" device was 50 tracks per inch and 800 bits per inch, giving a capacity of 610KB for each cell and 5MB on the table.

Between May and November 1966, Hursley took the idea from a concept to a set of partly working prototypes that could be linked together. Boxes were moved to and fro, and chips kicked up into the read station. The reading and writing of data were tried using circuits and components provided by San Jose.

There were plenty of problems, such as wear of the chips, sticking of the chips in the mechanism, and failure to recover the data (a common problem in storage devices as Hursley was to discover over the next 23 years). However, late in 1966, the advantages of disk files and tapes over all the competing alternative options were being recognized, and the chip file concept was dropped. Storage had a very interesting and practical lesson in the development of storage devices, but the vision was no nearer realization.

5.0 Dolphin (5444)

This section was written by Bill Case alone from memory.

First customer shipment (FCS)	September 1970
Capacity	5MB
Average access time	250 ms with friction drive, 126 ms with stepper motor
Tracks per inch (TPI)	100
Bits per inch (BPI)	2200
Data rate	200KB per second
Shipments to February 1990	60 000

Two 14-inch disks, one fixed and one removable, in a 5440 Cartridge.

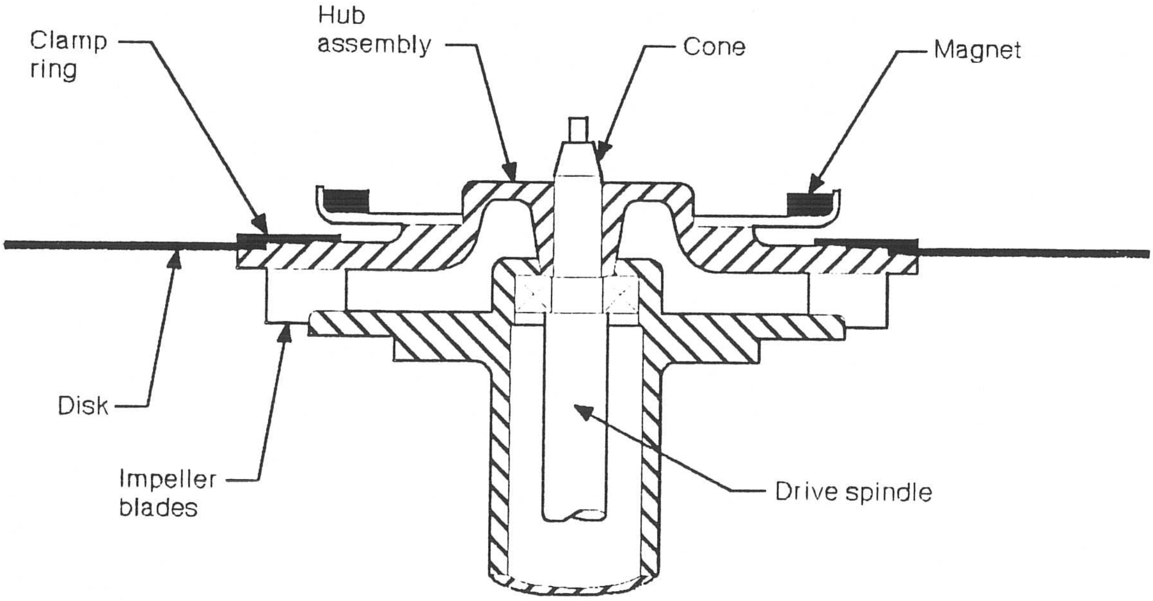

Figure 4. Fixed-disk assembly

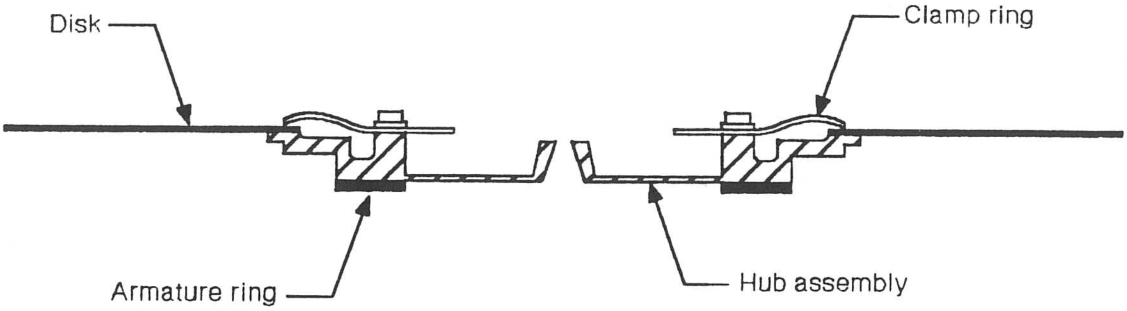

Figure 5. 5440 Cartridge assembly

Total capacity was 5MB with 2.5MB on each disk. The track density was 100 tracks per inch.

Both disks were driven by a single removable belt arranged as shown in the following figure.

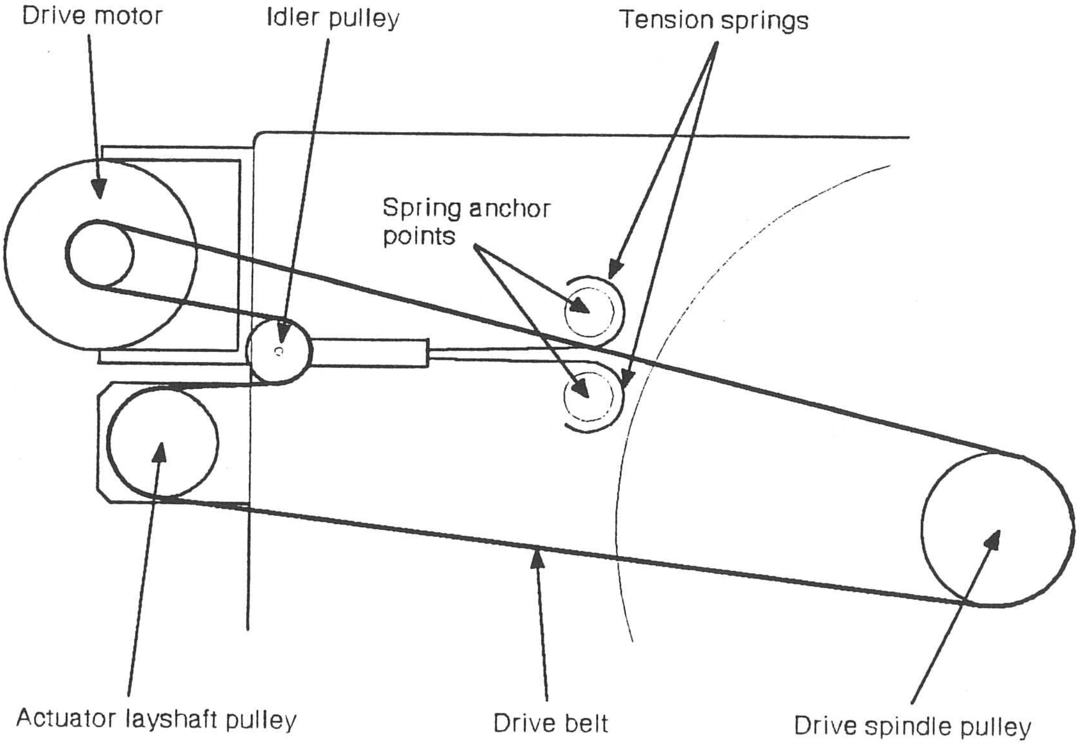

Figure 6. Drive-belt arrangement of 5444 Disk Storage Drive

The heads were the same as used in the 2314 Direct Access Storage Facility made in San Jose; a circular piece of ceramic material with a slot cut for the laminated read-write head and two circular air-bleed holes to reduce the air pressure at the flying surface.

All four heads were mounted on a block attached to a nut that was moved by rotating a leadscrew. This leadscrew was difficult to design as it was the most likely component to have wear problems. Any wear would mean that the actuator would not position the heads over the centreline of a desired track. The heads had to be correctly positioned over tracks by using a CE disk that had tracks accurately written on it. The position of the fixed-disk heads were never adjusted.

Access

The rotation of the leadscrew was by two different modes when the access time was reduced in a second release.

Friction drive with 250 ms average access time

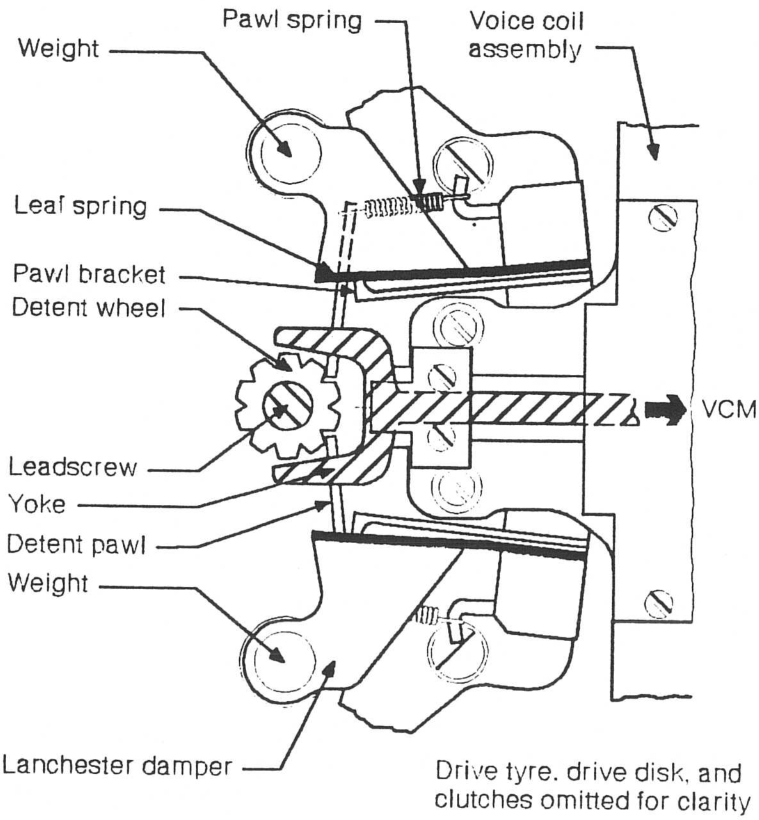

Figure 7. End view of friction-drive actuator showing detent wheel and Lanchester dampers

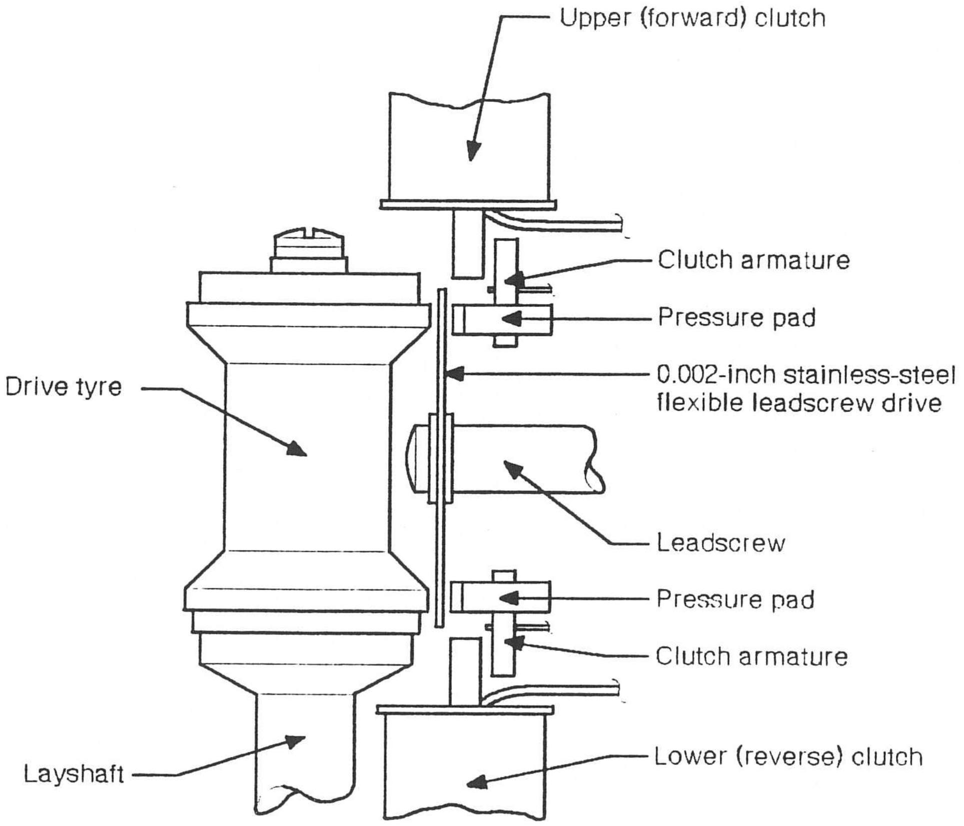

Figure 8. Friction-drive actuator detail

The leadscrew had a 0.002-inch-thick piece of circular stainless steel mounted at right angles to the leadscrew on which ran the actuator with the read-write heads mounted on it.

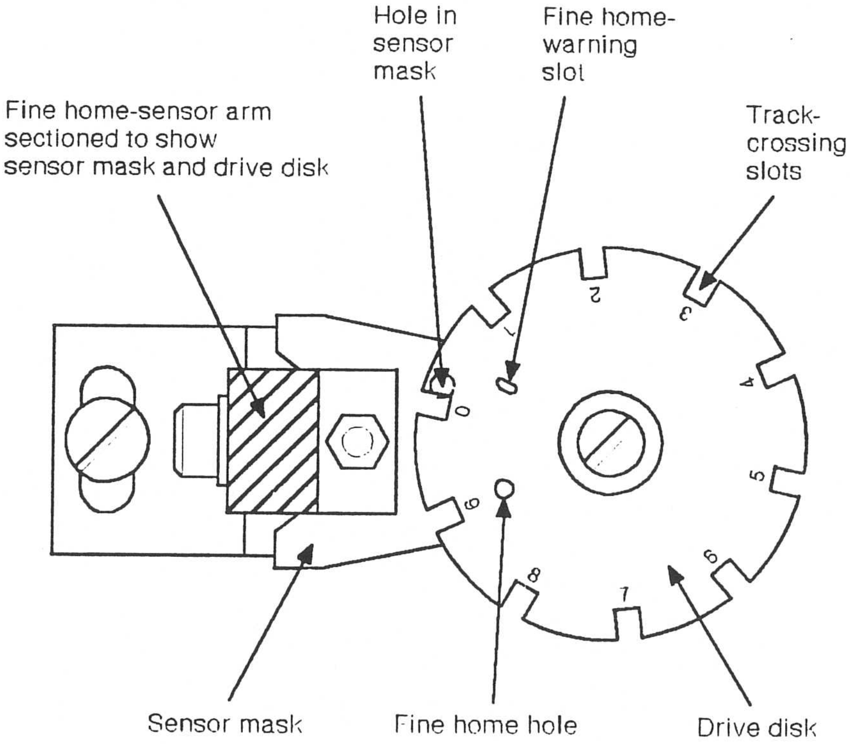

Figure 9. Detail of flexible disk in 5444 Disk File Drive

Closer to the leadscrew from the flexible stainless-steel disk was mounted a detent wheel with 10 positions that could be stopped (as Figure 6 on page 12). These positions were for each of the tracks on the magnetic disk. The track density of Dolphin was 100 tracks per inch, and so there were 20 turns of the leadscrew to move across all 200 tracks. Further away from the magnetic disks, was a spindle which was permanently rotated by the same belt that rotated the magnetic disks. Mounted on top of this spindle was a "cotton reel"-like drive tyre. There was a clutch above and another below the centre of the drive tyre. When either of the clutches was energized the stainless steel flexible disk was brought into contact with either the top or the bottom of the drive tyre. This rotated the leadscrew either clockwise or counterclockwise giving an in or out access.

The access procedure was to first disengage the detent pawls (which held the actuator stationary): this disengagement used solenoid driving a double hook arrangement that extracted the detent pawls. The top or the bottom clutch was activated to push the flexible disk into contact with the top or bottom of the drive tyre which then produced the required direction of rotation of the leadscrew to make an in or out access. The detent wheel that accepted the detent pawls made the actuator stop at the appropriate track. The flexible disk apertures were counted to determine the track number.

When the released pawls returned to the detent wheel one of them, sometimes, would bounce backward into the supporting leaf spring. When this spring returned to its rest position it sometimes caused the actuator to go back one track. This was cured by mounting a Lanchester damper on the spring.

The heads were held on the actuator apart from each other and had to be loaded into the flying position. The load force was 34 grams for the top heads and 42 grams for the bottom heads. The difference in load force was the effect of gravity. The loading action was done with by a Bowden cable connected to a solenoid.

Stepper motor with 125 ms average access time

A stepper motor was used to reduce the access time of Dolphin and to improve reliability. The 200-pole stepper motor was mounted on the end of another leadscrew, with five turns only. This number of turns was needed because the motor was arranged to be on the track only at one of four detent positions.

By advancing the phase switching to the motor it was possible to go far faster than the single-step speed of the motor. The head load for this drive was different from the friction drive. To load the heads, the actuator forced a bearing to run up an incline which caused the heads to be loaded. Once loaded they were held locked into the load position. The problem occurred when the heads were unloaded.

The heads were initially unloaded very quickly. Intolerable head crashes occurred because the heads would bounce back too far from their end stops. The solution was to reduce the unloading speed by including a simple dashpot. This system required no mechanical detent and lessened wear.

Data channel

The data channels on both versions of Dolphin were identical. The recording scheme for Dolphin was double-frequency recording: clock information was recorded in addition to the data.

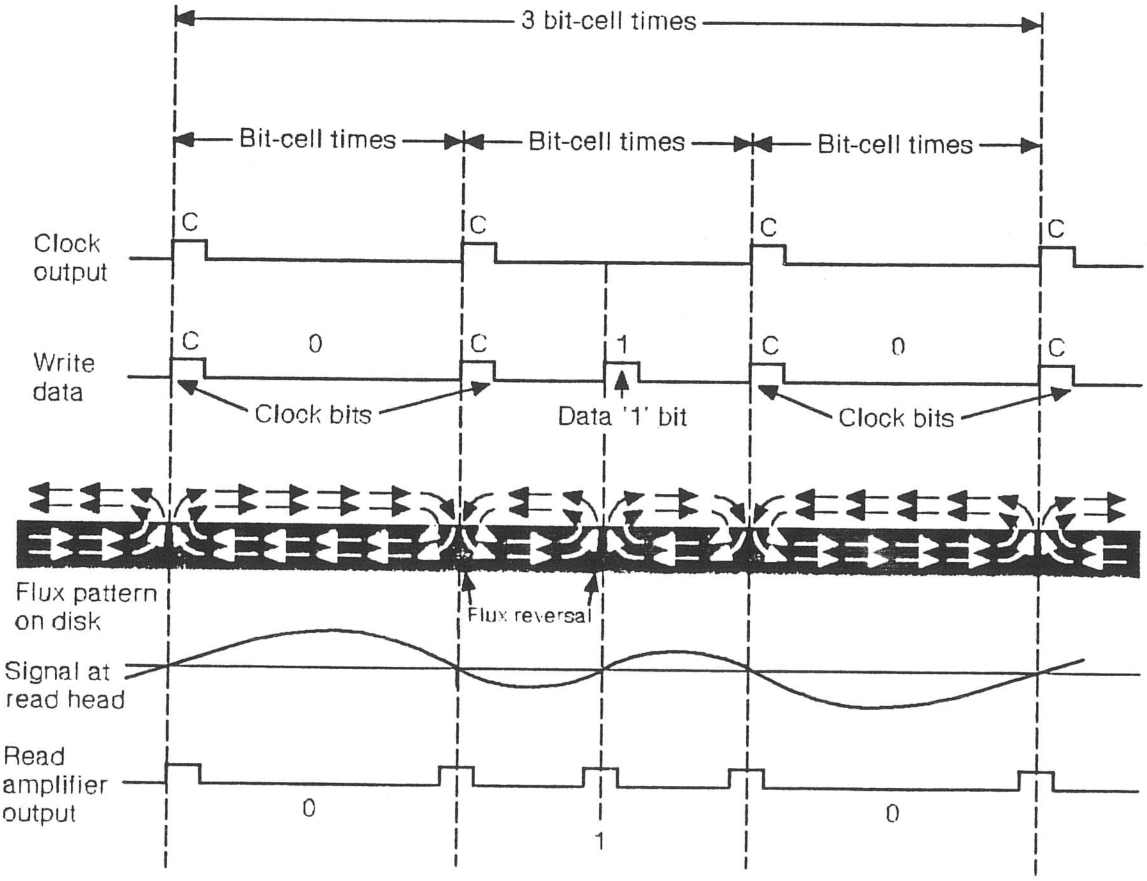

Figure 10. Waveforms in double-frequency recording

Imagine the track to be divided into equal lengths called bit cells. At the boundary of each cell a transition was recorded for a clock pulse to locate the cell, so that the data, which was recorded at the centre of the cells, could be extracted. A special area, called an address mark, was recorded in such positions to define the end of the fixed-length sectors. The detection of the data and the clock pulses was by amplifying and limiting the raw readback signal. This limited signal was then differentiated to give pulses from each edge of the limited signal. The pulses in one direction were detected as either clock or data pulses. The address marks were written with missing clock pulses that were captured with WOPAC, the invention of Chris Wallis and Charles Owen hence the name (Wallis Owen PACkage).

This worked as follows. A ramp was generated that was reset when a clock pulse was detected. A threshold was made above the centre of the ramp about three-quarters of the way through the cell. If a pulse was detected below this threshold it was taken to be a data pulse. A second level was used to measure the ramp and if it went above this threshold it was a missing clock pulse. If this occurred twice in adjacent cells it was then known to be an address mark (AM) and the data area in a given area after that. The sectors were of fixed length, 128 bytes. The output from the data channel was sent to the system where the clock pulses were removed from the data pulses to provide the customer's data.

During the development of the data channel, Hursley found an error in the design based on the 2314 Storage Control. The basic channel was in three parts: filter, limiter, and detector. The filter was a simple Bessel filter, to create maximum flatness over all of the pass band. The limiter was where the problem existed: the limiter was an increasing-gain stage made up of a series of long-tailed pairs. In the final stage, the long-tailed pairs would switch. This should have worked satisfactorily but the final stage was expected to run with no current in one of the legs, making the switching much slower than it should have been. Hursley redesigned this circuit to overcome the problem.

Surface analysis testing (SAT)

The SAT process was the first for Storage. The output from the SAT process was to mark the defective sectors so that they would not be used to store customer data. The SAT process made an average of the amplitude of the readback signals and then created a percentage (60%?) to use for detecting a missing bit. Also an extra bit was detected if the amplitude of pulses from the erased track was greater then n% of the track average. This would mark the sector as defective and thus not available for customer data.

6.0 Gulliver (62GV)

This section was written by Bill Case from memory with help from John Taylor's files.

First customer shipment (FCS)	May 1974
Capacity	5MB, 10MB, or 14MB
Average access time	70 − > 40 ms
Tracks per inch (TPI)	167 − 300
Bits per inch (BPI)	5650
Data Rate	0.9MB per second
Shipments to February 1990	177 000

Track following

The Gulliver drive was the first in Hursley to have servo-controlled positioning. It was decided to close a servo loop around the actuator, and a feasability trial was made. The idea that a rotary actuator would give much greater efficiency led to the first rotary actuator in the world being developed at Hursley. The first experiment used a cigar tube attached to a coil (forerunner of the VCM) at one end, and to a Moire-fringe assembly on the other side of a pivot. The Moire fringe was used to produce a position signal. The initial loop was closed at 9 Hz! This was a very low frequency compared with the present day, but was the first ever closed-loop rotary actuator. Much work was spent on producing the final version of the actuator, which had an open-loop crossover frequency of about 300 Hz.

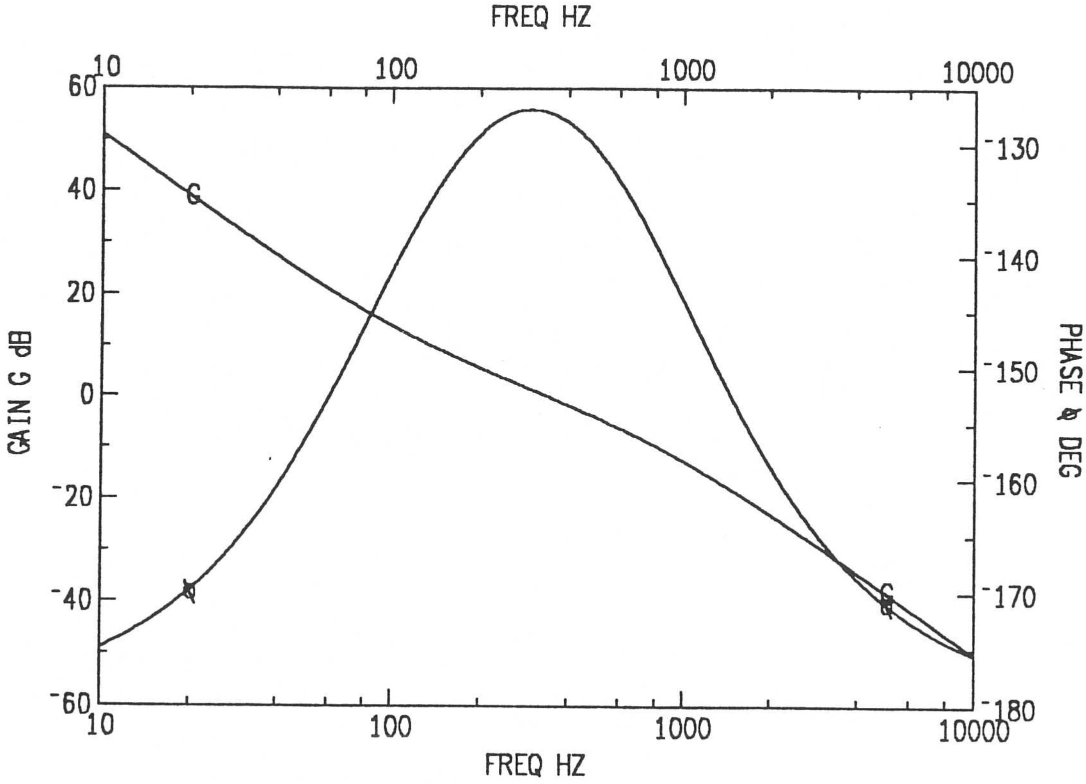

Figure 11. Gulliver open-loop transfer function

This initial idea of having a rotary actuator was so clever that the inventors received a corporate invention award, and the actuator has been copied by many other makers of disk files.

The final actuator was made from two flat plates with stiffening pieces. Each of the two arms supported two heads. The outer movable head on the lower surface was used as a servo head, to read specially written tracks to give a continuous error signal. The heads were what is now called "Winchester" heads, from San Jose. It was the first time that these heads had been flown at a skewed angle, as until then the heads had only been used on a linear actuator.

The track-following loop was developed using simple theory as expanded by Thaler and Brown in their Servo Design book. All that was needed was to close the loop to get a phase lead at the unity-gain point. This required design of a phase lead network that would have phase lead at the unity gain point. This point was chosen as 300 HZ and so the design was simple. The need for such an apparently high bandwidth was because of the need to keep the head over the centreline of the written tracks. All the track following and accessing was done with information from the servo tracks. The position loop was designed to position the data head at all times within 700 microinches of the track centrelines. However, many of the position tolerances were outside the servo loop and therefore only 170 microinches were allowed for positioning error within the loop. The external causes were:

- Demodulator drift
- Error amplifier drift
- Drive offset
- Demodulator noise
- Servo residual settling
- Servo signal modulation.

The position loop had also to be able to follow disk eccentricity with very small phase shift and be able to lock in from a wide range of velocities following completion of an access. The loop was a standard form of a type-two servo system. This meant that there were two direct integrations in the loop. These arose because the system was driven by passing a current through the coil of the VCM, which produced acceleration and hence motion of the data heads.

The phase shift associated with the two integrators is −180 degrees: hence, if direct position feedback were used with no compensation, the loop would have been unstable. A phase-lead network was therefore introduced into the loop to reduce the phase shift at the point where the open-loop transfer function had unity gain.

The initial capacity of the Gulliver file was 5MB. There were two heads on the actuator on the top surface and 8? fixed heads on the inner half of the lower surface. The rotation speed was 2964 rpm using a standard ac motor connected to the disk spindle by a belt.

Access

There were four equations for the access time for different seek lengths. In the following equations N represents the number of tracks moved.

1 track	14 ms
2 tracks	17 ms
3 – 7 tracks	$17.92 + 1.77x(N-3)$ ms
>7 tracks	$22.91 + 1.02x(N-7)$ ms

The average access time for the movable heads was 70 ms.

The position error was only measured with a simple position-error signal (PES) of single phase, so that when not close to the on-track position the PES was not a good representation of position. The differential of the PES signal was used, when close to an on-track position, as a measure of the head velocity, but was augmented by integrating the coil current and mixing them together to make what was known as the hybrid velocity signal. This hybrid signal was always used when close to the on-track positions, but during the rest of the time the integrated coil current was used on its own.

For seeks of more than seven tracks the following was the sequence of events:

1. Accelerate in the appropriate direction for three tracks.
2. Keep the velocity achieved at this point in a 6-bit analog-digital converter as the desired velocity.
3. Compare the actual velocity with the desired velocity during the rest of the constant-velocity phase of the access. (This velocity was about 1.02 ms/track.) This comparison was fed into a 'bang-bang' servo to drive the actuator faster or slower and maintain the desired velocity, only stabilized by the time constant of the VCM.
4. With only five tracks remaining, reduce the velocity to a level from which the actuator could be captured without undue overshoot onto the desired track.
5. On reaching the desired track, disable the bang-bang servo. The position loop then controls the current to the VCM.

To initially get to track 0, which is close to the position where the head is parked, the actuator moves the heads out until three track crossings are measured. These are measured as the PES being less than 10% of its peak amplitude. The actuator is then controlled to maintain the capture velocity until the third on-track condition is sensed when the position loop is activated with the head on track 0.

Data channel

Basic specification

Nominal rotation speed	2 964 rpm
Data rate	889.2KB per second
Data tracks per moving head	163
Track capacity (gross)	18KB
Sectors per track	60
Data bytes per sector	256
Customer-accessible capacity	2.5MB per head = 5MB

Encoding method

The encoding method used for the data channel was modified frequency modulation (MFM). If the data is considered to be constituted of bit cells then, in terms of magnetic transitions, the encoding is defined as:

- All "ones" in the data are encoded as transitions in a mid-bit cell.
- All "zeros" in the data (with the exception of a "zero" immediately following a "one") are represented by a transition at the start of a bit cell. The "zero" immediately following a "one" has no transition associated with it.

An example of a data sequence that shows all the laws is :

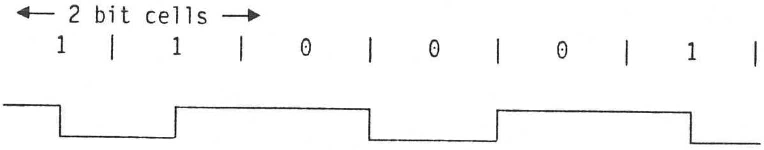

SAT

The disk surfaces underwent surface analysis testing (SAT) to mark the sectors or tracks that were defective. The defects could be defined as:

- Skippable

  Such a defect could be up to 16 bytes long. The contents of the sector would be rewritten into an alternative sector.

- Flaggable

  A flaggable defect was longer than those that were skippable. In this case the whole track was reassigned to one of two alternative tracks.

7.0 Piccolo (62PC)

This section was written by Bill Case from memory with assistance from John Taylor's files.

First customer shipment (FCS)	January 1979
Capacity	65MB
Average access time	27 ms
Tracks per inch (TPI)	450
Bits per inch (BPI)	8530
Data rate	1MB per second
Shipments to February 1990	360 000

Piccolo had a rotary actuator with six 8-inch disks. One of the surfaces was used as a dedicated servo surface. The total capacity was 65MB. The eleven remaining heads stored the users' data and had servo sample information stored with each sector of data. This meant that about 16% of the disk surface was used for storing non-customer data. The following four sections explain the various functions.

Encoding of sector and index

It was decided to encode sector and index pulses by having a sequence of missing clock pulses at the beginning of the quadrature cells. Counter 4 was synchronized with all the clock pulses, and counter 5 was counter 4 divided by 2. Counter 5 would then define the times for normal and quadrature cells. At this time the correct setting for normal and quadrature had not been defined.

The missing servo clock pulses were used as input to two shift registers 11 bits long that were clocked by counter 5. The output from these registers was decoded by five-out-of-six detectors to produce sector and index pulses. The detectors were protected for single-bit errors in any of the six pattern-locations. The detection of sector or index was used initially to reset the byte counters.

After they had been correctly reset they were used to produce a servo-protect signal, to disable writing onto the servo sample, made up from index, sector (both from the readback), and 'counter-mark' and missing sectors (both created by logic from the byte counters).

Power-on cycle

During the early development of the Piccolo drive, John Taylor gave a talk (February 1976) to the development engineers involved, about the power-up sequence and many other subjects.

Requirements

The following is the sequence of events to be completed at power-on:

1. Checking that disk rotational speed was correct before allowing the actuator to move
2. Synchronizing the phase-locked oscillator (PLO)
3. Distinguishing between normal and quadrature servo cells
4. Synchronizing byte counter for sectors and index.

Possible problem

A possible problem that might have been foreseen required attention. This related to the partial erasure of servo tracks under the gap and slider in the landing zone. When the motor starts the PLO must phase lock to servo clocks for at least a part of a revolution before the actuator is allowed to move. Counter 4 is in synchronisation and a 10 ms single-shot is started, which forces full "out" current to flow to drive the heads out into the data area. At this time, there is neither an error signal (PES) nor a velocity signal.

After this single shot time the head is in the data area and still going out. The following sequence is used:

1. Use missing clocks to preset counter 5 until a sector is detected.
2. Disconnect preset from counter 5 (there might have been several bad servo clocks on the surface giving false missing clocks).
3. Reset byte counter.
4. Use Behind Home detection logic to check that there is another sector detected 600 bytes later.
5. If not return to 1.
6. If Behind Home line goes away, counter 5 in sync line mode is active.
7. Allow demodulator gates to generate normal and quadrature error signals. This also allows generation of velocity signals.

By this time the velocity is much higher than suitable for stopping on a track. The desired velocity is set to be equivalent to more than 2 tracks per millisecond. The arm then slows down to this velocity and stops at the first even track after reaching the velocity.

After this stop, the actuator is sent back to the Home track during which time the velocity is reduced marginally every 4 tracks until the time to cross 2 tracks is ≥ 4 ms. The arm then continues at this velocity until track 0 is detected, at which time the drive goes into track-follow mode and sends a "Home and Ready" message to the system.

Track following

As the track density was higher than for Gulliver, it was necessary to have a servo positioning loop that did more than Gulliver. It was clear that the offtrack caused by wind force and gravity had to be removed. To do this, an integrator was added to the servo loop so that a position error was not required to produce the necessary current. The read-head position had to be very close to the centreline of the read-write tracks. This was done by having a hybrid position-error signal, made up from the low frequency from the data-head PES samples and high frequency from the continuous PES from the servo head.

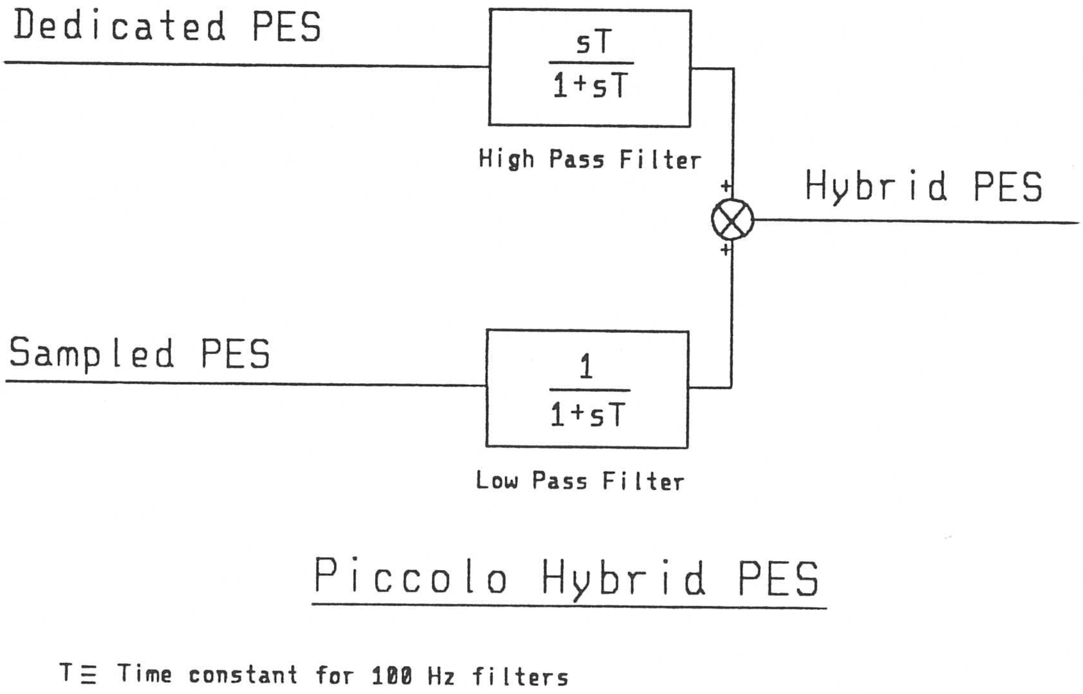

Figure 12. Method of combining dedicated and sampled PES to create hybrid PES

A change of read heads on the same cylinder had to allow a time for settling onto the position required by the new head. The basic loop was similar to Gulliver but the loop had to include the integrator. To do this, the ratio between frequencies of pole and zero had to be increased to 20, which gave a phase lead of 65 degrees at the unity gain point. This was considered more than necessary and the integrator could be brought in at a frequency where it would lag by 20 degrees at the crossover frequency. This frequency had to be as large as possible to give good settling. It was therefore chosen to be at 110 Hz.

It was very difficult to design the rotary actuator. It was never possible to make an actuator without resonances that made the closed-loop system unstable. To remove this difficulty, a notch filter was designed to remove the effects of the worst resonance. Having inserted this notch, a further phase-lead network had to be included to get the correct phase lead at the crossover frequency. This was when APL LISA was first used.

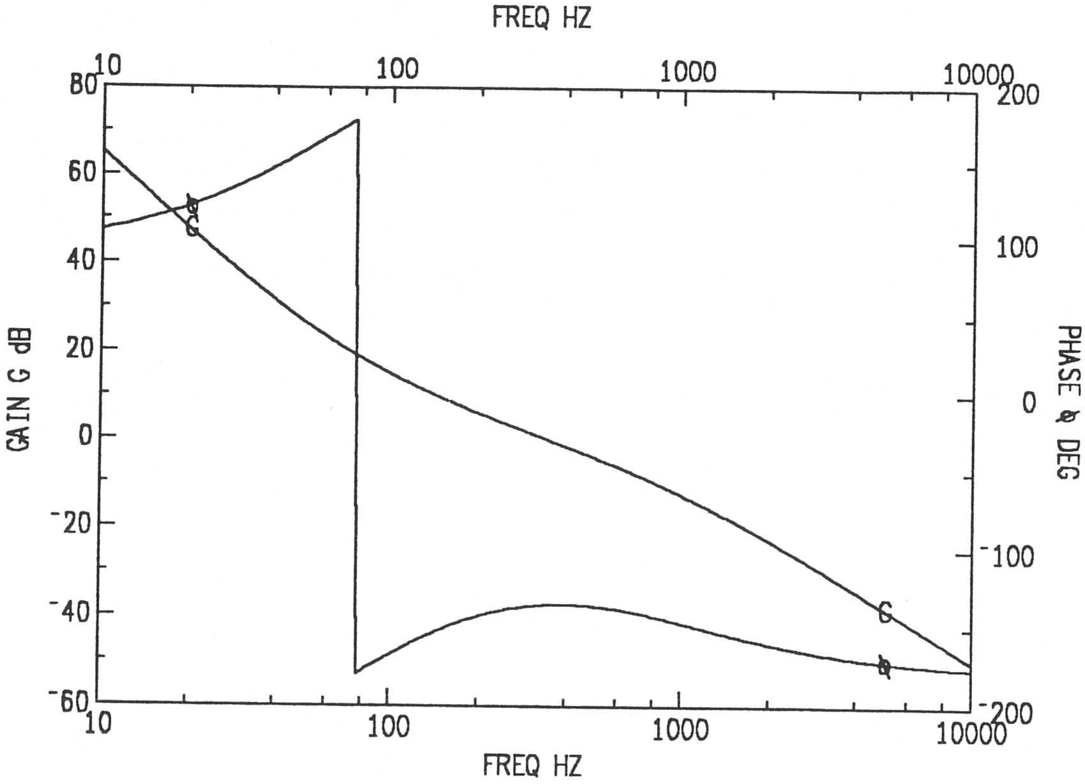

Figure 13. Frequency response of Piccolo position loop

Accessing

The accessing of Piccolo only used the dedicated servo until the head was settling on the desired track. The dedicated tracks used a quadrature error signal that was written in half-track positions. The two PESs from the normal and quadrature demodulation looked like two sine waves at 90 degrees to each other. They are illustrated below.

62PC Two Phase PES

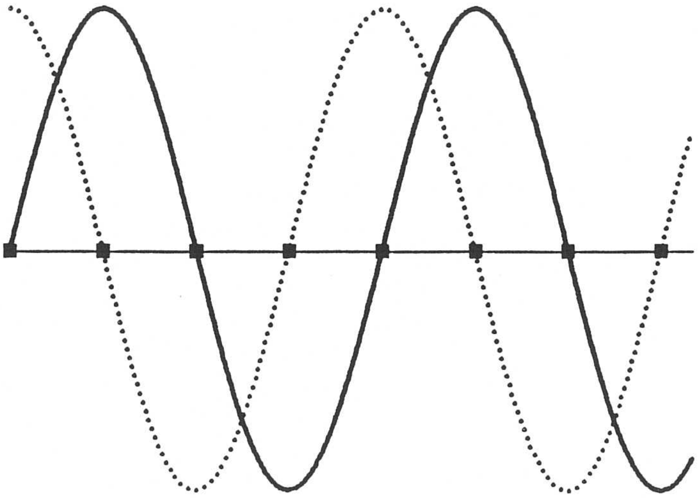

The demodulated signals were selected by the smaller one at the point where the servo head was at all times. This meant that a nearly linear error signal could be obtained at any location of the servo head by selecting normal or quadrature.

The velocity was measured by differentiating the selected PES and combining it with the integral of the VCM current.

The accessing was servo controlled by a loop that controlled the velocity of the servo head to a predefined locus. This locus was calibrated automatically at power-on. The velocity signal was generated by differentiating the smaller of the two error signals (normal or quadrature) and combining with the integral of the VCM current. The hybrid velocity was compared with a stored curve, which had been adjusted at power-on, and represented the ideal curve to reach the desired track in the least time with no residual velocity. The comparison velocity was fed into a bang-bang loop as for Gulliver.

Calibration

There were two values to be calibrated at power-on:

- Gain of velocity profile
- Offset of final velocity.

Gain

The actuator was accelerated for 10 ms to a target track sufficiently distant so that it would still be accerating after 10 ms. After 10 ms, the profile was reduced in steps until the desired profile was lower than the actual velocity, at which time the velocity loop would be activated to follow the profile. The velocity was not exactly correct for getting good settling on the desired track. This was corrected on the return to the home track.

Offset: This was the way the handover velocity was corrected. The initial handover was set higher than it should have been. The loop velocity followed this constant velocity and the time between tracks was measured. It aimed for 2 ms per track. A 3-bit counter, initially set to seven, was added into the desired velocity and was decremented until the track-crossing rate was at least 2 ms per track. This gave a correct velocity for track capturing.

Data channel

The recording used on Piccolo was MFM, as for Gulliver.

SAT

Surface analysis testing (SAT) on Piccolo was performed on the fully assembled disk enclosure (DE).

8.0 Sprat (later Bluegill then Flotilla)

Basic specification

Disk	6-inch removable
Capacity	4.6MB
Average access time	100 ms
Tracks per inch (TPI)	250
Bits per inch (BPI)	20 000.
Data rate	589KB per second
Rotation speed	1500 rpm
Power	dc supplies only

The following pages on Sprat and Bluegill are edited extracts from the only two development reports written at Hursley on Storage projects. Thanks are due to the authors Matt Taub and Chris Wallis.

Sprat at Millbrook January 1975 to December 1977

This section was written by Matt Taub.

Introduction

The recording of digital data on a magnetic surface dates from the earliest days of electronic computers. In the late 1940s, magnetic tape was used, followed in the early 1950s by drums, and in the mid-1950s by rigid disks. The 1960s saw the introduction of flexible disks, but because of wear caused by the disk and head operating in contact, they were suitable only for intermittent operations such as initial program loading.

The report traces the history of a project that set out to overcome this limitation and to exploit and develop the best features of flexible disks. Two of these features are particularly worth mentioning:

- The disk material was inexpensive, making it suitable for storing substantial amounts of information offline. This suggested making the disks removable.
- The flexibility of the disk material promised lower head-to-disk spacing than with rigid disks, leading to higher linear bit densities.

A further aim of the project was to make use of the sophisticated head-positioning servos developed for aluminium-disk stores, so as to get high track densities.

The project, named Sprat, was started at Hursley in the autumn of 1974. It was moved to the Millbrook annexe early in 1975 and continued there until the end of 1977 when, together with key engineering and management personnel, it was transferred to Boulder. During the final year, the Boulder Laboratory assigned some engineers to Millbrook to take over part of the work.

By the time of the transfer a great deal had been achieved. Much successful mechanical design had been done, the servo for track-seeking and track-following was working well, and a very flexible microprocessor system had been developed for the control of the store itself and its test gear. Important improvements had been made in the disk manufacturing process, much sophisticated test gear had been developed, and useful theoretical studies had been made. Some problems remained, however, notably head fabrication and control of head-to-disk spacing.

The time of transfer seemed an opportune one to trace the history of the project so far, to set out in greater detail what had been achieved, and to draw attention to the problems that remained.

Outline

In aluminium-disk stores the head is mounted resiliently to allow the creation of a stable air bearing between head and disk. The basic idea in Sprat, on the other hand, was to mount the head rigidly and to let the flexibility of the disk provide the necessary degree of resilience. To do this, the disk was arranged to rotate close to a rigidly-fixed "Bernoulli plate", the planes of the disk and plate being parallel. The head protruded slightly through a slot in the plate so that the disk "rode" over the head. Two air bearings were thus created, one between disk and plates, and the second between disk and head.

No one is certain where and when this idea originated, but a member of the Hursley Storage Development department, Chris Wallis, saw it in a development project at the Mohansic laboratory around 1967. A few years later, in 1972, several members of the Hursley Storage development including Dave Cuzner who was then Department Manager, saw it again in Emil Hopner's high-density recording work at Los Gatos.

The idea was not taken up at Hursley until 1974, when the Company was making plans to develop a successor to Igar, an in-contact flexible-disk store of the type mentioned in the introduction. The Rochester Laboratory had proposed a machine called Crystal, which was in effect an enhanced Igar, but Bob Evans thought this too unambitious. This prompted him to approach Dave Cuzner, asking him for his opinion on the type of product to which flexible-disk techniques lent themselves, and on the level of performance they would be expected to offer.

Dave's response of 8 May 1974 contained two proposals. One was for a multi-disk "archival" store based on Jim Weidenhammer's work at Poughkeepsie, and the second was for a single-disk store based on marriage of Emil Hopner's head-disk technology with Hursley's servo expertise. As an indication of the direction in which developments were going outside IBM, Dave cited a flexible-disk store with non-contact heads being offered at that time by Dynastor. The first of the above proposals developed into Boulder's "Breithorn" project, and the second, after some revision in the light of market requirements, developed into Sprat.

The initial Sprat proposal, in August 1974, was for a store of 5MB capacity with linear density of 20 000 bits per inch and 250 tracks per inch. The disk was to be chromium-dioxide-coated Mylar, housed in a removable cartridge. By mid-September, a non functional space model that incorporated a cartridge-loading mechanism and a functional novel actuator had been built to show feasibility.

Dave Cuzner and Jack Hockley took this model to the U.S.A. where they showed it to numerous people including Bob Evans and A. G. Anderson. Interest was high and Dave was pressed to undertake a formal product development programme. However, he felt that this would have been premature: there were several aspects of flexible-disk technology that needed more study, and until this had been done, to embark on a full-scale product development programme would have been too much of a gamble.

Bob Evans therefore agreed to fund a design study. It was to start in January 1975 and last nine months. The work was done by a team consisting initially of seven people drawn from different groups in the laboratory, under the direction of Jack Hockley.

The objectives of the study were:

1. To gain an understanding of flexible media technology as it applied to the Sprat proposal.
2. To investigate and quantify areas of the proposal containing technical uncertainties.
3. Using information gained from the previous two objectives, to put forward a product proposal that could be done in a normal product-development time-frame.
4. To produce a working model that demonstrated feasibility of the proposal.

To limit the range of possible parameters, the following were used as a basis:

Disk	6-inch Chromium dioxide (removeable)
Track density	250 tracks per inch
Bit density	20 000 bits per inch
Formatted capacity	4.6MB (256-byte sectors)
Rotational Speed	1500 rpm
Date rate	589KB per second
Average access time	100 ms
Environment	Class C
Power	dc supplies only

By the end of September 1975, good progress had been made in meeting the objectives of the study, and no insoluble problems had become obvious. However, several areas required further work before it was possible to define the product, such as the use of an analogue or digital approach to servo control, and a better understanding of the head-to-disk interface. The latter has an effect on the establishment of an acceptable bit density when trade-offs are made with error rates in the data-channel design. Work continued into 1976 to answer these questions and progress Sprat in all other aspects of the design.

Early in 1976, GPD, Boulder proposed Flatirons, a store similar to Sprat in concept and specification; it, too, had flexible removable disks and was based on technology developed for Breithorn.

Because of the difficulty for SCD to continue to fund Sprat and to supply sufficient headcount to enable a satisfactory conclusion for the project, and because of GPD Boulder's mission and interest, in August 1976 agreement was reached in principle to discontinue work on Flatirons and continue work on Sprat in Hursley with Boulder assignees. This would continue until a convenient point in the product development cycle, when the project would be transferred to Boulder with assignees from the Hursley Laboratory. The Document of Understanding was issued on 15 November 1976 as a joint SCD/GPO product development. After a short engineering test in Hursley to be done between mid-November and mid-December 1977, the project was to be transferred to Boulder at the end of that year with six assignees from Hursley.

The division of responsibilities was as follows:

Hursley

- Overall product design
- Cartridge design
- Development of read-write head profile
- Cost estimates
- Supply cartridges to Boulder for media test
- Supply precision data-channel spinner
- Design device-test and servowriter equipment

Boulder

- Develop data channel
- Supply media to Hursley specifications
- Provide head fabrication
- Perform technology work on media
- Perform technology work on media/head wear
- Provide business planning
- Provide user liaison
- Develop user liaison
- Develop attachment strategy and design.

The first assignees from Boulder arrived early in 1977, and when the project was at its peak the people who were involved directly were:

Hursley staff

J Hockley
D M Taub
S Y Rossiter
C Harding
L J Rigbey
C N Wallis
W P Case
I E Henderson
D J Craft
F D Dickin
M J Carmichael

J Farran
P Crumplin
M Saunders
M Werman*
R Chaplin*
J Holloway
T Brew*
R Pascual*
M Hatchett*
A R Cox
T Wright*

* From other groups

Assigned to Boulder

D A Evans
J Hartley

Boulder staff

R Bellamy
C Farel
J M Levine
R Cope
L Fangmeier
M Buxman
M Nettles
R Millington
N Feliss

Four models were built, together with the necessary test equipment and servowriter, but because of a delay of about six weeks during the debugging phase when all parts were brought together, it was decided that Engineering Test would take place in Boulder after transfer. To minimize the time spent on transfer, shipment of equipment and office effects was made over the Christmas – New Year break, so that the models were available when the bulk of the assignees from Hursley arrived. Jack Hockley continued to manage the projects in the U.S.A.

Head flying and wear

It was recognized from the outset of the study, that the head-disk interface was one of the most critical areas. Every project without exception experiences problems and especially when the head-disk interface is being developed in parallel with head fabrication, data channel, and media processing.

In April 1975 a simple rig was completed for measuring the flying height of the disk over the head by the use of white light interferometry. Difficulties were encountered in trying to fly uncoated, unburnished disks, which besides being unrepresentative, produced amounts of debris that changed the flying height. This was only partially overcome by cleaning.

The head profile determined by Emil Hopner was initially used, and later the air-bearing surface was increased. These heads all had the common fault of extreme sensitivity to penetration for a particular desirable flying height, until the partial or intermittent contact height was reached when fair stability was shown. Tests also showed that substantial abrasive wear was caused by such low flying.

The objective was to be able to define an easily manufactured profile with a flying height of 8 – 10 microinches over the data band, with as high a penetration tolerance as possible without significant wear. The first heads clearly did not approach the equipment, and up to February 1976 a large number of different types were tried, mostly in transparent spinel that allowed the correct coated and burnished media to be used.

In general, the results from these heads after short evaluation indicated that a spherical surface had the basic requirements, but needed some feature upon it to stabilise the close-flying area.

Up to this time heads that showed promise were evaluated in more detail by taking photographs of the fringes for different values of penetration and track radius positions. By laboriously plotting the results, the effects of the variables could be compared.

For each head the whole process involved 105 photographs and a whole week's work, and to reduce this time to manageable proportions a comprehensive flying rig was built by June 1976. By means of two TV cameras, and an x/y coordination system with mixer and monitor, the close-flying area could be semi-automatically plotted and the results stored for subsequent data reduction. A full test could then be done in two hours.

The profile showing most promise consisted of a spherical air-bearing surface 3 by 2 mm with a shallow annulus cut into the surface and centred at the magnetic gap. By adjusting the inside diameter of the annulus, the flying height could be controlled, and the effect of yaw, due to angularity of the actuator, was reduced. Although when flying at a particular radius on the disk, the penetration tolerance was 0.005 inches for a particular head, owing to roll effects this tolerance was reduced to 0.001 inches when movement over the data band was taken into account.

This tolerance was too small to allow for cartridge interchange, actuator setting, and head tolerances.

A small number of tests with fixed and accessing heads were made, and the worst results showed a wear rate of 30 microinches per 1000 hours, the surface of the head remaining polished.

In mid-May 1977, work was started on a miniature ring stabilizer. This showed promise, as it appeared to reduce the sensitivity to pitch and roll while keeping conditions constant for a stabilizer flying height over the range of track radius. The flying height of the head could be varied by changing the width of the air-bearing surface, omitting the annulus.

Great sensitivity to the profile of thè toroidal radius was noticed, some stabilisers producing a high negative pressure within the ring, which caused the disk to remain deformed and sealed to the stabiliser for several minutes after stopping. Startup with the disk in this condition stalled the motor, and the size of motor and the control system had to be changed to overcome the friction. Other stabilizers that produced a pressure only slightly below atmospheric did not control the stability sufficiently well for good signal output.

A study of stabilizer profiles and their effects on the disk was made in August 1977, keeping the toroidal radius and diameters of the stabilizer constant, but varying the apex diameter and the relative height difference between the stabilizer and the head. The result showed that the larger the apex diameter the more aggressive the stabilizer became, and that the stabilizer and head should be at the same height, within 0.0001 inch.

To meet the project requirements during debugging, stabilizers were added to the annulus head. Although unacceptable wear was created, the surface of the head remained smooth with a well defined gap. This "plan of record" head was later changed to a chamfered head within a stabilizer.

Although the flying height investigation indicated that this was a promising head, when preliminary wear tests were carried out in Hursley, and subsequently in Boulder, the chamfered head showed that catastrophic pitting takes place within 20 hours, even with a disk that is flying a nominal 20 microinches from the head. At this time all work on flying and wear ceased at Hursley and the rigs were transferred to Boulder.

In summary, it can be seen that the original objectives were not met and that considerable work remained to be done to produce an acceptable interface that was a compromise between the conflicting needs. With hindsight, it would appear that the wear measurements on an annulus-type head should have been repeated, and a wider range of annulus diameters and widths examined. On the positive side, the techniques of data collection and accurate measurement of the minute wear rates had made significant advances, while the early head work enabled other areas of the project to continue uninterrupted to Engineering Test.

Those involved in this work were Tony Evans, Mike Carmichael, Bert Feliss, Alan Cox, Rafe Pascual, Leo Rigbey, and Mike Nettles.

Backplates

The technique of flying a flexible disk over a "Bernoulli" plate is well known, and is used in IBM for test rigs. Emil Hopner also used this principle and it was included on the early Sprat rigs. Its use raised three questions: stability, wear during starting and stopping, and the specification of the finish.

Work in January 1975 tested various finishes ranging from mirror smooth to coarse sandblasting, and it was found that finish did not affect the disk-to-plate flying distance during running. However, the coarsest finish gave unacceptable wear during starting and stopping, and the smoothest finish, though it caused no wear, made the disk adhere to the plate and destroy it.

The solution developed was to produce small pits by photo-etching a smooth plate, leaving sufficient smooth area so that no damage occurred during starting and stopping. This also overcame the difficulty of specifying the finish so as to enable a repeatable product to be manufactured. The particular pattern selected was used for the whole of the period in Hursley.

The sensitivity of the height of the disk hub relative to the plate was also closely examined to give guidance on tolerances and help with the mathematical simulation of this interface. In practice with the adopted configuration, it was found that the hub height had little effect upon the disk flying height over the data band, and that the air flow rate around the hub was so small that normal clearances were adequate.

It was noticed that with high penetration of the head into the plane of the disk, when the head was moved towards the outer edge of the disk, instability occurred with the disk breaking away from the plate. In February 1975, a study was set up to measure the separation between the plate and disk with and without a head in a radial slot. From contour plots of these results, it was discovered that the slot alone, or any significant discontinuity, disturbed the disk for more than three-quarters of a revolution circumferentially, and that the "tent" caused by the penetration of the head exaggerated this effect and produced perturbations extending radially to the disk edge.

To minimize this effect and allow the head to approach the outer disk edge more closely, the plate was formed with two chordal bends adjacent to the slot. This had the effect of stiffening the disk over the slot area and confining the area of tenting to around the head. Some work was done to optimize the angle of the bends so as to reduce the thickness of the cartridge to a minimum.

These conditions were used throughout the project until April 1977 when a complete matrix of experiments on wing angle and position of chordal bends was performed. The results indicated in general, that even a small angle helped to confine the tent in the disk, but the sensitivity to penetration was then greater, and thus the dimensions established earlier represented a reasonable compromise.

In summary, the bent-plate approach was significant to the project in that it allowed data to be stored almost to the edge of the disk. Although the negative-pressure stabilizer may also help in this respect, at the time of transfer it had its own problems of head wear.

Those involved in the early work on measurement and methods were Tony Evans, Mike Hatchett, and Leo Rigbey. The bent-plate design was by Mike Hatchet and Leo Rigbey. The matrix of measurements was done by Bert Feliss and Mike Carmichael.

Mathematical analysis of disk-flying

Early in the project it was realised that, for an optimum design, a thorough understanding was required of the theoretical aspects of flying the disk over the backplate and over the head. Tony Evans was assigned to this work at the end of 1974.

The complete problem, taking account of the backplate with its slot for the head, and the head itself with its surface features, is extremely complex. Tony therefore attempted the solution in stages, starting with the case of a disk flying over a flat backplate, and ignoring the slot through which the head protrudes. Further simplifying assumptions were made, among them: (i) that the air between disk and plate was incompressible, and (ii) that the flow was axisymmetric. Navier Stokes equations for constant viscosity were then formed, together with the equation for the elasticity of the disk. These were solved numerically for laminar and turbulent flow.

The main result, reached in May 1975, was that the disk-to-plate spacing at the hub had little effect on the spacing further out. The spacing over the range of radii occupied by the tracks was calculated as 0.006 to 0.008 inch. Mike Hatchett checked the results experimentally and found that agreement was reasonable though not outstanding.

Tony then turned his attention to the problem of flying the disk over the head. To try to get a solution from scratch would have taken much more programming effort than was available. He therefore set about adapting a program originally written by Bill Langlois for the Boulder mass storage system (MSS).

The program in its original form was installed at Hursley during September 1975, but changes had be to be made to cater for the differences between MSS and Sprat. For instance, in MSS the stress in the medium is longitudinal, whereas in Sprat it is the bursting-stress characteristic of a spinning disk.

Unfortunately, agreement between theory and measurement was poor, probably because of additional differences between MSS and Sprat conditions. Also, the program was very expensive to run. At the end of 1975 therefore, this line of investigation was abandoned.

By this time, the bent backplate had come into use, and an attempt was made to analyze the new head/plate flying conditions. The work was pursued for several months but without success, and later attempts by others to solve the same problem fared no better.

The question of the natural modes of vibration of the spinning disk arose at about this time (early Summer 1976) and some calculations were done, but not checked experimentally.

In June 1976 Tony was assigned to Boulder, which gave him a better opportunity of keeping in touch with the analytical work being done there and at other U.S. locations. He found that some relevant work was being done at Rochester: the flying of a flexible disk with no backing plates over a head with a spherical surface, with either plain or with longitudinal slots.

It was several months before the program was made available to him. He received a copy of the object code in February 1977, and found on using it that in spite of the simplifying assumptions on which it was based, the results agreed well with Millbrook measurements. In April 1977 the source code became available and the Rochester people modified it to Tony's instructions so that it could be used on the annular-groove head. The program was run several times between May and September 1977, and gave consistently good results. It would have been valuable to go on to analyze the case of the same head mounted in a stabilizer but this turned out to be beyond the capacity of the program.

During the period July to October 1977, Tony, together with Armando Argumedo, designed and built a new rig for measuring head-to-disk flying height. Its main purpose was to check the Boulder-made heads, and unlike the earlier rig made at Millbrook, it could measure instantaneous flying height during track seeking. It had been feared that at the moments when sudden changes in acceleration were taking place, that is, at the beginning and end of the acceleration and deceleration phases, there would be perturbations in the flying height. However, measurements made between October and December 1977 showed no evidence of this.

To sum up, the problem of analyzing the flying of a disk over a backplate and over a head is a particularly intractable one. Nevertheless, progress was made. Reasonable agreement with experiment was found for the case of a disk flying over a flat backplate, and good agreement for a disk flying over a plain or annular-groove head. However, the problems of flying over a bent backplate and over a head/stabilizer combination remained.

Head fabrication

The interface between head and disk was discussed earlier in this report. The present section is concerned solely with fabrication processes developed by the Sprat head.

The heads used by Emil Hopner were modified Aries elements in which the gap length was reduced and the sides chamfered to produce elements of the required track width. As there were long delays in delivery times from San Jose it was decided to build up the facilities for slitting, within the Project, at Hursley. Head bars from the Winchester head line, supplied by Mainz, were then machined to provide individual elements.

Also, because of long lead times from other parts of the Company, it was decided to provide our own profiling capability, and in February 1975 a manually controlled machine using a flat lap was used to produce heads with a spherical contour. The width of the track was obtained by chamfering the mounted element to 0.0037 inch.

Calculations from the data-channel group indicated that a gap length of approximately 25 microinches was required to achieve the bit density, and it was decided that in the interest of the flexibility of design and speed of response, the skills required for glassing ferrites should be acquired. With help from Mainz this was achieved, without the need for sophisticated equipment, using a workshop furnace and Hursley-designed jigs. It also involved the setting up of sophisticated cleaning equipment, developing the process for absolute cleanliness, and making use of the SO_2 deposition facility in the Hursley Laboratory.

Early in 1976, improvements were made to the lapping equipment to improve the repeatability of the spherical contour and the surface finish of the air-bearing surfaces. A novel process was developed, in which the head and head mount were attached to a dopstick aligned in the ball and cup lapping tool. This allowed a spherical surface to be generated precisely on the axis of the dopstick.

Dummy heads made of spinel were produced as well, for the purpose of flying-height measurements. Spinel, magnesium aluminium oxide, was chosen in preference to glass or quartz because of its geater hardness.

From the time the study started until mid-1976 more than 35 varieties of head profile were produced. It became obvious that the only practicable profile, for manufacturability and repeatability, was spherical with an added feature to control the flying height. At this time the bulk of the magnetic heads produced incorporated small notches that limited the track width, and an annular groove centred about the magnetic gap. These heads enabled progress to be made in all areas of Sprat development, although it was known that wear was taking place and that the flying height varied over the range of data radii used.

When Boulder took over the responsibility for head fabrication in November 1976, advantage was taken of the processes already developed there, in which ferrite of the thickness equal to the required track pitch was sandwiched between pieces of barium titanate ceramic (BTC). This avoided the necessity for notching to define the track width, and brought the air bearing surface back to full size.

The first elements of this type were received from Boulder in December and more in January 1977. These all suffered from badly defined gaps, poor glass lines between the ferrite and BTC, and after profiling showed a very low output.

Further batches received during February and March 1977 showed large variability of signal output and a further batch of the earlier Hursley type heads was made to keep the project supplied.

To determine whether stress on the ferrite was generated during BTC glassing and causing the variable output, heads were made in Hursley using adhesive to attach the BTC side cheeks. These showed a consistent improvement before movement at the adhesive interface made the heads useless. However, they lasted long enough to show that stress played a major part in the low signal outputs and efforts were made to reduce mechanical loading in all jigs and fixtures at all stages.

As the year progressed, the head quality from Boulder improved to the point where satisfactory data-channel operation was possible.

In May 1977, work was being done on the toroidal stabilizer discussed in the section on head flying and wear. Because of the surface finish requirements, it was necessary to hand polish the turned surface of the toroid and this resulted in discrepancies in profile, which could be greater than 10 microinches, causing large variations in flying characteristics.

In November 1977, a "coining" technique was used to produce repeatable, low-cost, and low-mass stabilizers from a hollow eyelet. Successful stabilizers were made during January 1978, but further work on the tooling was required to prevent the gradual collapse of the dies because of the high local pressures.

For cartridge interchangeability, the azimuth of the head gap must be precisely controlled with respect to the actuator pivot. Tooling for use under high power magnification was developed enabling settings to be made simply, and to an accuracy better than 20 microinches over a length of 0.004 inch. This also involved a total concept of attaching the mounted head to the arm, providing a method of adjustment, and terminating the windings of the head with a strain-relieved thin tape-cable.

All the jig design, machining and fabrication were done by Alan Cox, Leo Rigbey, and Tom Wright. Rafe Pascual set up and performed the cleaning operations and provided measurements of the heads at various stages. Dick Millington was involved with the head azimuth jig work.

Data channel

Introduction

The data channel of a disk store includes the following:

A. Circuits for converting the data to be recorded into a suitable code such as MFM

B. The "write" amplifier which produces the bidirectional current for the recording head

C. The recording channel itself, consisting of the head and the recording medium

D. The data "read" amplifier

E. The data detector circuits, which recreate the original data expressed in the code used for recording

F. Circuits for converting the read-back signals back into the code used by the host system.

The outer portions, A, B, and F are relatively straightforward and are not covered here in any detail. The portion of the channel that presents the greatest problems is C mainly because of the difficulty of keeping close enough control of the spacing between head-gap and disk – the flying height. Allowing the two to operate in contact brings problems of head and disk wear, whereas too great a flying height degrades the recording and reading performance, setting a limit to the number of bits that can be recorded per unit length of track.

To some extent this degradation can be corrected by "predistortion" before recording or by including suitable networks in the "read" amplifier – post-read equalisation – but this becomes more and more difficult if the flying height, and therefore the degree of degradation, varies.

In Sprat, control of the flying height was the most difficult problem the project had to face, an associated problem being control of the magnetic properties of the head. Work on the rest of the data channel continued with these problems no more than partly resolved. Nevertheless, very useful progress was made and design techniques were developed that gave confidence for further progress with resolution of these problems and the design of data-channel circuits.

Concepts and terms

Information is recorded along the track of a disk by reversing the field produced by the recording head at appropriate intervals of time, producing reversals of magnetization on the disk surface. The recording-head field can be reversed very rapidly, in a small fraction of a bit period, but the resulting magnetic transition in the disk tends to be relatively spread out along the track because of demagnetization effects and head-to-disk spacing.

During reading, the effect of a transition passing the head is to induce a voltage pulse in the windowing, and again because of head-to-disk spacing, the spreading effect is increased still more.

Two methods of data detection are in common use: one is to detect the positions of the read-signal peaks – peak detection. The other is to integrate the read signal, producing a waveform nominally the same as the original write-current, and to sample it at regular intervals such that all the reversals are detached – sampling detection.

As linear bit densities become higher, the magnetic transitions themselves and the "head" pulses they produce, tend to merge into one another with two undesirable results:

- The magnitude of the "read" signal is reduced, and
- The positions of the peaks are shifted.

A measure of the first effect is the ratio of the read signal produced by a transition in every bit cell to the signal produced by transition in every second bit cell. This ratio is called the "resolution".

The second effect, peak shift, can be expressed in absolute terms or as a percentage of a bit period. If it is too great, the data is detected incorrectly.

Another measure of performance is the amount by which the head may be laterally displaced during reading, before the error rate becomes unacceptable. This is usually expressed as a percentage of the track width and is know as the "off-track" performance.

History

Progress during 1975

The initial aims for recording density were 20 000 bits per inch along track and 250 tracks per inch. Bill Case and John Hartley made some simulation studies during the summer of 1975 to check whether these figures were attainable. They used a model of the recording process presented in a paper by Talke and Tseng and came to the conclusion that, to reach the desired linear density, one had to use a chromium-dioxide-coated material such as Crolyn and a very low head-to-disk spacing, perhaps 10 microinches. They found that iron-oxide-coated material (Mustang) would give a linear density of half the desired figure: the reason is that the iron-oxide material has a lower coercive force and so, at a given linear bit density, the self-demagnetizing effect is greater.

These findings were tested experimentally soon afterwards, with encouraging results. Mustang appeared capable of giving 9000 bits per inch and Crolyn 17 000 bits per inch.

A problem that was particularly severe at the time, however, concerned the track width. In order to form an air bearing between the head and disk, the head had to be about 0.06 inch wide. In an effort to define the correct track width, slots were cut from the two sides of the head thus:

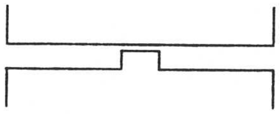

But the slots were of little help. During writing, high fields exist along the edges of the slot as well as across the gap, tending to make the head write to the side of the track; also, during reading, the portions of the head associated with the slots tend to pick up the low-frequency components of any information passing under them. Track-width definition at the time was therefore very poor.

Several other matters received attention during 1975, among them the choice of recording code and type of detection. In the interests of good off-track performance, so important with exchangeable-disk machines, it was decided to use MFM code and peak detection.

Another topic was the reduction of peak shift by post-read equalization. The equalizer used was one originally described by Vermeulen incorporating a lumped-constant delay line, open circuit at the receiving end.

New mechanical hardware was built (the data-channel rig), in which various adjustments to the head position could be made. It could be moved in a direction normal to the disk surface and its attitude could be adjusted relative to the three directions: along, across, and normal to the track. These adjustments are known respectively as penetration, roll, pitch and yawl. Sensitivity to these variables was one of the factors that had to be investigated. The final piece of data-channel work for 1975, completed in December, was the design of a low-noise read preamplifier.

Progress During 1976

Experiments during January and February 1976 showed that, with the head in use at that time, performance was very sensitive to change in pitch, roil, penetration, and track radius, showing that these variables had a marked effect on the flying height over the gap. There was a great improvement, however, when a new design of head became available the following March. This had an annular groove cut in the surface, and the track width was defined by V-notches cut in the two sides:

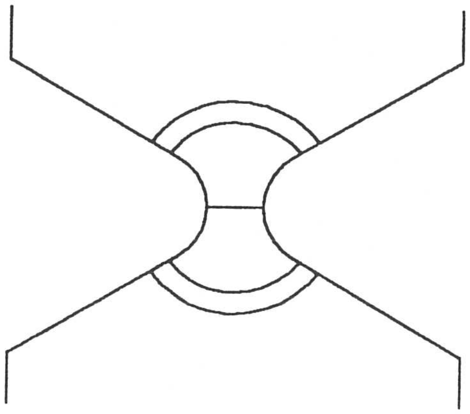

This was known as the "mesa" head, and in addition to giving more consistent flying height, it gave much better track-width definition.

The early equalizer had given poor performance, but an analysis by Matt Taub showed that this was largely a matter of incorrect component value. He also proposed a modified circuit that would give a lower signal-to-noise ratio of the same reduction in peak shift, but pressure of other work prevented this idea from being tried out.

Unfortunately, the same happened later in the year with an idea for improving the data-detector circuit. A property of the MFM code is that peaks in the "read" waveform representing binary 1 can undergo greater shift than those representing binary 0. The proposal therefore was to increase the proportion of the bit period during which a peak is recognized as a 1, and to reduce the proportion during which it is recognized as a 0.

A problem that received attention during the autumn was asymmetry in the "read" pulses: it is caused by a component of magnetization normal to the track surface, and acts as an additional source of peak shift. A correction circuit was designed and put into use in October 1976 with good results.

About the same time, John Farran was looking into the possibility of improving the peak shift compensation still further by predistortion, in other words, by deliberately advancing or retarding the current reversals in the head during recording. This method is used in Piccolo and was found to be helpful in Sprat as well.

By the year end it was possible to read at 20 000 bits per inch with an off-track performance of 10%.

Progress during 1977

At the end of 1976, work on the data channel was transferred to Boulder. The data channel was shipped across and the work was continued by John Hartley who was assigned from Hursley, and Jerry Walker.

The first important development took place in February when Ted Smathers provided the project with two new heads of an improved design. Instead of using a piece of ferrite wide enough to produce good flying, and notching it to define the track width, the head was formed as a barium titanate – ferrite – barium titanate sandwich, the track-width being defined by the ferrite portion:

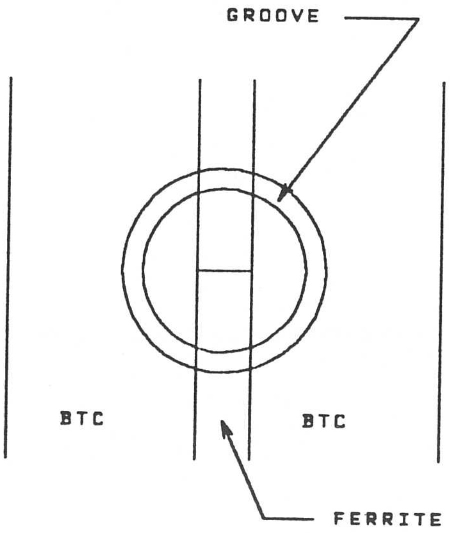

This head had the advantage over the earlier one of much better track definition and lower inductance.

A new equalizer was brought into use at the same time. It was of a type that had been developed at Boulder, the method of design being quite different from that used at Millbrook. In the Boulder approach the record/read process is regarded as if it were a linear communication channel. A pseudo-random binary sequence is applied to it, and the output waveform recorded, from which it is possible to deduce the transfer function: gain and phase as functions of frequency. The desired transfer function, giving less peak shift, is known, and therefore the required equaliser characteristic can be reduced. A circuit having this characteristic is then designed using known synthesis techniques.

When the new equalizer and head were tried together, off-track performance improved to 25%, though it should be noted that conditions during the test were not worst-case.

A topic that was reviewed on two occasions during the year was the choice of recording code and type of detection. During March and April, John Hartley made a comparison between MFM and 0,3 (8/9) run-length limited codes and between peak and sampling detection, and in November Jerry Walker repeated this using different detector circuits. On both occasions, however, the results were inconclusive.

Control of flying height had been improving as the project progressed but despite this, sensitivity to penetration, roll and pitch were still too high for an exchangeable-disk store. A way of improving matters was introduced which acted as a stabilizer.

Heads with the annular groove, mounted in stabilizers, became available in September 1977, but though the stabilizer was very effective, the mechanical and magnetic properties of the heads themselves were poor. A later design, in which the annular groove was replaced by chamfers on the sides of the head, became available in November, but its properties were even worse. It was this head that Jerry Walker used for comparing codes.

The remaining data-channel work done during 1977 consisted of the specification of the data-channel interface and the design, building, and testing of circuit cards for this and for the data channel itself. The work was done in preparation for the Engineering Verification Test being carried out early in 1978.

The Position in February 1978

Three factors held up the successful design of the data channel. The first was that flying-height had yet to be properly controlled. There is good evidence that during most, if not all, of the data-channel experiments, head and disk were running in contact. This gave good data-channel performance, but was unacceptable in the long term because the wear was too high.

The other two factors concerned the magnetic properties of the head. Calculations showed that with the existing head material, NZ4T nickel – zinc ferrite, the value of write current was enough to saturate the pole tips. The effect of doing so was to increase the spreading of the magnetization transitions in the disk, which, at the high bit densities used, degraded the data-channel performance. Alternative materials offered the hope of a very worthwhile improvement.

The remaining factor concerning the head material was permeability. Work carried out by Richard Thornley during the summer of 1977 showed that the effective permeability was a good deal lower than it should have been, probably because of stresses induced during the glassing and bonding processes. The effect of lower permeability was to diminish the "read" signal.

The correct head-flying choice of head material and the head fabrication process were matters that needed the most urgent attention.

Mathematical modelling of the record and read processes

Among the important characteristics of Sprat are:

- Its use of a very high linear bit density (20 000 bits per inch)
- Its use of a thick recording medium (approximately 180 microinches) such that recording takes place only in the surface layer
- The head and medium being out-of-contact
- The medium being exchangeable
- The impracticability of making a separate erase operation before writing, mainly because of the use of a pivoted actuator arm.

Many existing tape and disk machines had some of these features but none that were known of them all. The team of engineers therefore had to satisfy themselves that such a machine was a practical proposition.

The only certain way of doing so was by experiment, by building a fairly substantial number of machines and demonstrating that they continued to work reliably with disks being freely interchanged between them. During the Millbrook phase of development, however, this was out of the question: the head fabrication and head/disk interface problems had not yet been satisfactorily resolved, and too few machines had been built. The team therefore had to do the best they could by mathematical modelling.

As in the case of disk flying discussed before, a rigorous solution would have been enormously complex: the program would have taken a very long time to write and would have been prohibitively expensive to run. To make any progress at all, one had to make simplifying assumptions. Several people had worked on the problem over the years, each of them making different simplifications, but at the time the team became interested in modelling Sprat, during the spring of 1977, there was no program that the team could simply pick up and use. A program was developed by Matt Taub during the following few months.

The approach he used was based on Jim Smith's analysis, but various extensions were made. The resulting program allowed the user to choose a value of deep-gap flux, head-to-disk spacing, the pattern of transitions to be written, and the "transition distance" that is, the product of write-current transition time and the velocity of the medium. The program then computed the magnetization of the medium as a function of the distance along it and the depth into it. A second pattern can be written over the first with different values of the above variable, the two patterns having any specified phase relationship to one another.

One of the main simplifications concerned the self-demagnetization of the medium. During the actual process of writing a pattern this is ignored: the medium is then assumed to be suddenly moved from the influence of the head, whereupon self-demagnetization takes place.

By September 1977 the work had reached an advanced stage but had to be interrupted in favour of the track misregistration study reported in the section on track misregistration. A remaining piece of work was to extend the analysis so that a different head-gap length could be specified for each of the "write" operations and for the "read" operation, to cater for the case where all these operations are done on different machines. It was planned to complete and to document the work during 1978.

Microprocessor control

Background

Magnetic-disk stores contain a number of important control functions. For example, there is the servo system for track-seeking and track-following, drive-motor control, various interlocks, and interpretation and execution on command from the host system.

In the past, the usual practice had been to use analog techniques for the servo, and specially-designed logic circuits for the rest. Late in 1974, however, Dave Craft proposed implementing the control functions by using a single microprocessor time-shared between them. In this proposal, the analog servo would be replaced by a digital one, making the servo just one more digital control function, albeit a complex one, sharing the processor with the other control functions.

The main attraction of this scheme lay in the way it allowed large-scale integration to be used effectively: it offered lower cost without incurring the penalty of long development times. The lower cost per logic element justified building a much more sophisticated control system than would previously have been economic in stores of this size. For instance, it could keep track of the rotational position of the disk without any help from the host system. It could also include such features as automatic retry of commands, error diagnosis, and dynamic optimization of the servo. (The latter features have been considered but are not incorporated as yet.)

In designing circuits peripheral to the microprocessor, the objective was to design them to be as "universal" as possible. The microprocessor would then specify their exact function by sending them control information. In this way a modification to the overall control system, such as a format change, would require nothing more than a change in the microprocessor program — much quicker and less costly than developing a new special-purpose LSI circuit.

The above scheme was a natural follow-on from the SIAM Fellowship project on which Dave Craft had been engaged during the previous few years. When that project finished, Dave did some preliminary work on track-seeking algorithms and had designed a microprocessor architecture specially suited to their implementation and to the control functions needed in a disk store. The possible use of this machine, Inca by name, was being discussed in the Sprat project in August 1975. For economic reasons, however, it was not proceeded with, and instead, Dave made a comparative study of other microprocessors, some produced by outside vendors, and others under development within IBM, to find out which of them would be the best alternative. The study lasted from late 1975 into the first few months of 1976 and included the following:

```
IBM Atom 1, Atom 2 and Cornerstone
Intel 8080
Motorola M6800
Zilog Z80.
```

Other factors being equal, an IBM machine would have been preferred. However, Cornerstone had too poor a performance, Atom 1 was big, expensive, and would have required forced-air cooling, and Atom 2 was then no more than a technology study. With a strong interest in the M6800 elsewhere in the Company, the initial decision was to use this machine as a stopgap, in the hope of Atom 2 becoming available in time to meet the project schedules. The first Motorola hardware was delivered in February 1976. Towards the end of that year, however, the future of Atom 2 was still uncertain and so it was decided to retain the M6800 permanently.

M6800 software preperation

The initial method of preparing M6800 software was as follows.

Assembler language statements were entered into a 2741 terminal at Millbrook and transmitted to the Hursley APL system where they were assembled. (The cross-assembler program was completed in April 1976.) The resulting M6800 code, expressed as a succession of characters in IBM Correspondence Code, was returned to the 2741 at Millbrook where a local 1130 system had been arranged to "eavesdrop" on the information and produce punched cards. These were fed back into the 1130, which converted their contents into ASCI code and sent it serially to the M6800.

This scheme was satisfactory while the programs were small, but for large programs the low speed of the 2741 communication link created a serious bottleneck. The difficulty was overcome towards the end of 1976, by replacing the 2741 link by a faster link using Tektronix 4012 terminals.

It had been suspected from the start, that as the amount of programming work increased, a time would come when the APL system would not be able to keep up with the demands made on it. The Applications Development section of the DP Centre was therefore asked, in February 1976, to provide an 1130 cross assembler. The program was written by Ann Underhill and was being brought into use around the following November. There followed a phase-over period of a few months during which the APL and 1130 programs were used side by side.

Improvements to the 1130 program continued until about October 1977.

Software and hardware for track-following

Most of the microprocessor's time would be spent on track-following, and so the first task was to evaluate the algorithms already proposed, and if necessary to develop them further. As part of this work, Chris Wallis spent February to April 1976 building an electrical analogue of the actuator. It used operational amplifiers, and suchlike, and could be controlled digitally by the M6800. In the interests of speed, multiplication in the microprocessor was by means of stored tables.

By the middle of 1976 it had become clear that the simple algorithms on which the team had based their hopes were not capable of giving a good enough performance. More complex algorithms were found to be necessary, but with these, the M6800 turned out to be too slow.

The way out of the difficulty was to build a special algorithm unit to be used as a peripheral device in the M6800. Its main components were a hardware multiplier, two register tracks, and various special-purpose registers. Design had started in October 1976, and after various development problems the unit was working satisfactorily in August 1977. The delay between its receiving a position-error signal and outputting the desired value of actuator current was about 30 microseconds.

Control-system architecture and hardware

Late in 1976, potential users of Sprat were asked what architectural features appealed to them in a low-cost disk store intended for common subassembly use in entry systems.

In summary, the results were that the more function that could be taken out of the adapter and placed in the control system of the disk store itself, the lower would be the overall cost, and the more users we could expect to attract. In particular, the serializer/deserializer unit was awkward to accommodate in the host system, because to handle high-speed serial data often demanded a different logic family from that used in a low-cost system. Furthermore, function incorporated in the store has to be developed only once, whereas in general the adapter design has to be done anew by each user, and was a serious disincentive to new users.

The interface, it appeared, should be as simple as possible, with few lines, and loose timing tolerances.

Following this conclusion, Wallis devised and proposed the 1976 Sprat architecture and logic configuration, which exploited the LSI microprocessor configuration and attempted to meet the requirements described. It was derived from the concept outlined in Background on page 53, but the circuits peripheral to the microprocessor had grown into more complex pieces of logic, specialized to particular store and adapter functions that were either too fast for the M6800 or would have imposed too heavy a load on it. They are now referred to as "Functional Units" and a list of them together with their functions, is given here:

Functional Unit	Use
Interface unit	Handles interface protocol, serialization/deserialization, and command, and status buffering.

CRC generator/checker	Appends check-bytes to data being written; validates check-bytes in data being read.
ID checker	Searches for a match between sector address demanded and sector address read from the disk; buffers addresses read for input to microprocessor.
Data-channel control	Sequences controls to data channel; monitors progress of write operation; and reports to microprocessor if any errors are detected.
Format unit	This is programmed from the microprocessor to provide a variety of timing signals derived from counted-down clock.
Address decoder	Operates hardware control lines upon detecting presence of appropriate word on microprocessor address bus.
Algorithm unit	Performs calculation of actuator current from position-error signal, using coefficients supplied by microprocessor

The interface had a combined data, command, and status bus, having eight parallel bidirectional lines, parity, and three unidirectional strobe lines to provide dc interlocked operation. These lines are all common to however many Sprat machines are attached to one adapter. In addition there is a select line, an attention line, and a reset line for each Sprat.

The commands were more concise and powerful than had been common even in disk stores of much larger size. For example, records up to 256 sectors long, and extending across several tracks could be read or written by a single 4-byte command. Disk defects, once demarked, were imperceptible to the host except as pauses in the flow of data.

The clock head, used in all earlier designs having an interleaved servo, was eliminated, with significant savings.

Modifications to the original proposal made during 1977 were:

- The original method of driving the motor from the microprocessor clock was found to give inadequate speed control, and was replaced by a bang-bang servo design evolved by Craft, Case, and Wallis and implemented in software by Craft.
- Because of demand by users for a "write-protect" facility, provision had to be made for rewriting sector addresses in the field.
- The existence of a "write-protect" facility had made feasible a proposal (as yet (1990) untested) by Wallis and Case to use the flag byte in the address field to store a correction to be applied to the following servo sample, thus reducing the effect of variation in the magnetic medium of the track-following performance. It was planned to introduce this measure if servo testing showed it to be needed.

During 1977 the architecture was exhaustively checked and refined by Ray Larner who was handling user liaison in Boulder, and who wrote and maintained a functional specification that was most detailed and thorough for a product in early development.

During the same period, a model of the store, using vendor transistor logic, was designed and built for Engineering Verification Test. The main functional units were designed and built by Frank Dickin. Working from sketchy functional descriptions supplied by Wallis, he filled in a good deal of detail, and resolved many problems in the architecture. The channel-control logic was produced by Case, and the other functional units by Ralph Bellamy.

M6800 software

By mid-1977, Dave Craft's ideas for the structure of the software were firm and had been accepted. The tasks it had to perform fell into two categories:

- "Housekeeping" tasks, which had to be done at regular intervals
- Tasks connected with demands by the host system, which occured at irregular intervals.

The first category included motor-speed control and algorithm-unit operation; the second, command analysis and access control. The latter will be referred to as "executive" tasks.

The first structure to be considered was interrupt-driven, whereby executive tasks would be interrupted at a certain time in the sector, and the processor redirected to the housekeeping tasks, returning to the executive tasks at the point where it had left off. This idea was dropped, because although it made good use of the microprocessor hardware, it was thought to be difficult to debug. Errors caused by interactions between tasks would have been hard to reproduce.

Instead, the housekeeping tasks were assigned to a "superior" routine, which was called at fixed intervals by whichever executive program was in control, and which returned control to a point specified by that program. The supervisor contained two branches of slightly different execution times, and the branch taken was determined by the time of arrival of the address mark recorded on the disk. Thus the operation of the software was synchronised with the disk format.

The supervisor itself was written by Craft, the executives concerned with data-channel control by Dickin, access control by Ian Henderson, and command analysis by Mike Wenman. In addition, Wenman supplied routines to facilitate debugging.

Devkit microprocessor application system

Microprocessor systems differ from large computer systems in that, once they are fully developed, they generally have no console facilities and their programs are held in read-only storage. In other words there is no way for an operator to modify either the program or any other stored data. During development, however, these facilities are essential.

Microprocessor manufacturers allow modification by supplying special apparatus for use during the development phase. The device supplied by Motorola for their M6800 was called "Exorciser", and one was delivered to the Sprat project in February 1976. In use however, it was found to have certain shortcomings, and these led Dave Craft, shortly afterwards, to propose a scheme that came to be known as Devkit.

Devkit consisted of a set of modules that could be put together to form a console for a microprocessor system, or to form the microprocessor system itself. The modules were:

- A block of eight keys
- A block of eight LED matrixes for displaying hexadecimal characters
- An adapter card for the key blocks and display blocks
- A card containing an M6800 microprocessor, a small amount of random-access storage, and the necessary clock circuits
- A card containing 2KB of random-access storage
- A card containing standard serial-interface circuits and other circuits immediately associated with the SSIF (standard serial interface) card.

The SSIF card could, among other things, communicate with an ordinary audio cassette recorder.

The architectural design was done by Dave Craft and the detailed hardware design by Tom Winlow. Development took place between mid-1976 and mid-1977, and final release was in October 1977.

Devkit found several uses on Sprat and on other Hursley projects. In addition to being used on the store itself, it had applications on the disk surface analyzer, on the servowriter, the test robots, and the wear-test rigs.

Servo system

The servo system in Sprat had two main functions:

- Track seeking; positioning the read/write head until it was over the desired track
- Track following; keeping it there in spite of contrary influences.

Although development work on these two aspects took place concurrently, it is simpler and clearer to present them separately. Track-following is the more straightforward operation of the two, and is considered first.

Track-following

Background

The read-write head is attached to one end of the actuator arm, which is pivoted near its centre of gravity. At the other end is fastened an actuator coil (made by printed-circuit techniques) which moves in the field of a permanent magnet. The effect of passing a current through the coils is to apply a torque to the arm, and the resulting movement of the head depends on the waveform of the current and on the mechanical properties of the arm, including the effects of resonances.

At the boundaries between successive sectors of a track, the head picks up a signal whose magnitude and polarity are a measure of the lateral displacement of the head from its ideal position. This position-error signal (PES) thus has the nature of a sampled signal, and the servo system of which it forms a part is a sampled servo system.

The PES is taken to a device called the controller, which operates on it to produce the appropriate current waveform to apply to the actuator. Controllers can take two main forms. They can consist, on the one hand, of analogue circuits whose gain and phase shift are some required function of frequency. Depending on complexity, the controllers are known as second-order, third-order and so on. Generally speaking, the more difficult the specification a servo has to meet, the higher the order of controller it needs.

Designers have to ensure that the overall system they produce is stable. Two measurements indicate the degree of safety in this respect: the gain margin and the phase margin, generally referred to as the stability margins.

Early work with analog controller

The earlier system, purely analogue in nature, was built by Bill Case and Ian Henderson in March 1975. The PES in this system was in fact continuous, but it was sampled at regular intervals with a strobe pulse to produce the same type of PES as would be obtained later in the project. The samples were passed through a "hold" circuit and the resulting signal applied in a second-order controller built from operational amplifiers and active filters. The stability margins were found to be too low, and the track following not precise enough.

Later in the year (June, July 1975) the second-order controller was replaced by one with a fourth-order characteristic, and by August its performance was considered satisfactory.

Digital servo

By this time there was much interest in replacing the analogue servo by a digital one. The successive position-error signals would be converted from analogue to digital form and handled by a microprocessor.

The initial work was to test various control algorithms by mathematical simulation. This was done mainly by Ian Henderson and Bill Case during late 1975 and the early part of 1976. The simulation, written in APL, ignored the effects of computation delay, thought at the time to be insignificant.

The next step was to test the algorithm on an electrical model of the actuator arm. The advantage of an electrical model over real mechanical hardware was that initial conditions would be set precisely and repeatably, making it easier to debug the control program and to understand the cause of any instability.

The first electrical model, built by Chris Wallis between February and April 1976, was a particularly simple one consisting of two integrators. It ignored resonance effects in the arm. Ian Henderson later refined this model so that the arm resonance was included.

Development of control algorithms continued, account being taken of computational delay by using the technique of modified z-transforms. During part of this period, July to September 1976, the project had the assistance of Dr. V. Latham of Southampton University.

By October 1976 real hardware (the "servo rig") was being controlled. The program, a fifth-order one, was being run on the Motorola M6800 Exorciser using stored tables to speed multiplication, but before the system could be made to work, computation time had to be reduced still further. Ian Henderson did this by refining Dave Craft's original programs though even when that had been done, in October 1976, some concern over stability margins remained.

It had now become clear that, despite the various measures taken to increase speed, the M6800 microprocessor was still not fast enough. The solution to the problem was to build an "algorithm unit" to relieve the M6800 of the computation load. This replaced the Exorciser around July 1977, giving satisfactory results.

The only other notable point in relation to track following concerned the resonant frequency of the actuator arm. In February 1977 this was found to be too low, and a redesign of the arm by Mike Nettles increased it to 2.7 kHz. Further work (April 1977) showed how, if necessary, it could be increased to 3.3 kHz but the 2.7 kHz design was retained at that time. It is important in practice to keep the resonant frequency within well-defined limits so as to avoid the ill-effects of "aliasing", a phenomenon present in all sampled-data systems.

Track seeking

General method

The general procedure for moving the head from one track to another is:

1. Apply a current to the actuator coil for a certain period to accelerate the arm to the required track-crossing velocity, chosen to be 1 track every two sector periods.
2. Switch off the current, allowing the arm to coast.
3. Apply a reverse current to the actuator to decelerate the arm.

Early work with analogue controllers

The initial work, done in the spring of 1975, used the second-order controller. In this early arrangement, the controller was operative only during the "coast" phase; during the accelerate and decelerate phases control was "open loop".

With the chosen value of track-crossing speed (see above), position-error signals were available only half as often during normal track following. This fact had to be allowed for by switching the controller characteristic appropriately.

The scheme was working in May 1975, giving a seek time of 100 ms. Later in the year, the fourth-order controller that succeeded the second-order one was being used in a similar way.

Digital scheme

As explained, interest at this stage of the project shifted to digital control, though for much of 1976 the work concentrated on track-following rather than seeking.

Early in 1977, however, the appropriate algorithms were being developed. They were closely modelled on the earlier strategy used with the analogue controller, control remaining open-loop during acceleration and deceleration. Also, to allow for the lowest sampling rate of the position-error signal during coasting, the control algorithm used during that phase was third-order.

By March 1977 the programs were being successfully run using the M6800 Exerciser and the electrical model of the arm. The system continued to work when the model was replaced by read hardware during May and June 1977, though the tolerance margins were less than one would have wished.

A few months earlier, work had been started on adding a degree of sophistication to the track-seeking process; the aim was eventually to speed it up by a factor of about 3, that is, to about 30 ms. It would have been impossible to do this successfully with the servo information being recorded on the disks at that time, as the position-error signals were useful only while the head was within plus or minus one quarter of a track width from the ideal "on track" position.

A technique for overcoming this problem already existed. Originally invented by Bob Commander and John Taylor of the Hursley Storage Development department, it consisted of dividing the servo information at the sector boundaries into two portions, one following directly after the other. The information pattern recorded in the first portion produced the "normal" error signal, and was useful within plus or minus one quarter of a track from the "on-track" position as already mentioned. The pattern recorded in the second portion produced what was called the "quadrature" error signal, and was useful within plus or minus one quarter of a track from a head position exactly half-way between two tracks. Thus, by choosing either the normal or the quadrature PES, by means of a strobe, a useful PES was available at any lateral position of the head.

This immediately brought two important benefits:

- During the coasting phase, one was no longer restricted to a single PES sample every two sectors; a useful PES could be obtained once per sector, giving the same degree of control as during normal track-following.

- The servo loop could remain closed during acceleration and deceleration giving the same degree of control during the phases as during coasting.

An additional benefit was that one control algorithm served during track following and during all three phases of track-seeking, reducing the load on the microprocessor and the amount of program storage it needed.

The above scheme, known as "controlled seek", was being successfully checked by digital simulation around April 1977 using a fifth-order control algorithm. The electrical model of the arm was extended to give the quadrature PES around July/August 1977, at which time also the necessary microprocessor code was being written.

By then, as recounted in this report, the algorithm unit had become available, and it was being successfully used with the newly-written code and new arm model in October 1977. A few refinements were added later, and track-seeking with real hardware was achieved in December 1977.

Although controlled seek had the potential for speeding-up the seek process as mentioned above, it was decided to leave this improvement until later in the project.

Servowriter

Introduction

It was explained in other sections of this report that special patterns of magnetic transitions have to be written on every track at the sector boundaries, for the purpose of providing a position-error signal. Writing these servo patterns is part of the disk manufacturing process and is done on a specially-built machine called a servowriter.

Work on developing this machine started in April 1976 and, with many features being added to the specification as development proceeded, it continued up to the time the project was transferred to Boulder. The essential requirements of the machine were mechanical stability and accuracy.

Basic Scheme

Basically, the servowriter consisted of the following:

- A drive motor to rotate the disk on which the servo patterns were to be written. The drive shaft ran in an air bearing.
- Means for determining the rotational position of the drive shaft. This was done by mounting a Perspex disk ("clock disk") with a pattern of opaque and transparent regions rather like an optical strobe disk. The disk had three sets of "strobe" marks: the innermost with one mark per revolution, the next with 72 marks, and the outermost with 1800 marks. The marks were sensed by a lamp and photocells.
- A fixed head for writing a clock track on the disk.
- An arm carrying the servo-writing head, capable of being positioned very accurately. The arm was moved by a dc motor coupled through a leadscrew to a linearly moving carriage. The arm carried a quadrant, and the carriage was

linked to it through steel tapes, so that the angular movement of the arm was accurately proportional to the linear movement of the carriage. The position of the carriage was determined very precisely by means of a Hewlett-Packard type 5500A laser interferometer. Moving the carriage to a required position was done by means of a servo loop. The interferometer gave actual position in digital form; this was compared with the desired position to produce a digital error signal, which, after conversion to analogue form, drove the dc motor.

History

Work on the servowriter began in April 1976, with Pat Mulholland responsible for mechanical design, Bill Case for the phase-locked loop system for generating clock pulses, and Frank Dickin for logic circuits, and the remaining electrical matters.

The machine had been built in its initial form at about the end of 1976, at which time Bob Cope arrived from Boulder. He and Frank worked together until March 1977 when Frank transferred to other work leaving Bob carrying the main servowriter burden.

The initial machine suffered from various problems. One was that imperfections in the clock disk prevented the phase-locked oscillator from synchronizing reliably. This was overcome by taking greater care in manufacture, and by protecting the disk with a dust cover.

A more severe problem was stiction in the leadscrew, calling for a high starting torque from the dc motor. A temporary palliative was to add a small ac component to the motor drive, but for a long-term solution the problem was tackled along two lines:

- Mechanical measures to reduce the stiction itself, and
- Electrical measures to reduce the ill-effects of the stiction that remained.

Bob Cope discovered that the stiction was aggravated by the way the dc motor, leadscrew, and carriage were assembled. He greatly improved matters by placing shims between the rotor and stator of the motor during assembly.

The electrical measures were developed by Bill Case. As used here, the motor could have been regarded as an integrator: the effect of applying a known voltage was to produce a known torque, and the "output" of the motor that one could easily measure was its velocity, which ideally was the time integral of the torque. Stiction manifests itself as a degree of non-linearity in this element, and the way to reduce its effect was to place the non-linearity in a high-gain feedback loop. The velocity of the leadscrew was measured by using a tachometer, differentiated so as to negate the integration effect of the motor, and the resulting voltage subtracted from the "desired" motor-drive voltage as provided by the rest of the system. The difference signal was amplified, and it was this that formed the actual drive to the motor.

Using these two methods, the problem of stiction had been overcome by about February 1977. A substantial number of wiring errors had been found in the logic cards, and by the time these had been corrected, April 1977, the arm-positioning servo was working well.

While the corrections were being made, a start was made on adding various features. The first of these was the ability to write quadrature servo patterns, which called for the design of five more circuit cards. Work was started in February 1977 and the first disk with quadrature information was written about three months later.

In April 1977 the requirement arose for the servowriter to write, in addition to the servo information, a special information pattern at every address on the disk. (The process is known as formatting.) The information consists of :

- A preamble that includes several bytes of all-ones for setting the AGC circuits and synchronizing the phase-locked oscillator in the data channel, 5 bytes defining the address of the particular location, and 2 bytes of check bits
- Following the preamble, 256 bytes of all-zeros, 2 further bytes of check bits and all-zeros to fill the remaining space.

The first formatted disk was successfully written in August 1977.

During May, Bob Cope started adding a Devkit microcomputer. Its purpose was generally to control the servowriter, and to compute addresses for formatting. Four additional logic cards were needed; they were produced during May and June, and debugging started in July, with Lyle Fangmeier's assistance.

Though Devkit was well capable of computing addresses, it was not fast enough to compute the check bits. This was done instead with special logic designed during July 1977 by Ralph Bellamy. Devkit programming continued into the late part of the year with valuable assistance by Charlie Farel.

During July and August, two further features were added. The first was to make the servowriter write a particular servo pattern covering all the space between the outermost track (track 0) and the outer edge of the disk; its purpose was to ensure that if, during the actual operation, the head moved out beyond track 0, it would pick up a signal indicating the fact so that corrective action could be automatically taken. The second feature was to vary the write current with arm position. It had been found that over the outer 125 tracks the best value of write current was 30 mA, and over the inner 125 tracks, 25 mA. The change was arranged to take place automatically under Devkit control.

In addition to the electrical features mentioned above, 1977 saw a number of mechanical changes. During the early part of the year it had been found impossible to servowrite disks satisfactorily, mainly because of variation in flying height at different track radii. At that time the servowriter had only crude means for adjusting the head positions, but during April 1977 proper micrometer adjustments applied only to the servowriting head; the clock-track head, being less critical, was left with the earlier method of adjustment.

A second measure, aimed at obtaining more stable head-flying conditions, was to mount both heads in stabilizers as was done on the Sprat machines themselves.

Several other mechanical changes followed. They included modifying the machine so that disks could be servowritten after being mounted in the cartridges, building a jig for setting the servowriting and clock-track head positions and making a further modification to allow for the increase in disk diameter. The first two were done in May, and the third in July.

Finally, after the many improvements and changes described, the twenty disks needed for the Engineering Verification Test were written on 6 December 1977.

Burnishing

At the start of the initial feasibility study (1974), it was expected that burnishing of the disk surface would be necessary as all previous IBM media had required this smoothing process. A microscopic study was made to the "as received" surface of Crolyn, and later, Pegasus. This showed a large number of mounds of more than 40 microinches. Although the head-disk interface was known to be "forgiving", with a flying-height goal of 8 – 10 microinches, it was thought that mounds of these sizes would create wear of the head, generate particles, and produce low output signals as the mounds lifted the base surface of the medium away from the head.

Existing methods of burnishing applied to iron-oxide coatings were examined, and that used for the Igar process tried. This used a round, flat-faced, hard-steel burr which was slightly larger in diameter than the data-band width on the disk. The flexible disk was supported on a rubber-faced plate which was rotated. The rotating burr was positioned over the data band and impressed into the rotating disk, so that the sharp edges cut away the mounds.

On chrome dioxide media, however, there was very little cutting action, only smearing of the mounds, and the edges of the burr became blunt after every disk was burnished. Other cutting profiles were tried without success and the steel burr was abandoned.

Trials were made with a copper burr into which diamond particles had been impressed, but the surface quickly became loaded with media particles and scratched the disk surface badly.

Further trials were made with toothed burrs in which the teeth were made of Corundum (aluminium oxide); these tended to cut and smooth correctly, but still required frequent sharpening. They reduced the number of "anomalies" on a disk surface by a factor of approximately 10.

In July 1976 experiments were made with alumina-coated plastic films and, while a fresh surface was continuously offered to the disk, the mounds were reduced to the basic surface finish of the disk.

A machine was designed and made, and the process demonstrated to the Boulder media group, who eventually adopted it, applying it to all disks made of Pegasus material.

The above work was performed by Alan Cox and Leo Rigbey.

Disk surface analysis

Purpose

Disk surfaces generally have a number of small defective areas in which data cannot be reliably recorded. The way of coping with this difficulty is to mark as unusable any sectors in which the defects occur and to provide some spare sectors to be used in their stead. Another technique, useful if the defects are not too large, is to use error-correcting codes.

A feature of Sprat was that its disks were cut from inexpensive tape stock that could be expected to have a fairly substantial number of defects. This would call for a corresponding number of spare sectors. To gauge the magnitude of this problem, and to monitor disk quality generally, an analyzer was built for collecting statistics on the number of defects, their severity, and their detailed nature.

Defects fell into two types, according to the way in which they manifest themselves:

- Missing-bit defects

  As the defect reaches the head, signal amplitude falls, generally over several bit periods, to a value which may be too low to operate the detector circuit. As the defect passes, the signal amplitude rises again until it regains its full value.

- Extra-bit defects

  These are caused by sharp discontinuities in the magnetic properties of the coating, caused possibly by the inclusion of foreign material. The foreign material has a value of magnetization different from the coating proper, and as it passes the head a spurious output pulse is induced.

The surface analysis tester (SAT) built for Sprat, analyzed the disk surface for both these conditions. It examined the disk track by track, carrying out the following procedure for each:

1. Regularly-occurring transitions were written at a frequency some 30% below the highest that occurred in use.
2. The track was then read. At every type-1 defect, the apparatus recorded the rotational position of the disk at which the output passed the following thresholds: 90%, 70%, 50%, 30% and 10% of the normal value in the immediate vicinity of the defect. Thus, at a defect that caused the output to fall to say 60%, four rotational positions were recorded:
   a. The position at which the output fell to 90%
   b. The position at which the output fell to 70%
   c. The position at which the output recovered to 70%
   d. The position at which the output recovered to 90%.

   Other information collected during this read operation was the maximum value of the signal envelope, the minimum value ignoring defects, and the mean value.
3. dc was applied to the head for one revolution, to erase the information written during step 1.

4. The track was read a second time, and the rotational position recorded of any signal peak that exceeded a set proportion, generally 40%, of the mean value obtained during step 2. This gave the position of the type-2 defects.

History

The decision to build the SAT was taken at the end of 1975, and the work was done during the first three months of 1976. The overall design was by Bill Case, with John Holloway and Frank Dickin attending to the circuit details and Mike Hatchett to the mechanical design.

As originally built, the information collected was recorded on a tape cassette which was taken to Hursley and fed into the TSO system there. The analysis programs for operating on it were written by Chris Wallis during the rest of 1976.

The machine worked well apart from one problem: over the outer 100 tracks, the disks in use at the time (155-mm diameter, mounted in the prototype cartridge) would not fly over the head with the necessary stability. The problem persisted until the larger disk (166-mm diameter) was substituted. By the autumn of 1976, this and other mechanical improvements gave reliable operation nearly to the outermost track, but it was not until tolerances on the backplate were tightened up that operation was reliable over the whole disk. This was in January 1977.

A second problem that existed during the early days of SAT, was that of entering the recorded data into the Hursley TSO system. It was caused mainly by unreliable operation of the cassette machines. The solution was developed by Bill Case and Dave Craft, and was working well around the middle of 1976. Instead of recording the SAT data on a cassette, it was sent via the Motorola M6800 Exorciser and a data link directly to the Hursley APL system. This arrangement was used for about a year, at the end of which time the Exorciser was replaced by Devkit hardware installed within the SAT itself.

Various improvements to the software were made during the second half of 1977. Bill Case and Dave Craft added various checks to the Devkit software, and Chris Wallis and John Holloway improved the data-plotting procedures. Initially, the results of the analysis had been plotted on Cartesian coordinates, as if the tracks had been opened out straight. This was replaced by a plot on polar coordinates done at full scale, giving, in effect, a picture of the disk with the defect positions marked. By punching holes where defects were indicated, and laying the plot over the disk itself with the correct orientation, the defective areas were pinpointed and could be examined under a microscope. This was being done in October 1977.

Mechanical design — machine and cartridge

Early work was related to the various rigs for disk-flying measurements, data channel, and surface analysis testing from which the machine requirements crystalized. These were sufficiently firm by November 1976 to enable the cartridge location method, actuator layout, and drive motor concepts to be established, and four models with a base casting were built by February 1977.

The cartridge loading mechanism was then developed through four iterations with Hursley Industrial Design providing the detailed appearance of the bezel.

The design of the machine was further refined during June 1977 to incorporate the pre-amp – write-driver card-connection lead to the head, and the final cartridge loading mechanism. New machines were made and assembled in September 1977 for EVT debugging and test.

Early cartridges were machined from the solid, but with the design firm in December 1976, mould tools were started to provide the large numbers required later in the programme. Both cover halves were produced at a single moulding shot during June 1977.

During the whole of the project, Industrial Design was involved in the appearance and human factors aspects of the cartridge, finally producing a model in December 1977 that incorporated all the best features of previous designs and ideas such as labelling, handling, stacking, and finish for good wearability. These points were then passed to Boulder for continued effort.

Those closely associated with all the above work were: Allan Cox, Mike Hatchett, Nick Leon (Industrial Design), and Leo Rigbey.

Store-testing robots

Testing a new design of disk store called for a number of separate procedures. Many detailed mechanical and electrical measurements were made, but an important overall test was to operate the store under conditions that resembled, as closely as possible, actual use in the field. This entailed writing various information patterns to the various locations on the disk, and later on reading them back and checking that they were not corrupted. In a thorough test the number of write and read operations becomes very large, making a test robot a necessity.

The Sprat engineering verification test (EVT) took place early in 1978 and was on four separate machines, with four test robots. The task of designing and building them was allotted to Lyle Fangmeier who was in Millbrook on assignment from February to December, 1977.

His design consisted of a 5100 desk-top computer, a microcomputer constructed from Devkit modules and six specially-designed circuit cards. The Devkit was connected to the 5100 over a serial interface and to Sprat through the latter's normal customer interface. No connections were made to any internal part of Sprat.

In use, the 5100 was loaded with various test routines written in APL by the person specifying the test. For instance, a routine might call for the checking of every sector along a particular track. Several routines could be strung together to form what was called a Command List.

When a test was to be made, the specified command list was transferred from the 5100 to the Devkit microcomputer, which proceeded to execute it, checking the information as it was read back. If an error was found, the Devkit brought the 5100 into play again, obtaining from it a new command list whose purpose was, if possible, to enable recovery from the error to take place. Devkit then jumped to this error-recovery command list, and if it could be completed successfully the original command list was resumed at the point where the error was found. This method of nesting command lists could be extended to higher levels if further errors demanded it.

Design of the test robots took place during April and May 1977. The hardware was built between June and August, and the software written between August and November. Debugging was taking place during December but had not been completed by the time the project was transferred to Boulder. The system was fully working, however, at the beginning of February 1978.

Wear-test rigs

One of the aims in the design of Sprat was to use a head-flying height only just great enough to avoid excessive head and disk wear. Flying conditions during track-seeking are not in general the same as during track-following, and so it was important to make wear measurements for both situations. The wear-test rigs were built for this purpose.

To keep the amount of design effort to a minimum, the rigs were made as far as possible like ordinary Sprat machines. One important difference, however, concerned the servo. In general, the disks being tested would have no servo information on them, and so arrangements had to be made for producing a pseudo position-error signal for use during the coasting phase of track seeking. This was done by mounting on the actuator arm a Perspex arc carrying a sequence of opaque and transparent areas. These were sensed by a lamp and photocell, so that while the arm was coasting at constant velocity, a signal consisting of a dc component and a constant frequency ac component was produced. The dc component was removed and the residual ac sampled, nominally at the zero-crossing points, to produce a pseudo PES that could be applied to the controller. As explained in detail below, the wear-testing rigs were built in 1976, and all other details of the servo system – construction and performance – were exactly the same as on the Sprat machines of the time. In use, the head could be made to seek repeatedly over any number of tracks from 1 to 256.

The decision to build the wear-test rigs was made in June 1976. Bill Case and Ian Henderson were responsible for the electrical design, Mike Hatchett for the mechanical design, and John Holloway did some logic design and saw to the building. Five rigs were built at Millbrook during August 1976, and were complete and working in September. Mechanical parts and circuit cards for six further rigs were built around November 1976, and these, together with design information, were supplied to Boulder at about the end of that year.

The machines, in the form described, continued working satisfactorily up to the time of writing this report (February 1978). Nevertheless, in November 1977 John Holloway was asked to look into the possibility of controlling them with a Devkit microcomputer so that instead of repeatedly seeking over the same number of tracks, they could make a succession of different length seeks. John produced a scheme for controlling eight rigs with one Devkit, each of the rigs running a different program. Using a modified version of Lyle Fangmeier's test robot, the scheme was working before the project moved to Boulder.

Track misregistration study

A potent source of errors in a magnetic-disk store is failure of the head to register exactly with the track it is reading. There are several reasons why this happens:

- Mechanical imperfections in the servowriter cause the track, as defined by the servo information, to depart from a perfect circle.
- Imperfections in the disk drives cause the centre of rotation during reading to differ from that during writing, and both to differ from the mean centre of rotation during servowriting.
- Anisotropy of the recording medium causes the written track to become oval as distinct from circular.
- The head may be wider than the track it is reading, partly because of head-width tolerances, and partly by the possibility of the edges of the track being obliterated when neighbouring tracks are written (track squeeze).

The track-following servo system reduces most of these effects, but it does not eliminate them altogether, and it does nothing to reduce the effect of head-width tolerance. One therefore needs information on the degree of misregistration that remains, and on the error rate that it causes.

Matt Taub studied this topic in the Autumn of 1977 using a Monte Carlo simulation program specially written for the purpose. Ideally the input to this program would have been obtained by making accurate measurements on a fair-sized sample of machines, but at that stage of the project too few had been built. Instead, estimates of the mechanical imperfections were supplied by Leo Rigbey, Mike Nettles and Tony Hearn, and measurements of the servo performance, based on the available hardware, were supplied by Ian Henderson. Information on disk anisotropy came from E Bartkus at Boulder, and from a published paper by Greenberg and others.

The most significant result obtained was that most of the residual misregistration was caused by mechanical imperfections in the drive-spindle bearings, thát is lack of perfect smoothness in the ball races, and the balls themselves being out-of-spherical. These gave rise to spindle movements at frequencies up to several hundred hertz, and at the upper end of the range, say 300 to 500 Hz, the current design of servo did not help, in fact it accentuated the effect. Improvement of the servo performance in this range, even at the cost of a deterioration in the 25 – 50 Hz range, would probably have given an overall improvement.

A second result of the study was that, with the current servo performance and mechanical tolerance values, it would have been an advantage to decrease the electrical width of the head from 0.0038 to 0.0037 inch. At this figure, and with the error-rate information then available, overall error-rate caused by misregistration should have been well below 1 error in 10^{10} bytes.

Conclusions

The aim of the Sprat project was to develop a small inexpensive disk store with a performance closely approaching the limits of what was possible with the available materials and technology. The main body of this report dealt with the many different activities that had to be pursued, and has shown the high degree of inventiveness and perseverance that the team displayed, and the high degree of success that it achieved.

Even on the problems that were not solved completely, head fabrication and head-disk flying, much progress was made in understanding their nature so that solutions were then nearer. There was every reason to believe that concentration on these problems would soon bring satisfactory solutions, and that Sprat would develop into the product that many groups in IBM eagerly awaited.

Bluegill at Boulder and Tucson January 1978 to September 1979

This section was written by Chris Wallis.

EVT Experience

The previous report covers the period up to the arrival in Boulder in January 1978.

The original intention had been to hold an engineering test in Hursley in December 1977. It became clear towards the end of 1977, however, that the absolute requirement to transfer the project to Boulder on 1 January 1978 would lead to an unacceptably constrained EVT. It was therefore decided to perform the test on arrival in Boulder, as early as possible in 1978. The EVT was started on 3 February 1978.

The test was broken into four parts:

- Logic
- Drive system
- Servo and access
- Simulated customer operation.

In the logic test, the basic logical operations of the file were verified: namely, its ability to execute commands from the tester, to write, read, and access reliably in the absence of stress. This part of the test went very smoothly, once the interface to the tester and the tester microcode were debugged.

The organization of the file into microprocessor-controlled functional units made debugging of the file logic itself very easy. No hardware changes were required in test, all the faults being handled by microcode changes. This aspect of the Sprat design must be judged to have achieved its purpose very successfully.

The test of the drive system was also generally successful, in that the motor speed was found to be controlled well within the specified limits. Difficulty was experienced with the initial track acquisition on startup, because of mechanical variations between machines. Palliatives were applied in the microcode, but a final solution was deferred until AVT design.

The servo test was less successful. The difficulties were mainly caused by the quality of the servo information on the disks, which led to noisy position-error signals and poor track-following. Efforts to improve the quality of servo-writing were only partially successful, and the simulated customer operation had to be performed with reduced servo gain.

The effect of this shortcoming in the servo-writer was to delay exit from the test, to restrict the amount of testing possible, and to camouflage any other failures in the track-following and access performance that might have existed.

Another weakness in the EVT was the shortage of usable heads. At that time there was difficulty in fabricating heads with acceptable output and wear characteristics. The EVT was performed with a very small sample of heads, some of which had low output, and some of which inflicted mechanical damage on the recording medium.

Thus it was hard to know, at the end of the test, whether the errors experienced would be typical of the eventual product or not. (These problems in head manufacture were later solved, and the AVT results to date look very encouraging.)

During the test, a total of 18 348MB were read, including 5 259MB written and read on different drives. A total of 1 396 errors were recorded of which 45 were repeated on the following revolution. This corresponds to a "soft error rate" of 1 error in 10^6 bytes. It was not possible to say from the EVT results what the hard error rate was, because the only demarking technique available at the time was not reliable, and many or all of the repeated errors that occurred may have been on sectors, that ought to have been demarked.

AVT design

Before the EVT was over, it was necessary to start the AVT design. It was during the AVT design period that the move to Tucson began. The disruptive effect of the move and of dispersing the design team between two locations made it necessary to delay the scheduled start of AVT by three months.

The AVT design had to be organized differently from that of the EVT machine. For EVT, the microcode and the hardware design for each function was done as far as possible by the same engineer. This made for economies of design and few problems of communication within the design team. The AVT schedule and the complexities of LSI design ruled out this approach for AVT. The newly-joined team of logic designers, (in order of joining - Uma Reddi, Dave Cornaby, Debbie Collins, Ernie Gurule, Chris Pisciotta, and Randy Zwetzig) working from the rather sketchy EVT documentation and under the tutelage of Frank Dickin, managed to gain an understanding of the machine functions and start their design in a commendably short time. However, remarkable as this achievement was, it was to be expected that latent misunderstandings would persist and that there would be design errors that would not surface until debugging time, errors that would not have arisen given continuity in the design team. There were not many such problems, to the great credit of the group. A number of other handicaps, which will be touched upon in the relevant sections were, of course, imposed on the development through its being involved in the startup of a new site.

During the latter half of 1978, the LSI chips were released, and the format was defined. In September 1978, the microcode design was begun. Power was on the first test frame on 1 January 1979, and the first mechanical AVT level hardware was available for debugging on 6 January 1979. The first servo-written cartridge was produced 6 January 1979.

Throughout the design and release process, serious difficulties and delays were encountered in the use of the EDS system. They were caused by start-up difficulties at the Tucson D.P. Centre; shortage of staff in the EDS center of competence; software bugs in EDS itself; and difficulties in writing the EDS rules defining the logic and microprocessor in use.

Although the scheduled power-on date was met, there was by this time a considerable backlog of card debugging still to be done and, when the debugging of the machine as a whole got under way, it was found to be a very much slower process than had been anticipated. With hindsight, the reasons for this appear to this writer to be:

- Chiefly because of EDS delays, the bulk of the cards were yellow-wired rather than etched, with the consequent inevitable wiring errors. Those that were etched were of very poor quality, with many open and shorted lands.
- The card packaging scheme using 6-high 4-wide cards (for easy conversion to LSI) in a cooling shroud, had the disadvantage that plugging and unplugging of cards into the board was very difficult. Many faults were caused by dropping cards, bending them, or shorting them together. Furthermore, this made the debugging team chary of using extender cards and probing on the cards as they would otherwise have done.
- As mentioned above, the logic designers who performed the debugging had not had the opportunity to gain an understanding of the machine as a whole, however well they understood their own area. This made for confusion and delay in assigning problems to their real origins.
- In an effort to retain flexibility of function, Psalm PLA technology was used in some critical areas of the machine, notably in the format unit. While this flexibility proved of value in product test, the interface between the synchronous and asynchronous logic was a copious source of obscure problems, many of which would result in corruption of the content of the Psalm arrays. Again with hindsight, it would probably have been better to implement the format unit in asynchronous logic, and carried its program in ROS.

For these reasons (perhaps foreseeable, but not foreseen), the debugging process took very much longer than planned, in spite of the heroic efforts of the debugging team. In the course of the debugging, a more thorough understanding of the function of the machine was gained — particularly by Gurule and Pisciotta, who became expert in many areas.

Because of the protracted debugging phase, AVT entry was rescheduled from 2 April to 14 June 1979. The acceptance test was passed on this schedule. Because of a microcode bug, the data reliability could not be estimated. However, in the following acceptance test for the Sabino EVT, which used the Bluegill AVT models without change, the reliability statistics were as follows:

Bytes written	$1.52\text{x}10^9$
Soft write errors	13
Soft write-error rate	1 in $1.17\text{x}10^8$
Bytes read	$6.28\text{x}10^9$
Soft read errors	zero
Soft read-error rate	not measurable

Users

At the time of the Stage 1 review for Bluegill in May 1978, the first using systems were the Orbit Venus (8140 Processor). Both systems were scheduled to announce a Bluegill version in 9/79, with FCS in 9/80. Inca (4331) and Aztec (4341) were also listed as potential Bluegill users, as were Mistral, Store systems, the SPD 'D' Series, and a number of GSD systems.

Later in 1978, it became clear that the 4.5MB Bluegill would not meet the requirements for Inca and Aztec, which had an entry DASD requirement at the 65MB level. Bluegill also lacked sufficient capacity to be an off-load device for these systems, which typically had several hundred megabytes of DASD.

In early 1979 the 'D' Series of systems, as originally defined by SPD, were terminated. Mistral ran into cost problems in the second quarter of 1979 and had to remove function to save cost. A number of functions were eliminated, and the DASD requirements were redefined so that a Pilgrim diskette could be used.

In late 1978 GSD obtained a Corporate charter to develop their own DASD. As a result of this change in direction, GSD systems were committed to use GSD file products and would not commit to use Bluegill.

Store systems had also decommitted from the use of Bluegill because they felt that they needed upward capacity growth, which was not committed for Bluegill. Also Bluegill was physically too large to fit into their current package.

Alphaleonis (a low end 8100 scheduled to ship June 1982), the Freedom banking system, and the Trout/Sierra system became new committed users. However, it became clear at this time that the 8100 software (DPPX) which would be used for Leo and Venus would not fit on a 4.5MB Bluegill. It also became clear that the resources to fix this software problem would not be available in time. As a result, Leo and Venus decommitted as Bluegill users in July 1979.

With Leo and Venus gone, none of the remaining users of Bluegill (Alphaleonis, Freedom, or Trout) was scheduled to ship before mid-1982. Bluegill had been scheduled to ship 10/80. In addition, all of the current users and a number of potential users for Bluegill were expressing a requirement for a device similar to Bluegill but with two or three times the capacity.

The Bluegill program was therefore terminated and work began on the definition of Sabino. The initial objectives for Sabino were very similar to Bluegill, but the capacity objective was 15MB.

Staffing

Up to November 1977, the staffing plan had been to take with the project to Boulder six U.K. assignees, representing key skills, together with the eight returning U.S. assignees. From the time the Tucson move was announced, it was clear that there would be difficulties with staffing. Several of the U.S. assignees in Hursley were reluctant to move to Tucson. In the end, only two of them did, four electing to stay in Boulder with OPD, one moving to San Jose, and one leaving the company. To fill this gap, four more assignees were brought in from Hursley.

Some attrition among those recruited to the project from the Boulder laboratory was also serious. Of the 46 Boulder employees who were recruited to the program up to May 1978, six left the company, and 16 moved elsewhere in IBM. The situation was actually worse than these figures suggest, for three reasons:

- During 1978, the project was selectively recruiting people who were prepared to move to Tucson. This restriction excluded many of the more experienced and talented candidates.
- Of the six who left the company, four had critical skills, and one was one of the assignees returning from Hursley.
- Of the sixteen who moved to other parts of IBM, three had critical skills, and five were returning assignees.

Recruiting, and the problems of operating with serious gaps in the department's skills, imposed a heavy burden on project management at a critical time in the definition of the AVT machine. However, the project was very fortunate to pick up from a terminated program the services of four experienced logic designers, with their 3270 terminals, which proved to be in short supply in Tucson. Without this piece of good fortune, the AVT design would probably not have been done on schedule.

Once in Tucson, the problem became one of replacing the skills of the Hursley assignees before they returned. This was done in the end, but took longer than expected, and several of the assignments had to be extended. It was particularly difficult to find anyone with servo or servo-writer experience.

Architecture

Format

The EVT disk data format was not designed with attachment to a practical user system in mind. It consisted of 240 tracks, each having 72 sectors, each containing 275 bytes, giving a total data capacity of 4.7MB. None of the sectors were particularly designed as spares. The maximum track density was 250 TPI.

On 7 June 1978 the Bluegill Format committee was convened, under the chairmanship of Dick Scully. It made its recommendations on 10 July. The differences from the EVT format are:

Data field reduced to 256 bytes per sector
Gross sector length increased from 320 to 334 bytes
Sectors per track increased to 76, including spares
4 spare sectors per track
Recording band moved out to 50 – 75 mm
Reducing the bit density to 16.31KB per inch.

The increase in gross sector length was necessary to accommodate rotational tolerances in types of usage (such as recovery from particular defect distributions) ignored in the EVT design. The reduction in linear density was at the request of the channel group, and the provision of spares was calculated from a projected defect rate of 10^6 bytes between defects.

The bit and track densities chosen (16 300 BPI, 250 TPI) together represented a very great step forward in technology from any other floppy disk on the market, and put Bluegill among the leading IBM product developments in areal density.

There was intense debate over the issue of ECC. The use of ECC is conventional in tape products, but up to that time was not usual in low-end disks, although Piccolo has a rudimentary ECC to cope with small defects undetected during surface analysis of the disk at manufacture. Given the use of a verify operation following writing, only a restricted class of defects would cause loss of date ("sudden, hard, grown" defects). Moreover, the provision of ECC was thought to cost about 20 percent in storage capacity. These penalties would be imposed on all users, whether or not they required ECC.

A compromise was eventually adopted, whereby the benefits and penalties of ECC could be optional. Error correction was to be secured by the use of a procedure whereby the user would maintain on each track a sector (one of the spares) that contained in each bit position the odd parity of the corresponding bits in the 72 data sectors on that track. If a hard error occurred, detected by the CRC check, the data in the sector in error could be reconstructed by taking the odd parity of the other 71 sectors and the parity sector. The advantages of this idea were:

- Much longer errors (up to a whole sector) could be corrected than with any convenient ECC of practical length.
- Multiple errors in different sectors could be corrected if the same bit was not in error in more than one sector.
- No hardware was required in Blugill to implement it. It could be done entirely with user software.
- The overhead (the storage capacity required to be devoted to error correction) was extremely low: only 1.4 percent compared with nearly 20 percent for a conventional ECC.

The disadvantage of the scheme was its effect on write performance. Depending on record length, it would be between zero and 40 percent, if the user system could keep up with the processing required to maintain the parity sector. In practice, the processors could not keep up, and because of I/O bus limitations the effect was much more serious. As a result, none of the using systems elected to implement the procedure. It should be noted that this disadvantage would be minimized in any future application of the idea if the maintenance of the parity sector were done by hardware and microcode within the control logic, the cost of which would be comparable with conventional ECC hardware.

Command Set

Several commands were added to the AVT command set, and some were implemented that had been defined but never coded for the EVT machine. Significant improvements were made to the RAS facilities; the sense data was made more comprehensive; and automatic retry was added to the access and ID search routines.

Mechanical design

Drive

Most of the changes in mechanical design for AVT were made in the interest of reducing manufacturing cost. Casting changes simplified assembly, and the cartridge loading mechanism had fewer parts than the EVT version.

Some changes were made to remedy shortcomings either foreseen or experienced in AVT debugging.

The magnet structure was modified to increase the clearance available to accommodate warping of the actuator printed-circuit coil.

A small flywheel was added to the motor shaft to improve speed constancy under unfavorable power-supply conditions.

The actuator arm was changed from a channel section to a tapered tube (made as an investment casting) to improve the stiffness-to-inertia ratio and maintain out-of-plane stiffness.

Cartridge design

The AVT cartridge was extensively changed from the EVT design. The Industrial Design group recommended a sharper-edged design with an inclined front surface. The number of parts was reduced, the top and bottom mouldings were designed to snap together and the shutter was made of PTFE-impregnated polycarbonate to improve durability and avoid generation of debris. It was found that the addition of ribs to the inside of the top moulding had the effect of preventing rotation of the air mass inside the cartridge and the consequent reduction in pressure above the disk, which in the EVT cartridge had been the cause of the disk moving away from the plate. The vent needed on the EVT cartridge to inhibit this could therefore be eliminated.

To improve the penetration tolerance, the cartridge location method was changed. The locating pins on the base casting bear directly on the plate, instead of against an intervening layer of plastic. The holes in the bottom moulding that allow the pins access to the plate are closed by the shutter when the cartridge is not mounted.

Towards the end of 1978, it became apparent that vendors were having difficulty in producing the precision needed in bending the bent plate, and that, even when it was accurately bent, the long-term stability was doubtful. It was found too, in early debugging, that the plastic mouldings were capable of exerting enough force on the bent plate to deform it and cause penetration of the head into the medium. For these reasons it became urgently necessary to review the justification for the bent plate. It was found in tests done by the Tribology group that with the stabilized head, adequate control over the air-bearing could be maintained over a flat plate. The cartridge was modified to accept a flat plate, and it was shown that this design was easier to manufacture, and more tolerant of material instability than the bent-plate design. The servo-writer was modified to take the flat plate, and flat plates were used in the AVT cartridges.

Microcode

Control code

A significant difficulty faced by designers of the Bluegill microcode was the need for the microcode processing to be synchronized to the disk rotation. There were particular operations that the microcode had to perform within rather narrow windows in the format. This was met in the EVT design by imposing a rule that all paths through the monitor and executive task routine had to be of the same length. Although the EVT microcode was successfully written on this basis, the incoming microcode team felt that for the second pass, with the job better defined and understood, it was unnecessarily restrictive. An alternative was to use an interrupt-driven system, with interrupts generated by a hardware timer started from sector marks. An objection raised to this idea was that malfunctions might be hard to reproduce in an interrupt-driven system. The advantages of the proposal were: code would be much faster to write and to change if path lengths did not have to be equalized; efficiency of some essentially asynchronous operations, such as communication with host, would be improved; and a good deal of ROS could be saved by removal of padding.

The advantages were thought to outweigh the disadvantages and the AVT code was changed to interrupt-driven.

Another significant difference in the AVT code structure is the linking together of the various tasks involved in the processing of a command by means of a list processor. A command analyzer routine selects a list of tasks representing the command received, and a command dispatcher processes the tasks in the list, branching as required by return codes handed to it on completion of each task.

A substantial advantage of this approach is that the task lists are a compact representation in the microcode of the desired behaviour of the machine; its progress in processing them is easily followed in debugging; and the logic governing such things as retry procedures is easily understood and changed. However, the approach proved to be costly in microprocessor power, and should probably have been modified for SVT.

The microcode team was not assembled until late in September 1978; October was spent in organization of tools and in architectural design of the code; and coding did not start in earnest until November. Enough code was in place to support the start of the hardware debugging early in January 1979, and then the production of code was paced by the hardware debugging until early April, when enough hardware was operational to allow the microcode debugging to be completed to test entry level by 11 June.

Chris Wallis managed the microcode effort up to April 1979, when Jack DeLancey took over. Cliff Williams served as chief programmer and controlled the change level and integration of new microcode; Dave Patton wrote the command analyzer dispatcher, and most of the task code; and Joe Luciani wrote the Monitor and the Access code.

Microcode tools

The methods of assembling and loading microcode that were used up to EVT had worked well, but the team was worried about their continuing dependence on the 1130, and there were difficulties in controlling and maintaining the security of microcode stored on punched cards. Accordingly, Dave Craft and Paul Hundley designed a new system for the AVT coding cycle.

The code was to be entered from 3270s; assembled and stored under VM/370. Object code would then be transmitted to an IBM Series/1 located on the test floor and stored locally on disk, thus providing stand-alone capability in case of poor availability of VM. The Series/1 would drive a "microcode bus", a cable that could be routed around the floor to the various application microprocessors. The act of connecting to the cable would cause the Series/1 to interrogate the microprocessor and send it the appropriate code module.

In the end, only the VM assembler and library were implemented. Paul Hundley devoted four months to an effort to program the Series/1 part of the job, but progress was so slow that the team eventually decided to cut our losses and stay with the 1130 and audio cassette EVT system for microcode distribution. The difficulties with the Series/1 included: inadequate and inaccurate documentation of the operating software; hardware unreliability; and unannounced changes in formatting of data transmitted from Boulder via the "Codex" TDM unit.

As a consequence of the change to the 1.5 MHz MPU card, new console programs had to be written. Initially the old 1.0 MHz consoles were retained, but to conserve maintenance effort and economize on hardware, Dave Craft designed a system that allowed the console program to run in the application microprocessor. At the same time, he designed new display hardware to be driven serially from the special interfaces on the 1.5 MHz card. This "mini-console" system was installed on the AVT machines in test by 12 June 1979.

Control logic

The original SPRAT architecture was founded upon the idea that the control logic would be implemented in a high-density, low-cost technology, which might have a long design and process turnaround time, such as Emerald V. At the start of the AVT design, it was apparent that to use Emerald would not be practicable: there was no support in place for it in the PD group in Boulder; there was reluctance to put in that support; and, anyway, the schedule demanded by the primary user, Leo, would not allow the long design cycle. The decision was therefore taken by Jack Wells to forego the cost advantage of Emerald (thought to be about 60 dollars), and use whatever technologies would permit the schedule to be met. There was considerable experience in Boulder in the use of PLA technology, and this seemed to offer the rapid design cycle needed, but would not serve for all the functional units.

A task force was convened under the chairmanship of Uma Reddi, and met on 14 March 1978 to decide on what combination of technology to use. Technologies considered, recommended, and finally used were:

Considered	Recommended	Used
STREAKER	STREAKER	STREAKER
RMS	RMS	RMS
DUTCHESS	DUTCHESS	DUTCHESS
VTL	VTL	VTL
PSALM/PRO		PSALM
GOLF		GOLF
PLAs		
DUTCHESS-F		
EMERALD		
TANGO		

The logic design was started in March 1978. The project was fortunate to obtain the services of Ernie Gurule, Chris Pisciotta and Eric Halvorsen, bringing with them design skill in Golf and Dutchess. The move to Tucson occurred in July 1978, while the logic design was under way. Disruption was inevitable, and some allowance had been made in the schedule, but in the event the problems encountered in Tucson were worse than had been anticipated. TSO service was very poor: slow, with long turnaround and low availability.

Tucson Physical Design group also had difficulties. They, too, used TSO, they were understaffed, and the people who had moved to Tucson were generally the less experienced. There was a lot of learning to be undergone on both sides. A program bug in the EDS system caused a delay of several weeks in the layout of the LSI card.

During the design, problems were solved, but there were further delays in procuring the chips. The RMS Algorithm Unit chip yield was zero on the first pass. The Streaker ID checker chip gave zero yield twice, probably because of the wrong Test Data File tape being used in Burlington. One of the Dutchess chips gave zero yield once.

The main problem in debugging was not design errors but card quality. In general the quality of the logic design was excellent, and microcode errors were common but usually easy to fix. But the quality of the first etched cards received from Boulder was bad, with many open circuit lands and short-circuits at the via holes. The resulting faults were obscure, and were of course different on each machine. The debugging team was at first inexperienced, and progress was very slow. The situation was greatly improved when cards produced in Endicott began to arrive, but the entry to AVT was delayed by six weeks.

The initial design was done largely in VTL, and the packaging scheme was devised to allow piecemeal substitution of the LSI parts as they became available. This strategy allowed the team to avoid debugging a machine containing multiple untried LSI chips. It was planned that the AVT should begin with a VTL machine, and that whatever regression testing was necessary would be done to allow the LSI version to replace it later. In fact, because no changes were required to the first pass LSI design, the team were able to enter test with all the LSI in place except the ID checker and the Data Channel Control logic. The first complete LSI machine was ready for test at the time the project was terminated.

The control logic design was managed by Dave Bailey. Responsibilities within the group were allocated as follows:

Functionl Unit	Engineer	Logic family
ID Checker	Dave Cornaby	STREAKER
Algorithm Unit	Dave Cornaby	STREAKER
Format Unit	Ernie Gurule	GOLF/PSALM
Interface Unit	Dave Cornaby	STREAKER/ DUTCHESS
DCCL	Uma Reddi/Gurule	RMS
Address Decode	Dave Cornaby	DUTCHESS
Encode/Decode	Chris Pisciotta	DUTCHESS
CRC Gen/Check	MSS Part	

Microprocessor hardware

Matt Taub said in his report that the M6800 microprocessor was chosen for the EVT design largely as a stopgap until a suitable IBM microprocessor should become available. The intention was to assess the field again before the start of AVT design, in early 1978. By that time, the obvious candidate was Woodstock, a single-chip microprocessor produced by the Functional Products group in Kingston, of very similar performance to the M6800. An edict had been issued by Jack Kuehler (former GPD Vice-President, Development) that without his specific approval, no other microprocessor could be used where Woodstock could do the job. Two independent studies of the ability of Woodstock to do the job were commissioned; one in the machine group, and one by Jack DeLancey and Paul Paulsen, who were at the time working on the Prospector microcode. The two studies agreed that the microprocessors were comparable in performance in the Bluegill application. However, the decision was made to press for retention of the M6800 for the following reasons.

- Although either machine would handle the EVT job, it would be fully stretched to do so, and it was estimated that for the AVT application 26 percent more power was desirable. The version of the M6800 being qualified by IBM was capable of operating at 1.5 MHz, a 50 percent increase in power over the EVT version.
- There would be a two-month schedule slip because of the drastic restructuring and rewriting of the microcode that would be required.
- Extra resource and schedule slip would be required to produce Woodstock tools.
- There would be additional schedule risk introduced by becoming an early user of a new machine.

For these reasons (chiefly the first two), a case was put to Jack Kuehler for retention of the M6800, which he accepted, and AVT design went ahead using it.

Changes to the Devkit prototyping system used in EVT were inevitable, consequent on the change to the 1.5 MHz M6800. It was decided to condense the four kinds of cards in the Devkit set into one 6-high 4-wide card, compatible with the packaging of the rest of the AVT electronics. This was a mistake, as it turned out: the team were unprepared for the difficulty of processing large new cards containing nonstandard components through the EDS system in a new site.

Dave Craft was the architect and Ralph Bellamy designed in detail the MPU card. It was conceived as a selfcontained, general-purpose minicomputer, providing as much external interfacing capability and RAM as could be fitted onto the card: 8 bi-directional ports, 4 output-only ports, 1 input-only port, 20KB of static RAM packaged on the card, with the ability to address another 12KB off the card. There was also 512 bytes of ROS for IPL purposes, and the Devkit-style 'A', 'B', and 'C' standard serial interfaces. Consequently, the card was exceptionally tightly packed. The card was designed in this general way so that its design could be started before the functional units and microcode were defined. It was planned that, at SVT design time, the excess function (such as debugging facilities and unused ports) would be pruned away.

The logic design (BDLS) was delivered to the EDS group on 23 July, with a forecast delivery date for the card deck, from which the prototype card could be wired, of 1 August, and the etched card to follow as soon as the team could specify the changes needed following debugging. In fact, the wiring list was not produced until 1 November, and the first etched card in early February. The reasons were:

- The card was a particularly difficult one to wire, and demanded the use of a "Dynamic PNEUC" for customized voltage-distribution plane, which had never been tried in Boulder before.
- It contained several components that had not been used before, necessitating writing new EDS rules. Following the move to Tucson, there was only one person left in the EDS center of competence who had the skills to do this, and he was overloaded.
- In the early days at Tucson, TSO reliability was extremely poor, which slowed down the efforts to exercise unfamiliar parts of EDS.
- The EDS center of competence was short-handed on arrival in Tucson, and those who were there were not experienced.

The effects of this delay were serious. It had been intended to build one yellow-wired card, because it would be easier to make changes to it to correct for design errors. The following cards were to be etched, however, because the size of the card and the volume of wiring made it unwieldy and mechanically vulnerable and unstable. The effects of any unreliability in these cards on the progress of the machine debugging was expected to be severe. However, by November, the delay was critical to the schedule, the servo-writer needed a card to continue its bring-up, and it was eventually necessary to build four such cards before the etched version was ready. They were indeed unreliable, and did cause additional delays to the machine debugging. When the first etched cards finally arrived though, they were actually worse, suffering from open nets due to poor etching, and shorted nets caused by damage in drilling. Some cards had as many as fifteen faults. The only way to test the cards was to assemble them and try to run diagnostics, because Boulder (where the cards were made) had no test facility. The second batch of cards was made in Endicott, which had such a facility, and were much better.

It is hard to estimate the effect on the schedule of the difficulties with the MPU card, because they were so widely diffused. Effort was absorbed by debugging that should have been devoted to honing the design, and the schedule damage caused by this did not appear till later. The writer's personal opinion is that it cost about three months.

Once debugged and in good supply, the MPU card concept was a success. The surplus capability and flexibility has been useful in performing auxiliary development functions, such as loading the Psalm arrays, and running keys and displays. It was also useful in the servo-writer.

Data channel

At the beginning of the AVT development in Boulder, the decision was taken by the channel group to change the coding scheme from MFM to a '0 – 3' run-length-limited code. The reasons were primarily development and manufacturing convenience. It was thought that in performance there was not much to choose between the two codes but the 0 – 3 code was being used by the Prospector and Ocotilla projects, and it was hoped, by using common parts, to economize on development resource.

The change of code necessitated some change in the file control logic. The write clock now had to run at 9/8 of the bit rate, and serial data was no longer available to the ID Checker; but these complications did not appear to be sufficient reason for the machine group to object to the code change.

The performance of the channel gradually improved during the AVT development. At the start of 1978, with an average head, the channel would run error free when 8.5 percent off track. In August 1978, the AVT level heads were delivered and improved this figure to 11.5 percent. By test entry in June 1979, improvement in equalizer design and write current had raised it to 15.5 percent, and today off-track performance is typically 17.5 percent.

Motor drive

The term "brushless dc motor" is often used to describe the motor used in the EVT version of Bluegill. This is confusing, because that term usually refers to a combination of ac synchronous motor with Hall-effect angular transducers driving a static inverter of some kind. In fact, the EVT motor was simply a two-phase ac synchronous motor with permanent-magnet rotor.

A disadvantage of this kind of motor is that to start it requires an ac source whose frequency increases from zero to the desired synchronous frequency at a controlled rate. In the absence of static position information (such as is supplied by the Hall-effect transducers in the "brushless dc motor"), any interference with the motor acceleration (such as might be caused by high static friction) is liable to cause loss of synchronism and failure to start.

To avoid this problem it was decided to use instead a hysteresis synchronous motor for AVT. This kind of motor has synchronous running characteristics similar to the EVT motor, but has the advantage that it is capable of producing torque while running asynchronously. To start it, therefore, it is only necessary to apply the running ac drive. This offered a potential increase in reliability and simplification of microcode, at the cost only of a decrease in running efficiency.

Because it was thought important to have a uniform torque, the driver circuits were made to have low output impedance and, to minimize dissipation in the motor, a bridge configuration was used, driving between plus and minus 12 V instead of the alternative of two plus 12 V drivers and a centre-tapped winding. In retrospect, this was probably a mistake. It turned out that the torque and dissipation were not critical, and conventional open-collector saturation drivers would have been cheaper and more reliable.

A simulation program was written in APL to estimate the precision of speed control to be expected, and this indicated that the same system of phase control used in EVT, implemented in microcode, would suffice for EVT except under worst-case power supply conditions. Some slight modifications of the switching strategy and the addition of a small flywheel to increase the system inertia was an effective palliative. A more complete cure, which would have made the flywheel unnecessary but would have required more extensive change to the microcode, was modelled, but the schedule did not permit its implementation.

Two vendors were found who were able to produce the motor to our specifications at an acceptable cost. There was difficulty in writing a specification for the motor that would reflect our usage but would be comprehensible to the motor manufacturers, who are used only to specifying their products in terms of sinusoidal drive voltages. However, a combination of simulation and testing served to establish a correlation between performance on square-wave drive and sinusoidal drive, and the specification was written in terms of an equivalent sine-wave drive.

Read-write heads

The May 1978 Sprat review gives a concise description of the evolution of the head design up to that time.

The EVT heads were of the "Annulus" design, with stabilizer. They suffered from low and variable flying height, which in some cases led to scratching of the recording surface. By late 1978 the candidates for the AVT head were reduced to the '4-4-4' penta-rail, and the '4-8-4' penta-rail designs. Both gave acceptable flying and wear performance with sufficiently carefully controlled contours, but the 4-8-4 design suffered less from erosion of the edges of the core, because they were protected by an additional layer of ceramic. However, although the problem of edge erosion was crucial, the 4-8-4 head was not judged an acceptable solution, because of the high cost associated with lapping to a controlled throat-height, a head on which the throat was concealed. The solution adopted was to protect the edges of the core by sputtering a layer of aluminium oxide over each side of the core before joining it to the ceramic. This effectively supported the corners of the core and prevented chipping and grain pull-out, which had previously led to width erosion. A detailed report summarizing 40 000 hours of wear testing over 15 months was written by Paul Losee.

The problem of low and variable output remained, however, until after the AVT heads were manufactured. Thermal stresses locked into the core at the time of deposition of the protective alumina layers were suspected, but only while the AVT was in progress was the cure hit upon. It was found that by mounting the cores vertically in the sputtering chamber, and sputtering simultaneously from both sides, the asymmetry introduced by sputtering first one side and then the other was avoided. The resulting heads had dramatically better and more consistent output.

Degradation of head output over time has been reported in AVT, but by 1979 it was not clear whether the cause was the accumulation of contamination or changes in the structure of the core material.

A significant development during 1979 was the "drop-in" core design, which was the idea of Allen Cox, and offered considerable simplification of the head-assembly process.

Recording medium

The "Pegasus" recording medium planned for use in Bluegill cartridges was to be obtained as a by-product of the fabrication of Prospector disks. The 24-inch "Jumbo" rolls were to be slit into two parts – one to be used for Prospector disks, and one for Bluegill. The disks were to be cut from the sheet eventually by a laser, but at first by a steel-rule die. The laser process was used for some of the AVT disks. Pegasus was a chrome-dioxide medium, coated on a 0.0015 inch Mylar base. The ink, produced in a batch-operated sand mill, was applied to both sides of the base, but calendered on only one side. The disks were then burnished using the one-pass abrasive tape method.

Initially it was planned to perform a surface analysis on each disk before assembling it onto a backplate. However, the saving in yield at the servo-writer stage was insufficient to justify the cost of the test, and it was abandoned. Instead, the Surface Analysis Test was performed on the servo-writer itself. Each sector was tested for amplitude drop-outs, and if any worse than 60 percent was found, the sector was demarked. The criterion used on the AVT disks was that not more than two sectors on any one track were to be demarked: any more and the disk was rejected. This criterion led to a yield at the servo-writer of about 70 percent.

After the EVT in March 1978, an engineering test of the medium was done in Boulder. No significant problems were found.

An AVT of the medium was in progress at the end of 1979.

Track-following and access

The mechanical design of the actuator was little changed for AVT: the only changes of significance being the introduction of a cast aluminium arm (to improve resonant characteristics), and redrawing the printed-circuit armature card to decrease maximum resistance and improve manufacturability.

The use of LSI (RMS and vendor EPROM) for the algorithm unit significantly improved the computation time. The complete calculation of the fifth-order algorithm was reduced to 27 microseconds. The A/D and D/A functions on the AVT models were performed by a Dutchess logic module and Zircon D/A module. The 8-bit A/D function (performed by successive approximation) took 16 microseconds. Sharing the same hardware between A/D and D/A functions economized on hardware, and provided some useful tolerance cancellation.

By microcode changes, the average access time was reduced from 100 ms in EVT to 33 ms in the AVT model. This was done by using the total available current drive to the coil for an extended period to reach a maximum velocity of 2 tracks per sector for seek lengths of 26 tracks or greater.

A method of demarking severely defective servo samples and correcting the less seriously erratic ones was incorporated in the AVT design. When the disk is servo-written, directly after each track has been written, it is read back and the servo sample stored. The average of the stored values is subtracted from all of them, to remove the effects of servo-writer head width tolerance, and then the modified values (if they are less than 32 times the value of the least-significant bit of the digitized PES representation in the file) are written into a 6-bit field in the "flag byte" – part of the ID field of the previous sector. If this correction exceeds the value that can be written in the 6-bit field, a seventh bit, the Servo Demark bit in the flag byte, is set.

When the file is track-following, each correction field is read, and the value stored there is subtracted from the following PES before it is used. If the Servo Demark bit is set, it is taken that the following sample is useless: it is discarded, and the stored value from the previous sector is used instead.

This system, due to Bill Case and Chris Wallis, provided a useful improvement to the defect resistance of the track-following servo.

No change to the servo algorithm coefficients was needed, except for a correction to compensate for the different gain of the actuator driver.

The track-following and access design was the work of Ian Henderson, except for the microcode implementation, which was by Joe Luciani.

Servo-writer

The EVT servo-writer had undergone many changes by the start of EVT. Several different people had made contributions to its design, and each had made changes to what already existed. Moreover, when it was shipped to Boulder it was decided not to ship the granite baseplate, but to procure a new one in Boulder. This had not in fact been done by the start of EVT, so the servo-writer had to be assembled in a makeshift fashion on a surface plate.

The result was that the disks written for EVT were of poor quality, suffering from poor phase accuracy (within each track and from track to track) caused by mechanical vibration and by errors in the design and set-up of the PLOs used to process the clock. The machine was hardware controlled, and the documentation was not in good order following the many hasty changes that had been made to meet EVT problems and schedules. It was foreseen that in making the modifications to the format required for AVT there was a serious danger of loss of engineering control.

For these reasons it was decided to build a new servo-writer for AVT from scratch, rather than try to refurbish the EVT one. The AVT design incorporated the following improvements.

- Great pains were taken to improve the speed constancy and phase accuracy. A low-inertia dc motor is fixed rigidly to the baseplate, and drives the air-bearing spindle through a flexible coupling. Accurate speed control is maintained by a servo from speed information derived from an optical shaft encoder.
- Air-bearing linear slides are used to support the carriage, to reduce static friction and its destabilizing effects on the positioning servo.
- As on the EVT machine, control is by microprocessor, using the file MPU card. On the AVT version, the format, too, is controlled by the microprocessor, affording the ability to alter the format by microcode changes only.
- By taking advantage of the microprocessor, it was possible to incorporate a number of development and debug facilities. For example, it is possible to provide a scope trigger simultaneous with any chosen bit position on the track; making it very easy to check the format or investigate defects on the medium.
- Comprehensive logging facilities are included, using an attached IBM 5100. (It was originally planned to use a Series/1, but after experiencing difficulties in using the Series/1 for microcode loading, it was decided to fall back on the 5100, which was attached with great ease in a few days).

Pat Mulholland was assigned from Hursley to manage the task.
Other people involved with their tasks were:

Bill Case	Overall system and logic design
Ken Hales	Microcode
John Farran	Demodulator design
Russ Chaplin	Analog and check circuitry
Vince Jaques and Walt Poplewko	Mechanical design

Because the servo-writer at that time determined the critical schedule path for the whole project, its own schedule was cut very fine. In February 1978 the decision was taken to build the new machine. Debugging started 21 November, three weeks behind schedule because of delays in production of the MPU card; but the first disk was nevertheless written on schedule on 16 January 1979. Very few format problems were found, and those were quickly corrected. Addition of demark and check facilities continued up to the end of 1979. The design has been found an unqualified success.

Track misregistration (TMR) studies

Up to and including EVT, all the TMR work that had been done was based on a Monte Carlo simulation program written in APL by Matt Taub. In Boulder, the TMR work was taken up by Rahmat Aziz, who at first used the APL simulation also. However, the APL simulation had drawbacks. It was expensive to run, for the sake of simplicity it used worst-case assumptions that would have been costly to the design, and the number of variables it could handle was limited. The FORTRAN IV program "OPTIMAC" had been written in San Jose for use in diskfile TMR modelling, and was a good deal more capable. It was used for all the succeeding AVT TMR work.

In the middle of 1978, concern began to be felt over the possibility of head-width erosion in the field because of grains of the ferrite crystal structure pulling out of the edge of the head at the gap. This process had been seen to occur in head wear testing, and would have a serious effect on field performance. The danger was that the machine concerned would go on reading and writing apparently satisfactorily, but the data it wrote might not be readable by any other machine.

The OPTIMAC studies showed that only a low probability (2.5 percent) of occurrence of grain pull-out could be tolerated if the data reliability specification of 1 error in 10^6 bytes was to be met.

Apart from this problem, the error rates were very satisfactory, even with the off-track performance current at that time (about 10 percent).

The model was also used to identify those tolerances in the machine design that most critically affected the unwanted component of read signal. They were found to be the tolerance on the eccentricity of the disk as mounted on the hub, and the tolerance to dc offset on the servo. The first of these had to be carefully controlled in the mechanical design, and was the reason for an engineering change to the cartridge at a late stage to improve centering. The second needed to be very tightly controlled in the servo design at that time.

When grain pull-out was soon cured, the TMR situation would have been secure were it not for the question of minimum tolerable head output. In April 1979 there was anxiety over the low and variable head output, and its sensitivity to temperature. This problem, too, appears at that date to have been put right, but in June 1979 it was thought worthwhile to try to model the head behaviour and incorporate it in the TMR model.

Throughout the AVT development, the TMR studies showed satisfactory margins to failure, in the absence of degradation of head performance with life. In test, the source of errors had always been drop-outs rather than marginal track-following performance, suggesting that, if the head flying and medium quality would support it, track density could have been increased.

Testers

Very little in the EVT testers had to be changed for the AVT development. Only trifling changes were needed to accommodate changes in the command set and interface specification.

The development testers were unsuitable for use in the AVT itself, however, and it was decided to build an adapter to enable Bluegill to be attached to a universal test system (UTS) for the test.

John Holloways was assigned to Product Assurance to design and build the adapter, and assist in debugging of the test system. The adapter was ready in January 1979, and the test system was largely debugged using an EVT level file, by 1 May, when the AVT level files were supplied for attachment.

Cost estimates

Product cost

It could be seen that the total cost inclusive of contingency fell by 10 percent from Stage I to Stage II, largely because of savings in the mechanical part of the machine, and in spite of offsetting increases in the electronics cost.

Following Stage II there were some upward revisions of the Tucson Labour and Burden rates that would have restored the cost to its Stage I level. These rates were under review at the time the project was terminated.

Cartridge cost

It has always been recognized that the cartridge cost was crucial to the success of the program, and continuous effort was applied to reducing it. Between November 1978 and 22 February 1979, the Tucson Labour and Burden rate increased by 84 percent. At constant Labour and Burden rate, the cartridge cost was estimated (by the Media group) to have decreased from an initial estimate (Hursley, March 1977) of $16.09 to a final value at the time the program was terminated, October 1979, of $14.93. At the end of 1979 however, the cost was $19.27.

Maintenance cost

Overall there was a 47 percent decrease in CE cost per unit between Stage I and Stage II.

9.0 Swallow (62SW)

This section was written by Bill Case.

First customer shipment (FCS)	June 1984
Capacity (MB)	130
Average access time (mS)	27
Tracks per inch (TPI)	850
Bits per inch (BPI)	8770
Data rate MB per second	1.0
Shipments to February 1990	11 000

Track-following

The track-following was with the same loop as Piccolo but the production of the Hybrid PES was as shown below.

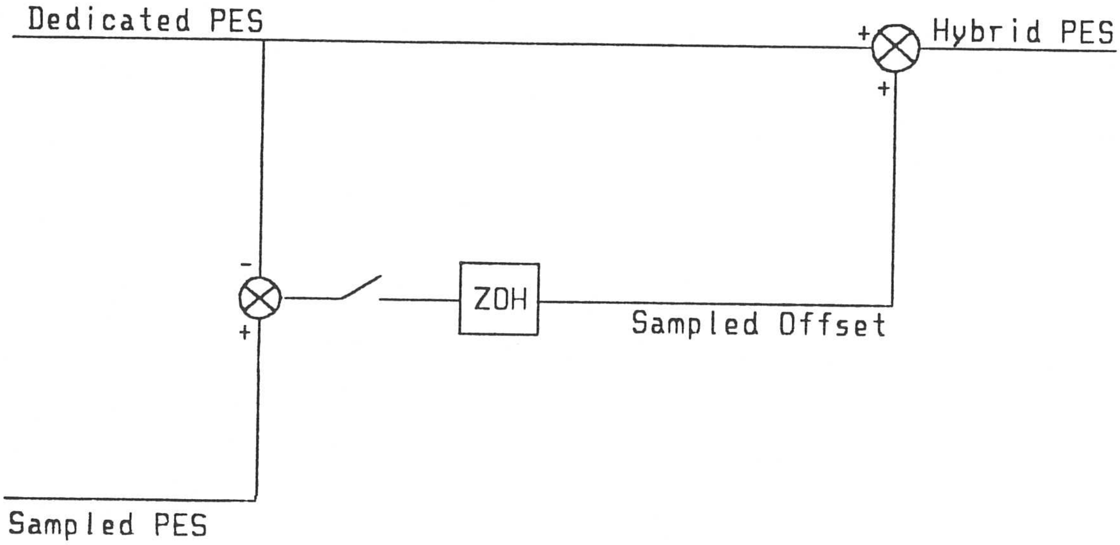

Swallow Hybrid PES

Accessing

This was the same as Piccolo.

Data Channel

This was the same as Piccolo.

SAT

This was the same as Piccolo.

This was a file that was to only use the basis of the Piccolo HDA, but doubled its capacity. This was achieved in a variety of ways.

The head read-write width was halved to allow double the track density. This was not simple, as for Piccolo the side faces of the element were at 45 degrees, which allowed lots of side reading and writing and meant a lot of interference with the adjacent tracks. To reduce this problem the sides of the read-write element were converted to 70 degrees.

A new method was invented to combine the position-error signal (PES) from the dedicated servo head and the sampled PES from the required data head. This is shown in the diagram to create a new hybrid PES. All the logic of Piccolo was put into PLAs using the Streaker module family. This family was one of the first LSSD chip families. There were several passes through the system before error-free logic was created.

The price fixed for Swallow was the same dollars per megabyte as Piccolo. Because of this, very few were sold and all in the development team were very disappointed.

10.0 Woodwind (3310)

This section was written by Bill Case. There are four models of 3310 available as shown below:

Model A1	One Piccolo drive with 64.5MB capacity
Model A2	Two Piccolo drives with total capacity 139MB
Model A2 with Model B1	Three Piccolo drives with total capacity 193.5MB
Model A2 with Model B2	Four Piccolo drives with total capacity 258MB

The 3310 device is connected to an IBM 4331 Processor. Each disk enclosure appears to the 4331 as 126 015 blocks of 512 bytes each.

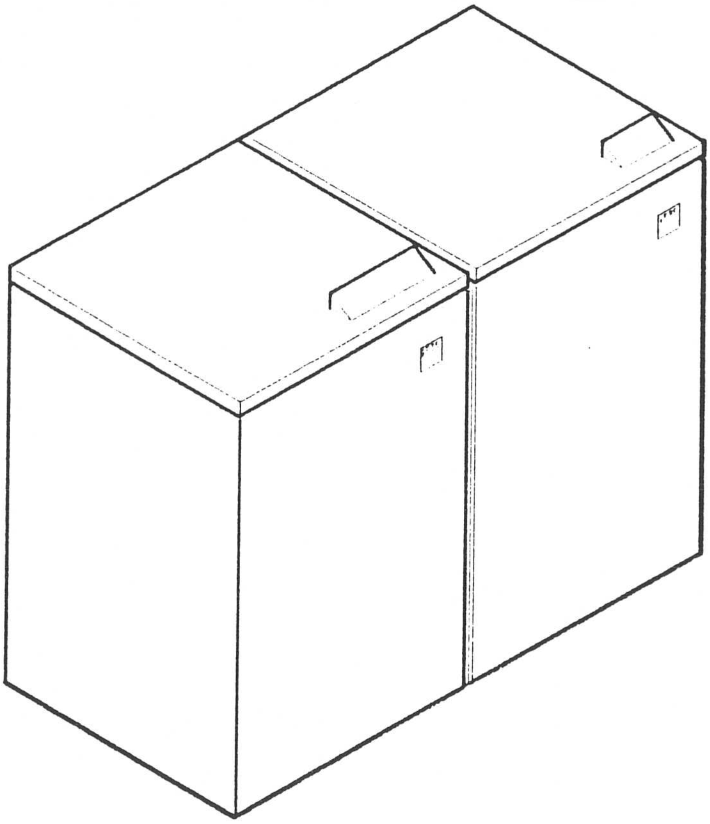

Figure 14. The Woodwind (3310) cabinet

11.0 Flotilla

This section was written by Bill Case. The Sprat-Bluegill project was resurrected after work had finished on the Swallow device. It was managed by Bob Avgherinos with as many of the people who had experience of Sprat and Bluegill as was possible.

Disk	208 mm flexible
Formatted capacity	20MB
Recording density	27.8KB per inch
Track density	333 tracks per inch (TPI)
Number of data tracks	573
Head	Ferrite in ceramic slider
Access performance (ms)	
Average	35
Maximum	65
Single track	10
Latency	20
Actuator	Swing arm
Servo	Samples only
RPM	1500
Data rate	1.19MB per second

The development was stopped after about 18 months work as it was not considered to be a product with sufficient guarantees of sales numbers and profits for the corporation or to be a valid technology concept.

A very clever scheme was invented, by Chris Wallis, to make the actuator follow the ovality of the Crolyn disks with no error. It was required to learn the ovality during an initial revolution and then update that value on each revolution thereafter. This value was then used to make a "feed-forward" current which did not need input from the track-following servo.

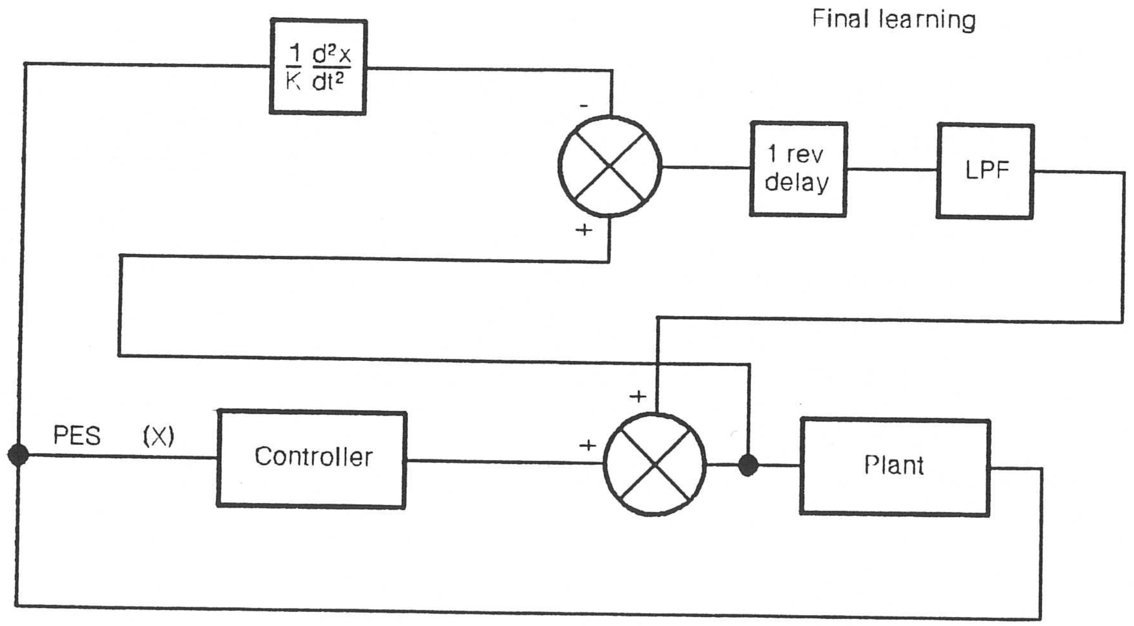

Figure 15. Method for following ovality

12.0 Harp (later Heron then Falcon)

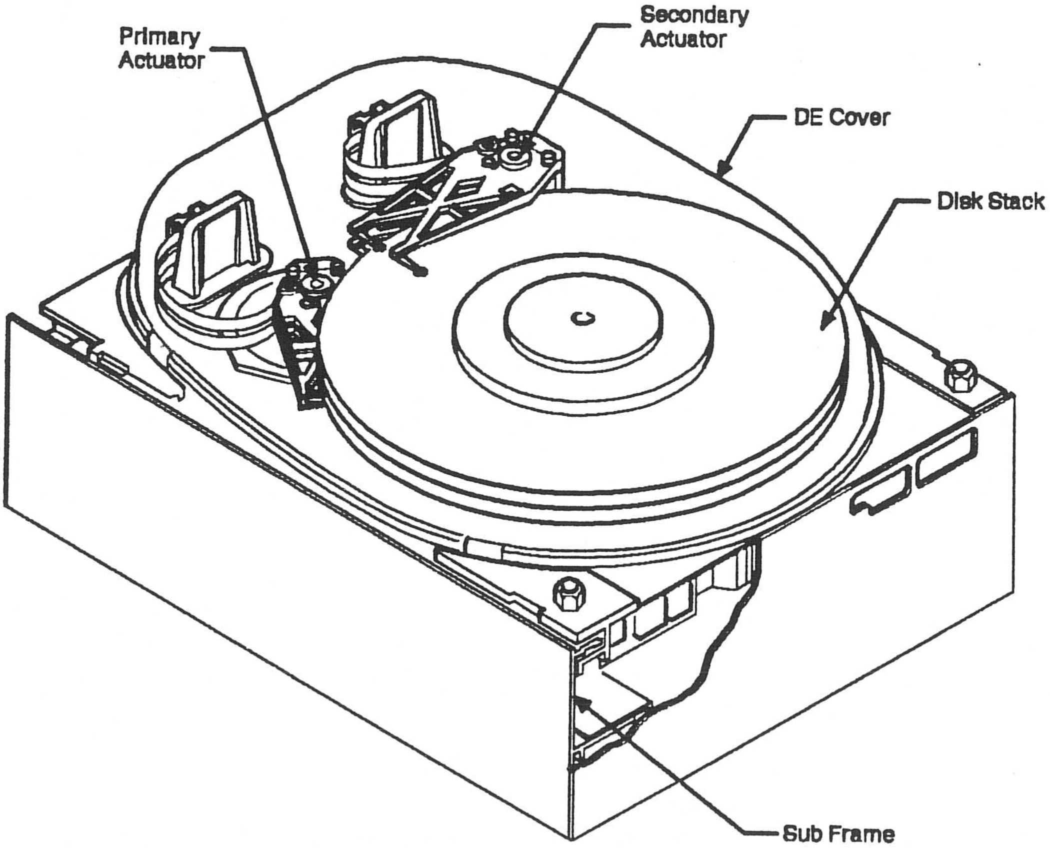

Figure 16. Inside view of HDA

Introduction to the Falcon servo

This subsection was written by Pete Baker.

The selected head position is derived from information written on the data tracks once per sector. The servo is therefore sampled, and the sample rate is 4350 Hz. The servo error signal is held digitally, and many of the access controls are digital; however, the compensation for both position loop and access loop is analog. The compensator is a dual lead – lag with a parallel integrator that acts as a low-frequency integrator. A filter is used to stabilize mechanical resonances.

Accessing

Because the Falcon file has position information at sector time only, a normal velocity servo loop for controlling accesses could not be applied. This is because up to 13 tracks are crossed between samples and the resulting position information would be too coarse to control the loop, also the resulting quantization would present differentiation difficulties in generating velocity.

The problem was solved by designing an electronic model of the file actuator, that given actuator motor current, could estimate the resulting head position. See Figure 17 on page 98.

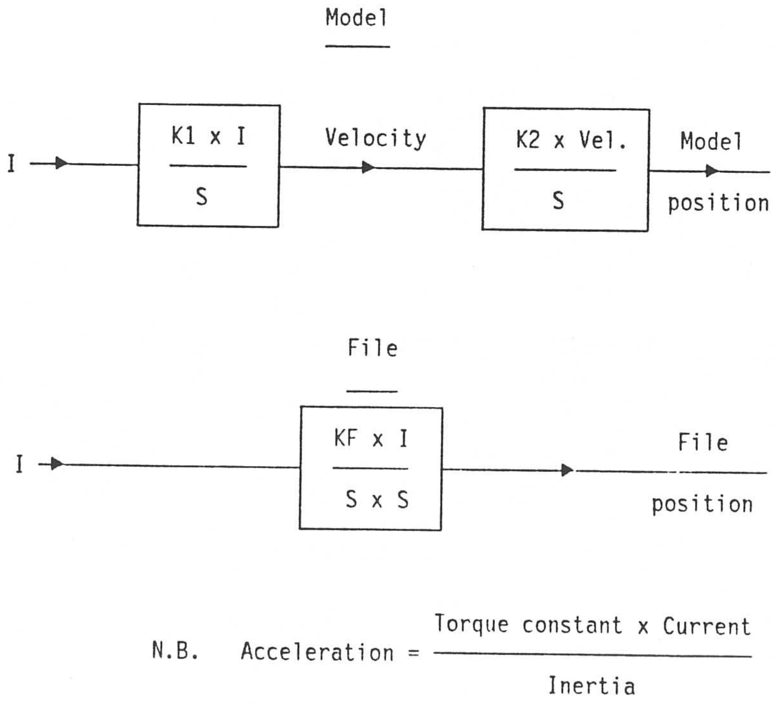

Figure 17. Simplified model

If the model is a perfect representation of the file actuator and equal currents are applied to each, then model position tracks file position. The continuous position and velocity information, unobtainable from the file, can now be taken from the model and used to control file accessing.

It was found that a simple double integration of acceleration was sufficient to represent the model, provided the constants K1 and K2 were chosen such that K1 x K2 = K, the file acceleration constant. The first integration is achieved by analogue means, thereafter the velocity signal is digitised and digitally integrated.

Any errors resulting from the mismatch between the model and file can be corrected by closing a servo loop around the model. SeeFigure 18 on page 99.

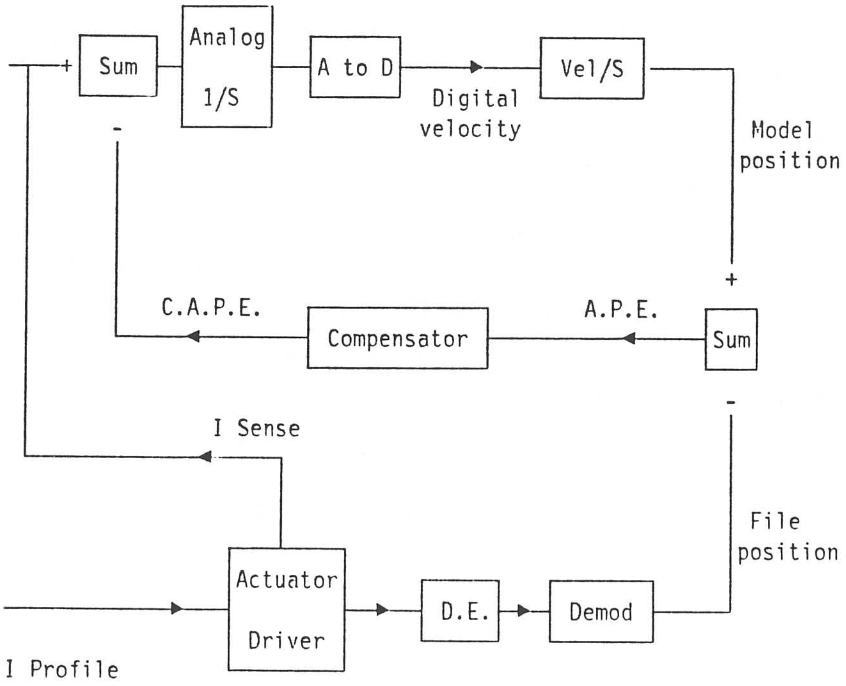

Figure 18. Model loop

To close the loop around the model it is necessary to generate an error signal. This is done by sampling both the model and the file at the same time. The file position information is then digitized, and the difference between the model and the file is computed; this difference is known as APE (access position error). Because the loop forms a second-order system, a compensator is required to ensure stability. APE is passed through the compensator to produce CAPE (compensated access position error signal). This is used to modify Current Sense, such as to drive the model towards the file. The model is servoed to the file.

By using the CAPE to modify file, current the file could be made to follow the model (file servoed to model). Both methods are used during accessing.

During the accelerate part of the access, current from a current-profile generator is fed into the file and model. The model is then servoed to the file as in Figure 18.

During the decelerate part of the access, the current-profile generator produces current of the opposite polarity which is fed to the file only. As model velocity is the integral of the current, and because the integration constants are known, it was not necessary to do this computation during the access. The velocity is precomputed and stored in ROM to be addressed by model distance from the target position. This velocity is fed into the second model integrator to produce model position. Thus the model is forced to follow a well defined velocity profile with respect to the distance from the target track, which is zero at the target position. See Figure 19 on page 100.

Since the model is well controlled, the file is servoed to the model to bring it to zero velocity at zero position. The current profile is set to zero when the model position is zero.

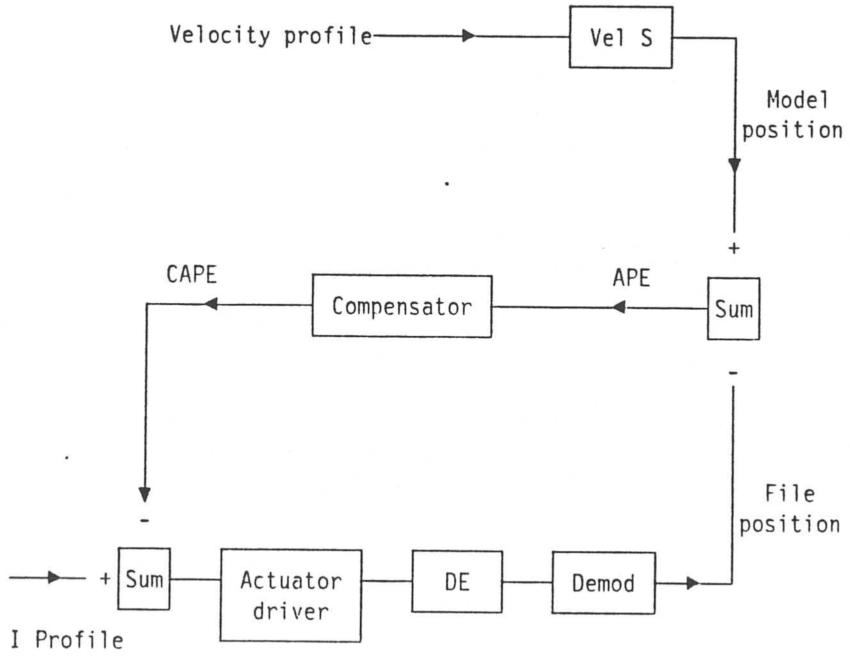

Figure 19. Decelerate model loop

Access types.

The access lengths are divided into five groups according to length. This takes into account the current slew rate of the actuator motor coil and the need to allow the servo sufficient time to control the file during each phase of the access.

All seeks of less than one track, and all one-track moves without head change, are done as a position-loop settle. The model is loaded with the position of the new track centre, and the resulting error signal, APE is used to drive the head to the track centre. The remaining four groups are done under the access-loop control.

One remaining type of access exists. This is the access that is only used during power-on or a re-sync operation to get from the actuator stops to the track before the home track. It is like a normal access without an acceleration phase. At the start of the move, the model differentiator, creating velocity from position, is fed with the least significant bit of velocity. This causes the model to move slowly, towing the file along with it until the track-minus-one centre is reached. Here the velocity bit is switched off and the model and file stop.

Track-follow

In track-follow mode the file is configured as in Figure 20.

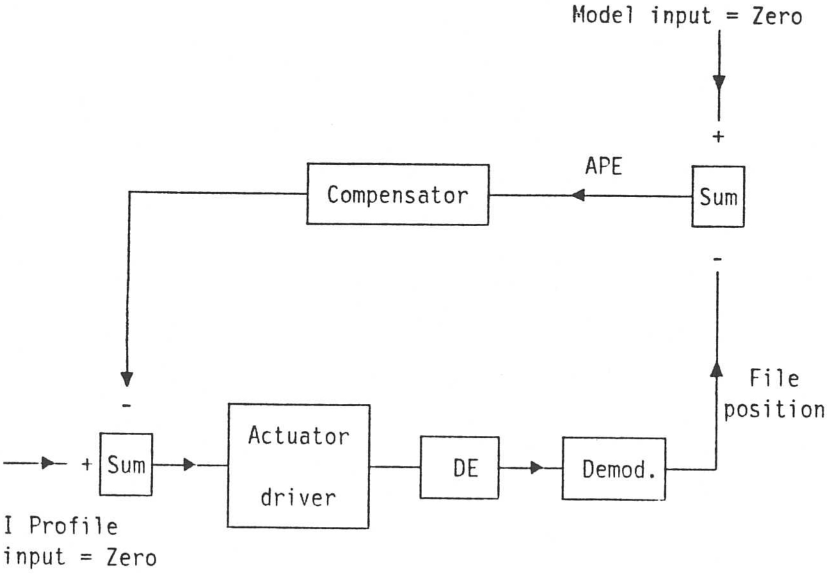

Figure 20. Position-loop configuration

Compensator

The dual lead-lag consists of two zeros and two poles: see Figure 21 on page 102. During track-follow the zeros are Z1 = 135 Hz, Z2 = 200 Hz, and the poles are P1 = 350 Hz, P2 = 2000 Hz. The dc gain of the network is 0.1. This gives an open-loop crossover of 330 Hz and phase margin of 40 degrees.

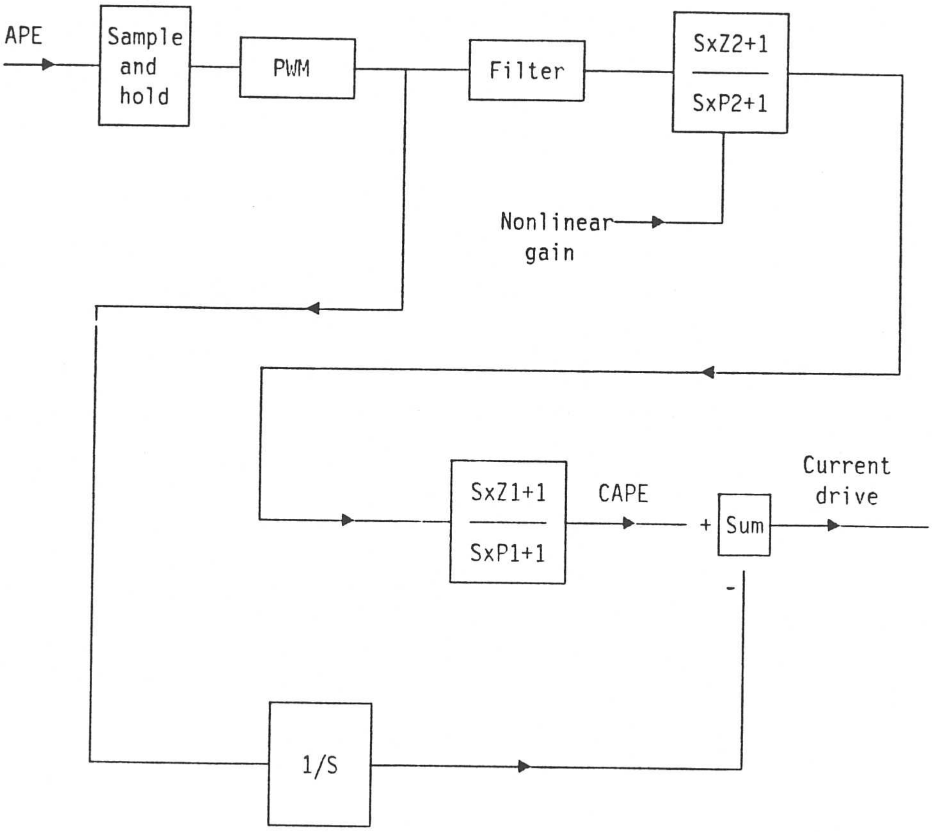

Figure 21. Block diagram of compensator

To speed up settles and reduce access errors, a switch is closed during seeking and settling to boost the gain of the stage and modify the values of Z2 and P2 with the aid of two diodes. The effect of the diodes is to allow signal amplitude to control Z2, P2, and the compensator gain. For 1.5 track error, the nominal values are Z3 = 380 Hz, P3 = 2810 Hz, and gain = 0.174. For a 15% track error Z3 = 241 Hz, P3 = 2042 Hz, and the gain is 0.118.

Low frequency integrator

The integrator removes any position errors caused by dc forces on the actuator. It integrates only during track-follow. During access, its last value is held and updated with a voltage proportional to predicted windforce variations across the disk. This force is derived from analog velocity. The integrator is never reset to zero unless a servo error occurs.

During an access, the wind and spring forces on the actuator change with position. The relationship between force and position is approximately linear. By allowing the integrator to integrate a signal derived from velocity during an access the integrator can be preset to an estimate of the force it has to compensate for at the new track. The signal used, "Feed Forward Wind", is rectified analog velocity.

Filter

The filter desensitizes the servo to resonances above 1200 Hz. It provides 10 dB of attenuation at the actuator main torsional resonant frequency of 2 kHz, and 18 dB at the first arm in plane frequency of 2800 Hz. It also adjusts the PWM switching frequency to reduce the effect of the switching frequency by at least 16dB, the fundamental frequency of which is 32 kHz.

The filter frequency response is shown in Figure 22.

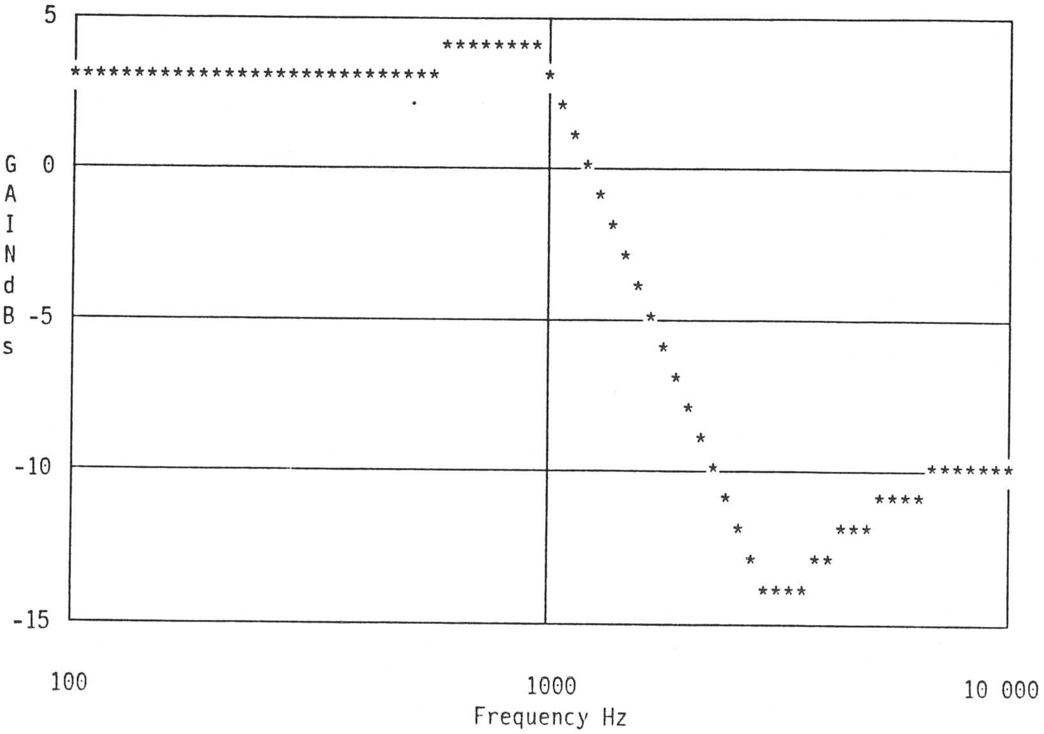

Figure 22. Filter frequency response

The power-on cycle and calibration

Before power is applied to the file, or after a servo error, the actuator is retracted to the outer stop by a spring. After the motor has come up to speed, or at the start of the re-synchronization operation, the demodulator is synchronized to the chosen head. Position samples are produced and the servo moves the head to the track minus one where demodulator calibration data is read.

The average torque-inertia ratio is measured by performing a series of calibrating accesses. During the accelerate phase of these moves, the gain of the model is adjusted to minimize the positional error between model and file. The gain of the model is controlled by pulse width modulation (PWM) and the amount of adjustment required is a measure of the actuator motor torque-inertia ratio that is used to match the servo and model gain to the actuator.

Servo control-loop block diagram

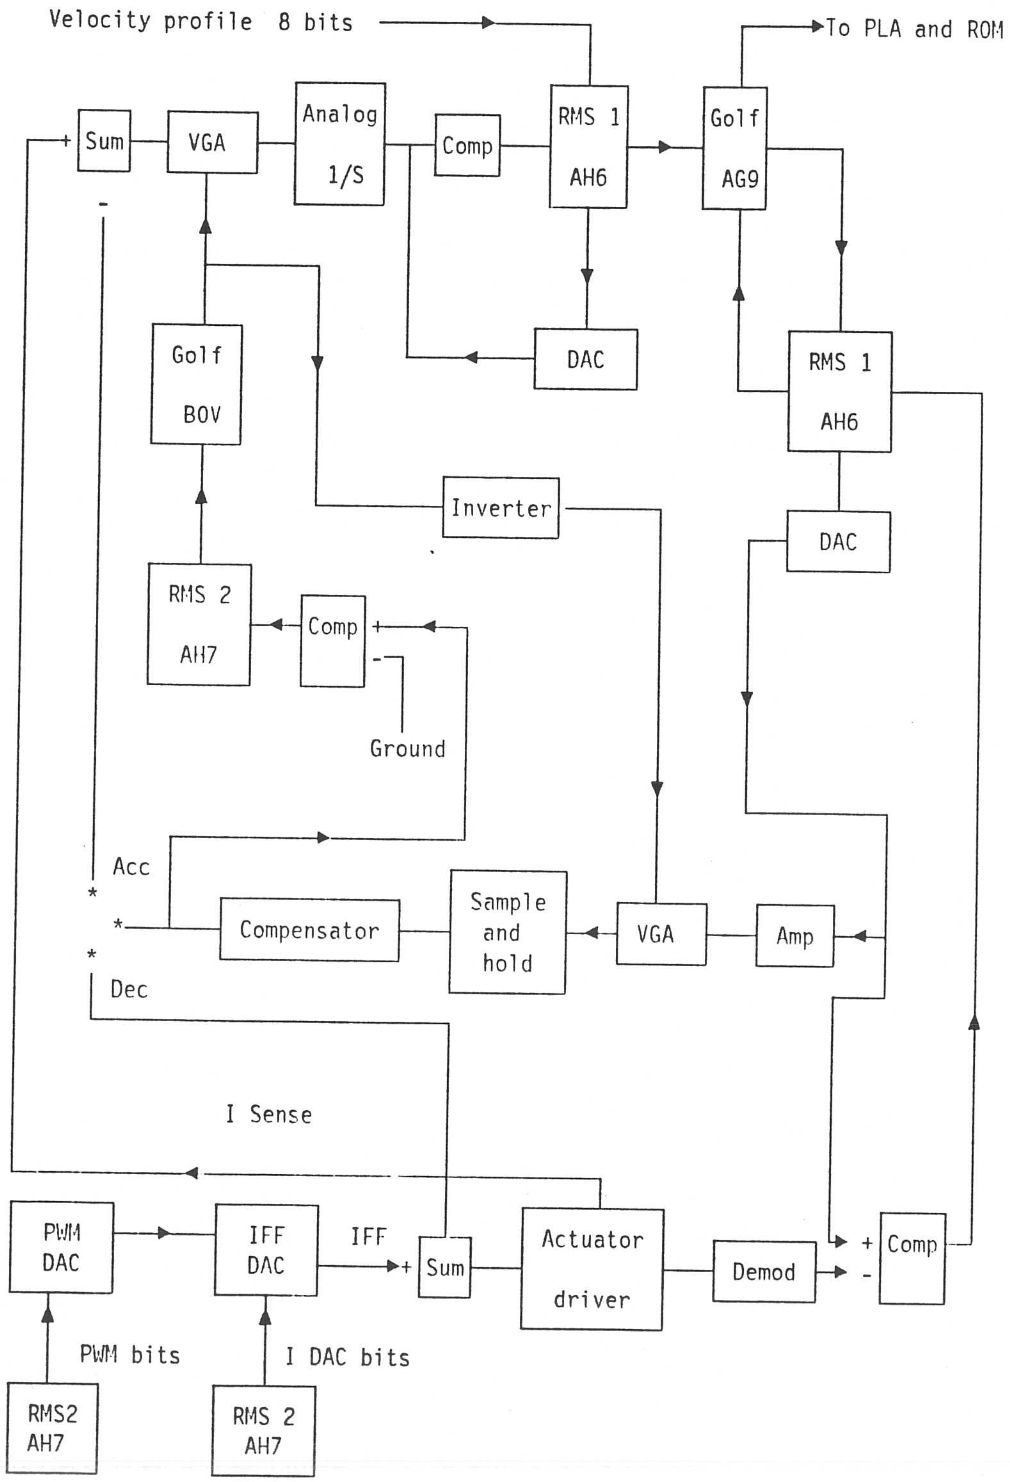

Figure 23. Functional servo drawing

13.0 Falcon GT

This section was written by Bill Case.

When Flotilla was cancelled, Storage was still in the GPD Division. John Newman headed a small group of engineers including Chris Wallis, Bill Case, and Andrew Cameron. The initial idea was to make a device capable of using a count key data (CKD) interface that could be be used in most of the systems that had previously used San Jose files.

The servo system was planned to use the system designed in San Jose by Mike Workman using a specially designed DSP (Digital signal processor) called DCP (Digital control processor). This device was somewhat expensive so the team was going to use an a comercially available device which was much less expensive.

It was believed that the Falcon HDA would be adequate for the greater accuracy required for CKD. For a CKD machine, very close control of the spindle speed was required. This was achieved by Chris Wallis with a very minor modification to the Falcon spindle drive.

Falcon GT (Falcon Prime) was to have been a version of Falcon, providing:

- Faster access
- Lower cost
- Fast head-switching, making it suitable for CKD emulation.

Its primary market was to have been at the low end of System/370. There were several significant technical innovations:

- A digital servo was adopted, using a signal-processing chip set developed by Mike Workman's servo-architecture group in San Jose.
- Because the servo was digital, it could be multiplexed to control both actuators, providing the significant cost reduction.
- The increased bandwidth of the digital servo permitted an increase in track density from 1600 tpi to 2000 tpi.

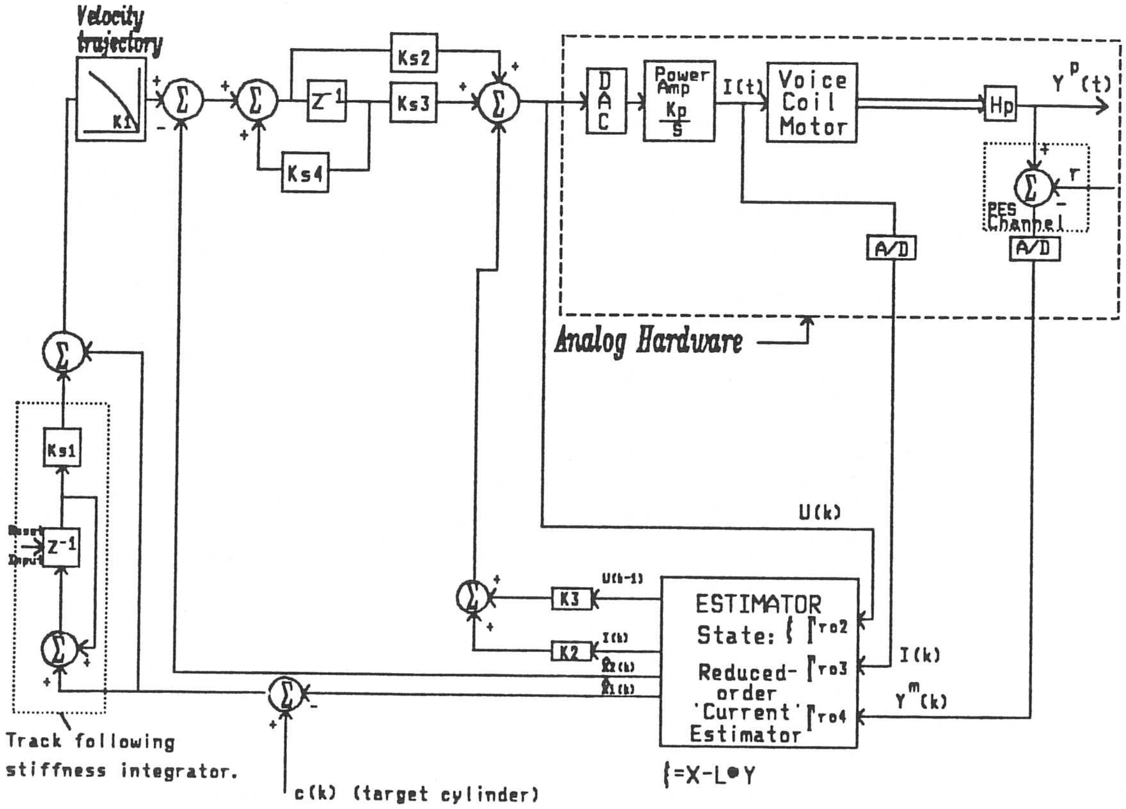

Figure 24. The design of the servocontroller by M. Workman

- In the San Jose servo architecture, the microprocessor controlling the servo contained a mathematical model of the actuator (as was conventional by 1990). An additional innovation (only in 1990 becoming conventional) was that the temperature of the voice-coil was to be modelled, allowing the loading on the actuator to be limited automatically to a level that avoided damage.
  The electrical design of the voice coil could therefore be less conservative, and the average seek time improved from 24 ms (then the Falcon target seek time), to 14 ms.

- To allow the servo to be multiplexed, the position-error samples from the two actuators had to be synchronized. For reasons connected with the geometry of the actuator, servo samples from the inner of the two heads on each surface had a phase relationship that varied with arm position, making this synchrony impossible to achieve.
  To overcome this problem, a scheme was devised to operate the servo always from the outer head. Thermal motion of the inner head relative to the outer was corrected by a "Wedge" servo: a low-bandwidth servo loop controlled by an occasional (once per revolution) error sample from the inner heads. To allow rapid settling onto an inner head following a head-switch, the thermal correction for each head pair was stored in memory, and applied before it would have been available from the single sample per track.

  Thus the major delay in head switching in Falcon, the time to resynchronize the servo PLO when switching between inner and outer heads was eliminated.

An extension of this scheme (never implemented, because the market requirements group in San Jose thought it unnecessary), was to arrange for the head-to-head correction to be zero at operating temperature, allowing a virtually instantaneous head switch. The only penalty would have been occasional missed revolutions at head switch under CKD usage during warm-up. Late in the programme, the planners changed their stance on this point, and the inability to perform instantaneous head switch became one of the reasons for the abandonment of the project.

- Another innovation that has since become conventional was the accurate control of disk speed, relative to a crystal-controlled oscillator. Because of this tight speed control, the same crystal could be used to generate write-clock. Previously, write clock had been generated from a PLO driven from the sector pulses. Phase jitter on this PLO was a major contributor to the error rate, and was eliminated by use of a crystal oscillator.

Apart from technical innovations, the project was remarkable for the organizational model it adopted. It was the first time in Storage Products that the role of Project Architect was defined. This was a technical professional without management responsibility, reporting to the Project Manager, whose role was to do the broad-brush overall design of the electronics, and specify the consequent requirements to the managers of the various functional groups. In this way, he was in a position to make intelligent trade-offs between the difficulties of the component tasks. He also reviewed proposed design changes, and followed their effects through to other affected groups.

This project organization appeared to be successful in producing a well-integrated design.

A special invention by Chris Wallis meant that the very large ovality seen with the flexible medium could be tolerated. The idea is shown below.

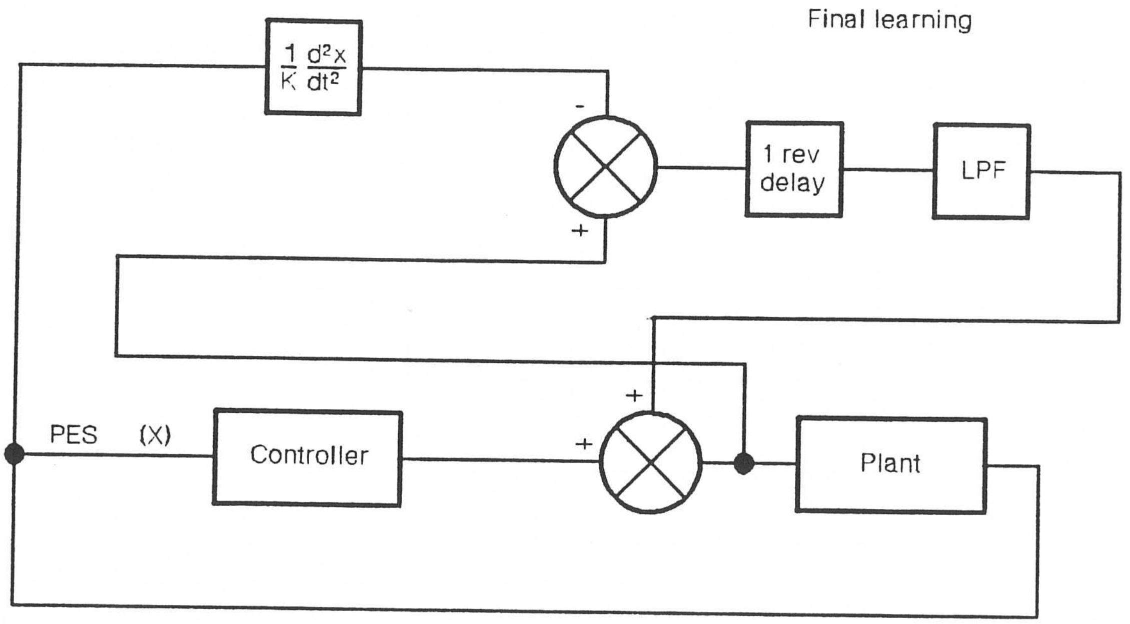

Figure 25. The scheme used to overcome ovality

In April 1984, resource difficulties in the attachment group in San Jose caused them to abandon their work on Raven, the attachment for Falcon. It was decided to stop work on Falcon GT, so that the Hursley resource could be applied to pick up the work on Raven. This was successfully done, and the Raven group thus created was the nucleus of the Storage Subsystems group that exists today (1990).

14.0 Kestrel (9335)

This section was written by Bill Case.

First customer shipment (FCS)	August 1986
Capacity	850MB to 3.4GB
Average access time	18 ms
Tracks per inch (TPI)	1616
Bits per inch (BPI)	16 000
Data rate	3.0 MB per second

The Kestrel tower has two types of device in it:

• A Disk Controller	A01
• A Disk Storage device	B01

Configuration

Every configuration has to include one A01 controller, which can have between one and up to four storage devices attached to it. The interface of the controller to a system is IBM IPI-3, which means that up to eight controllers may be attached on a single IPI-3 line. A configuration is shown in the following diagram.

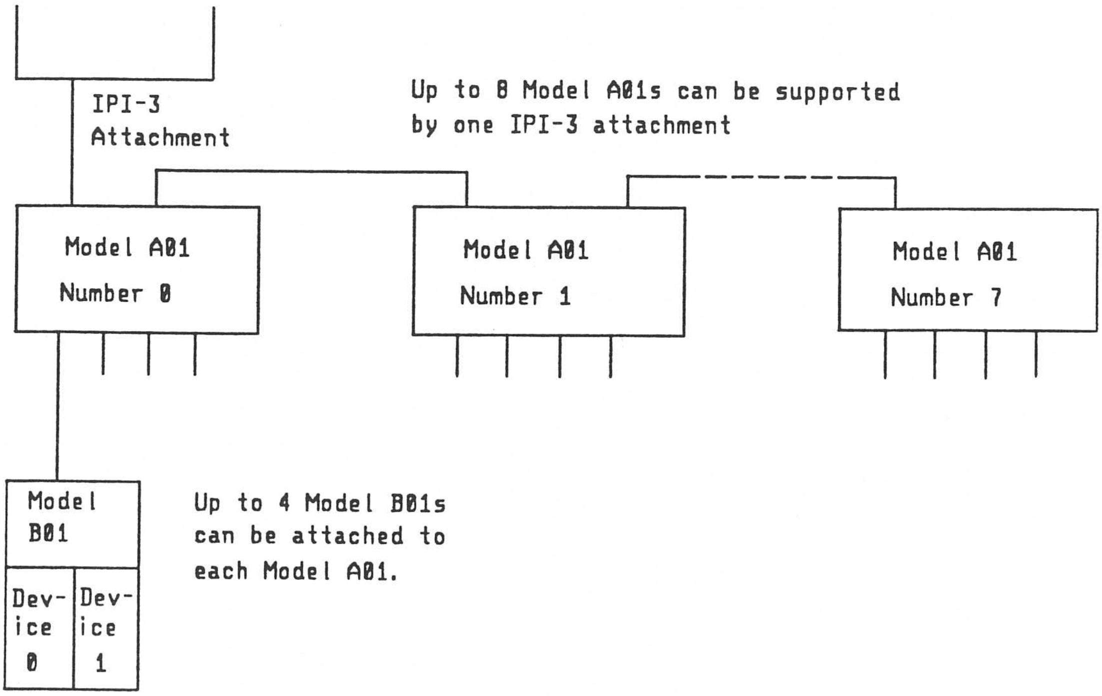

Figure 26. Configuration for one IBM IPI-3 attachment

The heads in the DASD drawer, B01, are arranged as below.

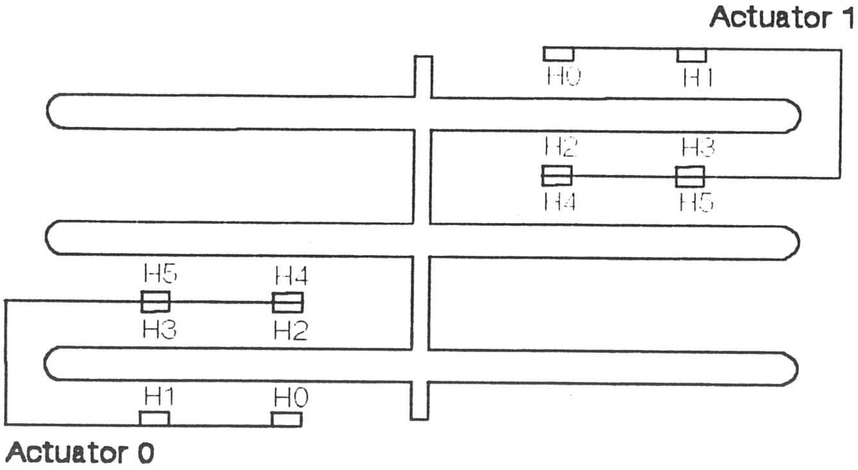

Figure 27. Cross-sectional view of the recording surfaces and read-write heads

As shown previously, there are two independent actuators each with six heads. There are two bands of 1965 data tracks, with an area outside each band where the heads land when the disks are not rotating. A diagram of this follows.

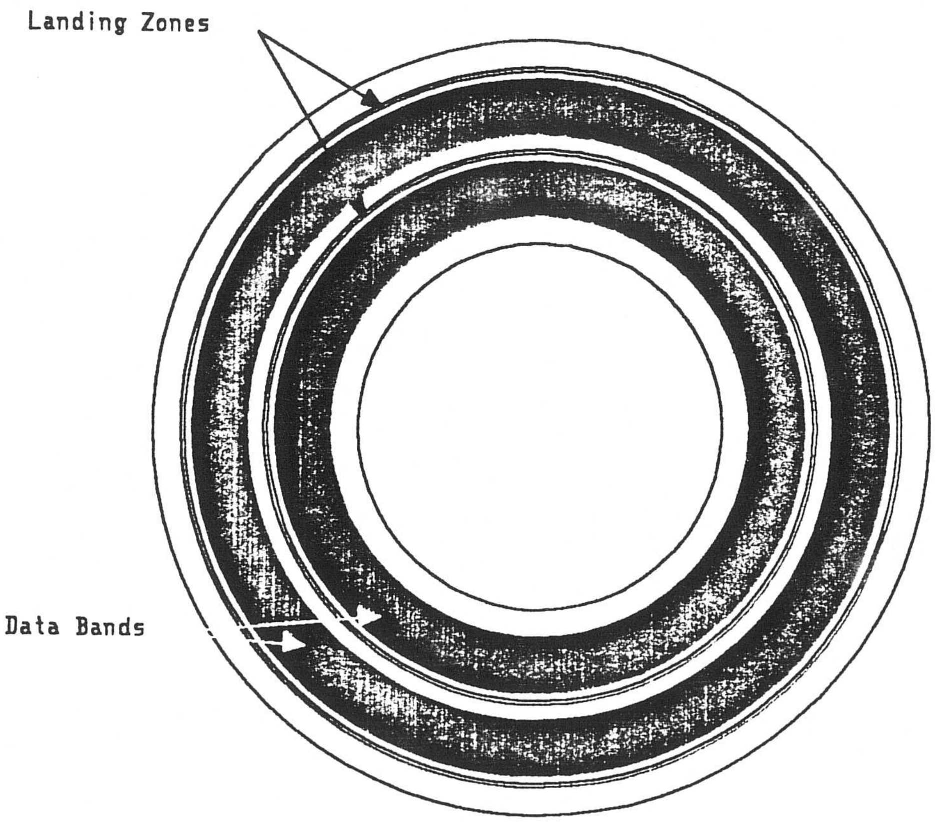

Figure 28. Plan view of recording surface

15.0 Simulator and Matchbox (1984)

This section was written by Bill Case.

When Storage was first involved with servo systems around actuators, it was thought that it would be much quicker to develop control systems if a simulation of the mechanics was available.

A design of such a simulator was got from San Jose, where it had been designed in Hal Stevens' group. The group had no formal (IS) documentation, so Bill Case brought back the hand-drawn documents and used them with Nigel Ansel in Dave Brown's group in Physical Design. They worked on the hand-drawn documents and produced EDS-compatible drawings of the circuits and of the card assemblies. This work was essential for all the simulators to be made with etched cards.

(It was at this time that Bill Case worked with Ron Wilks to write a package to enable board wiring to be done from ALDs produced from EDS. The basic EDS did not allow for board wiring without full details of the cards to be mounted on the boards. To overcome this, card macro diagrams were used to enable full board ALDs to be produced. The file produced of the ALD was then input into the APL package which would then produce a shortest route correctly with no "daisy chains", that is wires which go from a lower wrap to a top wrap. The output from APL was used by any group, inside or outside the Lab, to wire the boards.)

The functions in the Simulator and Matchbox included the following:

- Electrical noise generator
- Index and sector generation
- PES generator
- Position ramp generator
- Cylinder counter and display
- Simulator of the power driver
- Electrical simulation of the VCM
- GPIB bus decoder
- Mechanical parameter selection
- Display of these parameters
- Count and display of the Behind Home position
- Interface to the HP 5451 Fourier analyzer (FA)
- A/D converters to interface with the FA
- D/A converters to interface between the Matchbox and Simulator.

There were two cabinets: one contained an electrical simulator of an HDA, and the other contained an interface from this simulator to a Fourier analyzer (HP 5451). LEDs on the front of the simulator displayed all the parameters that could be set up from switches or by direction of the FA. The PES generation had to be changed for the device being simulated, this was because of the various PES types, for example; Piccolo quadrature pattern and Heron tri-phase pattern. The cards in the Simulator also had to be changed for the same reasons.

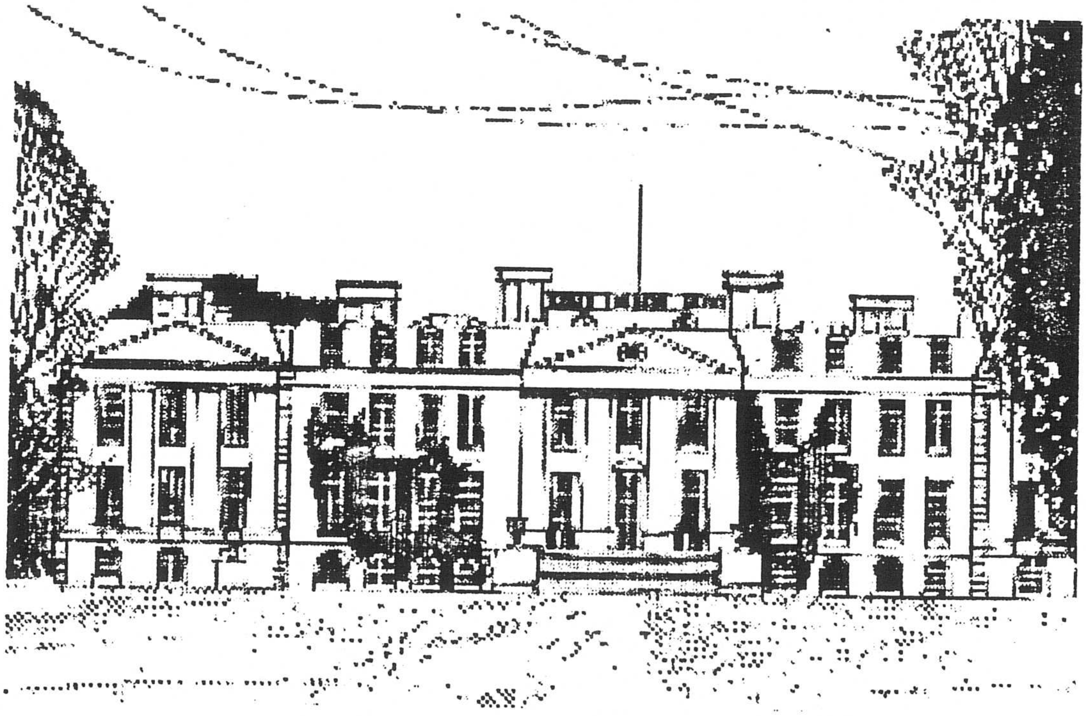

Figure 29. Back of Hursley House