Digital Archeology with Drive-Independent Data Recovery: Now, With More Drive Dependence!

[ELEN E9002 Research Project Final Report – Summer 2011]

Christopher Fenton
chf2110@columbia.edu

Introduction

The goal of this project was to recover the data from an 80 Megabyte CDC 9877 disk pack that potentially contains system software for a Cray-1 supercomputer that may be of some minor historical interest. It is quite challenging to recover data from obsolete digital media for a variety of reasons – functioning hardware can be difficult to come by, as well as difficult to interface with even if you have it, and magnetic media can degrade over time, especially if not stored in an archival environment. The target media for this project is a disk pack containing three double-sided 14”-diameter platters containing data – five data surfaces and one 'servo' surface, which provides alignment data for the other five surfaces.

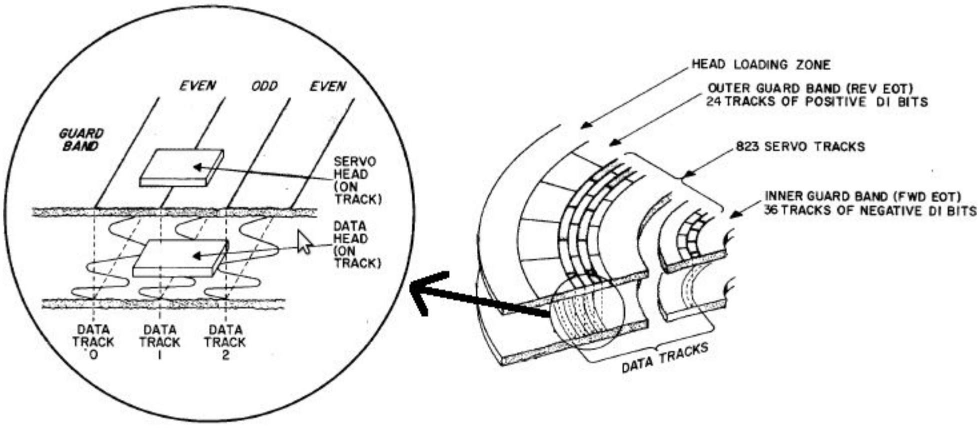

Figure 1: How data is stored on the disk pack (from pg. 74 of [5])

The initial plan for this project was to attempt to build a custom magnetic sensing platform that would allow me to recover the data without a working CDC 9762 disk drive. Research from the University of Maryland [1] had suggested that this might be a feasible approach for data recovery. Unfortunately, this scheme presented a number of difficulties which eventually proved overwhelming.

The primarily challenge was the relatively high data density. The disk contains data that is stored with a maximum linear density of 6000 bits per inch, on 823 concentric data tracks that are 2.5 mils wide. This means any particular bit might be a mere ~50x4 microns wide – a fairly tiny target that would require extreme precision to sense. A magnetic sensor was located [2] that actually had adequate precision (an active sensing area of only 1x2 microns), but then the problem (which was eventually determined to be insurmountable) became one of actually positioning the sensor. Nearly all magnetic disk drives work by allowing the read/write sensor head to 'float' above the surface of the disk. If a disk is rotating quickly enough, a thin layer of air will 'stick' to the surface of the platter. A magnetic read/write head in a disk drive effectively acts like a wing, floating above this thin layer of air – allowing it to float a few microns above the surface of the disk, as well as automatically adjust to minor variations in the surface height of the disk.

Unfortunately, the initial plan had been to use stepper motors and gears to rigidly position the head over the platter, with the platter mounted to a turn table. The turn table could then spin relatively slowly, while an analog-to-digital converter quickly sampled the data. It quickly became clear that it would be impossible to vertically position the sensor close enough to the surface to accurately sense bits while maintaining enough clearance to avoid collisions. Additionally, due to the way that servo data for all five data surfaces is contained on a separate surface, both the servo surface and the targeted data surface would need to be sensed simultaneously, which also meant leaving the disk pack intact and working within incredibly constrained physical dimensions.

An Exercise in Disk Drive Rehabilitation

At this point in the project, it became obvious that a multi-head sensing assembly that was engineered specifically to 'fly' above the surface of the disk was really needed. This also meant that the disk needed to be mounted securely and spun quite quickly (a few thousand RPM), and the analog-to-digital sampling needed to be performed that much quicker. Given unlimited resources and time, these are surmountable problems. Given the time and resource constraints of this project, however, it meant that I needed to find a working CDC 9762 disk drive.

I contacted Gil Carrick, who is the Director of the fledgling Museum of Information Technology at Arlington, in Arlington, TX, and whose website happened to mention that they had had a few of these drives in storage. After some lengthy logistical discussion, Gil agreed to lend us two CDC 9762 disk drives (in unknown condition), a CDC TB216-A Field Test Unit (FTU) designed for testing and calibrating the drives, as well as a spare disk pack for testing. We also acquired a Customer Engineering (“CE”) Pack from John Bachellier[1] with a company called MBI-USA that specializes in vintage computer equipment. A CE pack (as well as the FTU) is needed to align and calibrate the disk heads in the event that a head needs to be replaced, or the drive has become unaligned somehow.

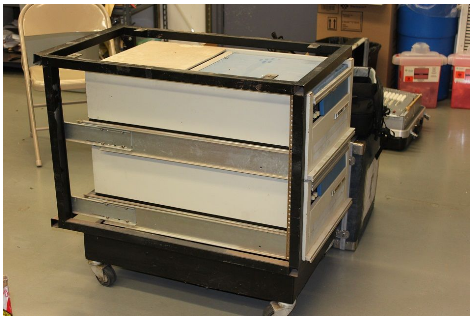

Figure 2: The two CDC 9762 Disk Drives shortly after arrival.

All of the equipment finally arrived on July 21st, allowing me to begin work. The first setbacks occurred almost immediately. Both drives had been sitting in some form of storage for at least two decades, and had acquired a fairly thorough coating of grime and/or filth. Disk drives are extremely precise, complicated electromechanical systems that effectively can't tolerate any kind of particulate contamination, so cleaning alone was going to be a challenge. Additionally, I had initially been working under the assumption that I had full documentation (including electrical schematics) for these drives [3], which would be an immense aid in debugging and repairing. Unfortunately, it appears that CDC produced multiple versions of the drive under the “CDC 9762” label. Both of the drives I was working with were manufactured in 1976, and appear to be CDC's earliest version of the drive. The documentation I had available belonged to a later version of the same drive being manufactured as late as 1985. Although the drive's mechanical parts were virtually

1 MBI-USA initially had a CE pack that was compatible with our disk drive that had been in their inventory for a decade or more, but it was apparently purchased by a customer from the US Navy while I was in negotiations with them. John Bachellier was able to contact a personal friend of his that happened to own one, and was able to sell it to us.

identical between versions, the newer drives contained a nearly completely reworked electrical subsystem (each drive is controlled by a 'logic cage,' containing sixteen circuit boards connected through a wire-wrapped backplane, as well as a handful of other boards scattered through the machine).

Both drives, when powered on, immediately asserted their internal 'fault' signals. The machine with the lowest number of hours on its lifetime counter (a mere 38,000 or so) was chosen for serious cleaning and debugging. A week or so of cleaning ensued before any serious electrical debugging was attempted. One of the largest problems encountered with the cleaning process was that the entire case of the drive was lined with 1/4” thick noise canceling foam that had degraded over time. Any contact with the foam would cause it to crumble into dust, something potentially disastrous if it were to contaminate the disk cavity, and ultimately all of it needed to be carefully removed. Additional problems were encountered from the large number of spiders that had taken up residence inside the disk drive, as well as a 3”-diameter (thankfully abandoned) “mud dauber” wasp nest [4] that had been constructed within the drive.

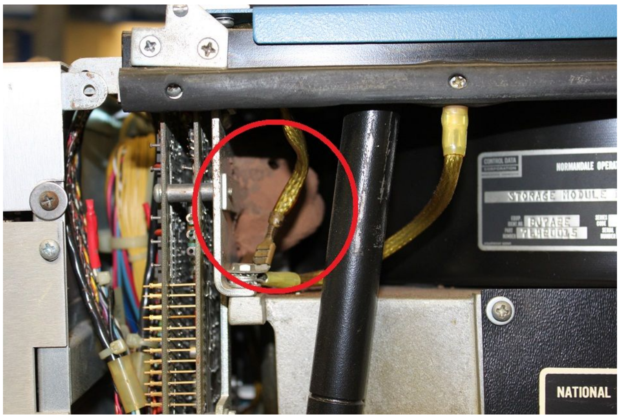

Figure 3: The spacious former home of a family of computer-savvy wasps

During the cleaning process, an internal status panel was located within the drive that indicated the 'fault' signal was being generated due to an internal voltage fault. The disk drives internally use +-42V, +-20V, +-12V, and +-5V, and the problem was eventually tracked down to a short circuit on the +20V supply. Through process of elimination, the fault was determined to be on a logic card located in slot 1 of the logic cage, although there were no obvious faults visible on the card. A replacement card was taken from the 'spare' machine which cleared the fault and allowed the machine to continue its boot process.

At this point, the FTU was setup and appeared to pass all of its internal diagnostics (thankfully, documentation for the FTU was available). When the FTU was connected to the disk drive, however, the drive remained unresponsive to querying. The same process was repeated with the spare disk pack installed in the drive, following which the drive spun up the disk and, following a 30 second delay, promptly burnt out a fuse on its +42V power supply and re-asserted its internal fault signal. Consulting the documentation available, it appeared that the primary use of the +42V supply was to drive the large voice coil responsible for positioning the head assembly. The head assembly, requiring extreme positioning precision, is constrained to only move in one direction via a system of bearings and guide rails. Some kind of lubricant appeared to have dried out and congealed on the rails and bearings, effectively cementing the head assembly in place. When the drive attempted to power the coil to load the heads as part of its initialization process, the coil was unable to move and a power surge resulted, blowing the internal fuse. Extensive cleaning of the rails and bearings ensued, but movement continues to be significantly stiffer than intended, potentially causing positioning errors.

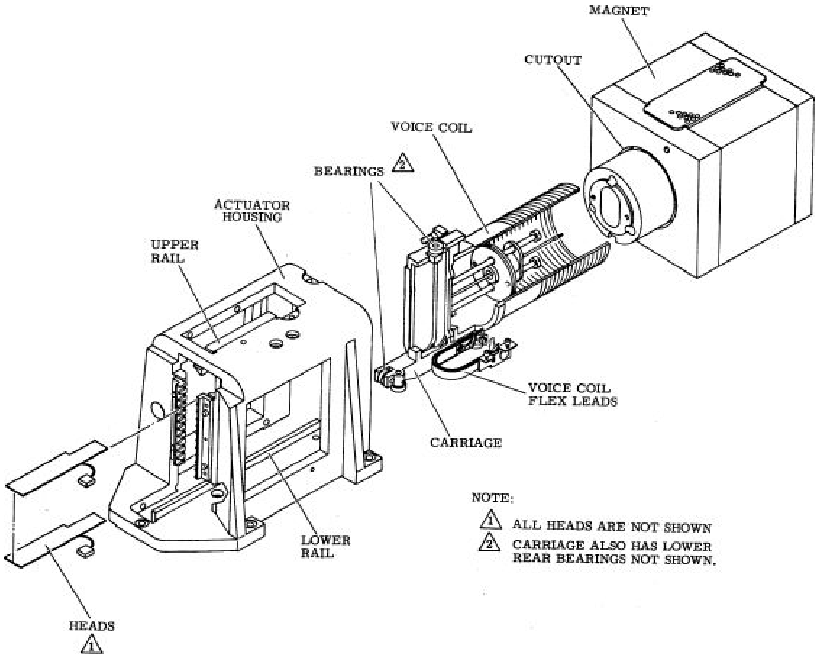

Figure 4: The coil and head assembly for a similar model of drive (from pg. 49 of [5])

As a debugging feature, the coil and head assembly can be disconnected from the power amplifier and manually positioned over the disk, so long as the disk is spinning faster than 3000 RPM (the minimum speed required to allow the heads to fly). This procedure was attempted with the spare disk pack installed, and the drive actually asserted its 'ready' light, which I believe means it had successfully sensed valid servo data and completed its initialization process. Unfortunately, within 30 seconds of the heads being loaded a high-pitched whining noise began to be emitted from the drive, implying a potential head-to-disk contact was taking place. The drive was then powered down and the disk pack and heads were carefully examined. Thorough examination revealed that Head #4 on the drive (which reads the bottom surface of the lowest data platter) had 'crashed' into the disk surface and scraped away a concentric ring of oxide material, permanently damaging the platter. This is a good time to point out the advantages of not experimenting with your primary source material when performing digital archeology experiments!

The offending read head was removed from the drive, carefully cleaned to remove the layer of oxide that had been deposited on it, and set aside until further notice. At this point, the spare disk pack was once again loaded into the drive (now with only four read heads) and spun up, and the heads were then able to be successfully loaded without further incident.

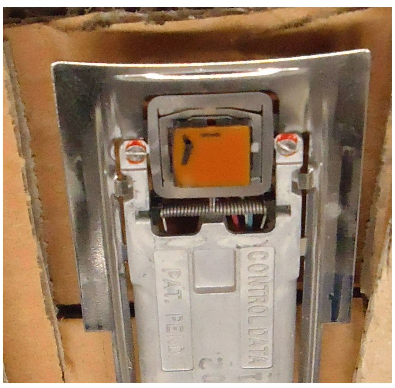

Figure 5: Exposed read head following cleaning

Reconnecting the coil to the power amplifier and attempting to let the drive continue initialization on its own, the drive would now progress to the point where it would spin up the disk and attempt to seek out the first data track (Track 0), before quickly retracting the heads and re-asserting its internal fault light. According to an initialization flow chart belonging to a different drive model in the same family [5], which appears to be identical across machines thus far, the drive appears to be reaching a 350 millisecond timeout without locking onto the start of the servo data while attempting to perform a 'load seek' operation. This could potentially be due to a number of factors, but the current most likely explanations seem to be:

- Due to friction in the rail and bearing system, the coil can not move quickly enough to lock onto the servo data before reaching its timeout.
- The disk and/or servo read head has suffered damage due to a head-to-disk contact, and is unable to function properly.
- The magnetic servo data on the disk pack being used has degraded over time, and the signal is not strong enough for the drive electronics to sense it properly
- Due to the large number of electrolytic capacitors used in the system, and their tendency to 'dry out' over time and suffer from somewhat unpredictable failure modes, the analog sensing electronics could be behaving improperly (this is the likely cause of the +20V short mentioned earlier).

Drastic Measures

With time rapidly running out on this project's end-of-summer deadline, it became apparent that debugging the myriad potential failures of the disk drive's electronic control system would lead to little but frustration and heartache. A more direct approach was needed – as much as possible of the disk drive's electronics needed to be bypassed. As mentioned earlier, schematics were not available for much of the drive's electrical subsystem, but as fate would have it, schematics *were* available for the drive's internal analog “read amplifier” (a fairly simple circuit that amplifies the weak magnetic signal coming directly from the read head sensor itself). If the read-head assembly could be appropriately positioned, the low-level analog data could be recorded directly from the disk and post-processed off-line in order to recover the underlying data.

To test this hypothesis, our poor test disk pack was once again installed and spun-up, and an oscilloscope was used to observe the (remarkably intact!) analog data signal coming directly from the read amplifier.

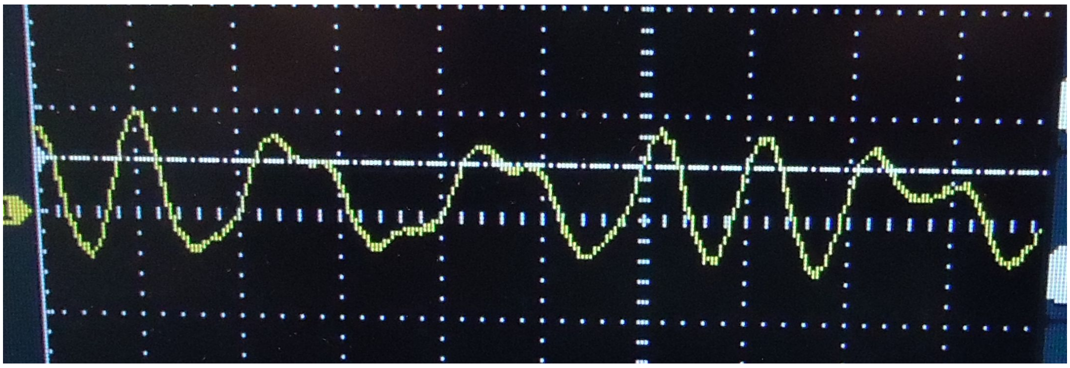

Figure 6: Analog data snapshot clearly showing MFM-encoding pattern

With confirmation that the amplifier was intact and working properly, a plan was formulated to quickly implement the necessary positioning and data logging system, completely bypassing the rest of the drive's problematic control system. For a more modern system, this would be a daunting design challenge. Fortunately, 35 years of technical progress have provided a number of useful tools for tackling such a problem quickly. A high-speed, Field Programmable Gate Array (FPGA)-based data logging system, along with a high-precision stepper motor and controller were chosen to provide ample (some would say overkill) margin.

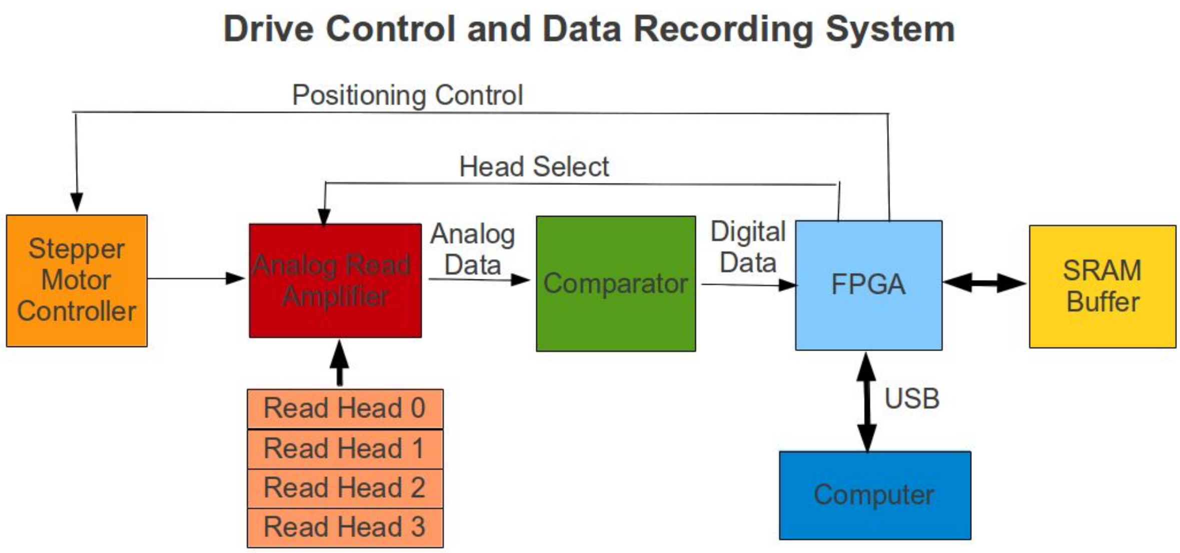

Figure 7: Proposed block diagram of drive control and recording system

Positioning Sub-System

The actual data on the disk is recorded with a track density of 400 tracks per inch. Feedback from the disk's servo sensor allows the drive to know exactly when its sensors are centered over the intended data track. Without the drive's control electronics working (including any feedback from the servo mechanism), a completely 'open-loop' control system would be needed. A mechanism driven by a stepper motor would be mounted directly behind the voice coil, and used to slowly 'step' the entire coil-and-read-head-assembly forward, across the surface of the disk. If the linear resolution of the positioning system is sufficiently high, one can guarantee (if somewhat inefficiently) that they accurately sense each data track by severely oversampling.

The positioning system was built from a modified Makerbot Thing-o-Matic [6] Z-axis positioning stage mounted on a custom, laser-cut acryllic frame. The frame was designed to mount securely to the rear of the disk drive and sit snugly behind the voice coil. The stepper motor has a resolution of 200 steps / revolution, while the acme lead-screw it is driving contains 13 threads/inch, and has four 'starts,' (which means that it requires 3.25 revolutions to advance the nut one inch). This would only give us a linear resolution of 650 steps/inch, insufficient to guarantee that we appropriately over-sample the data stored at 400 tracks/inch. Fortunately, the Makerbot Industries stepper motor controller thoughtfully supports 1/8 'micro-stepping,' so we can effectively increase the resolution of our motor by a factor of eight. This brings us to a total of 5200 steps/inch, allowing us to record 13 samples per data track, and effectively guaranteeing we get at least one accurate sample per track.

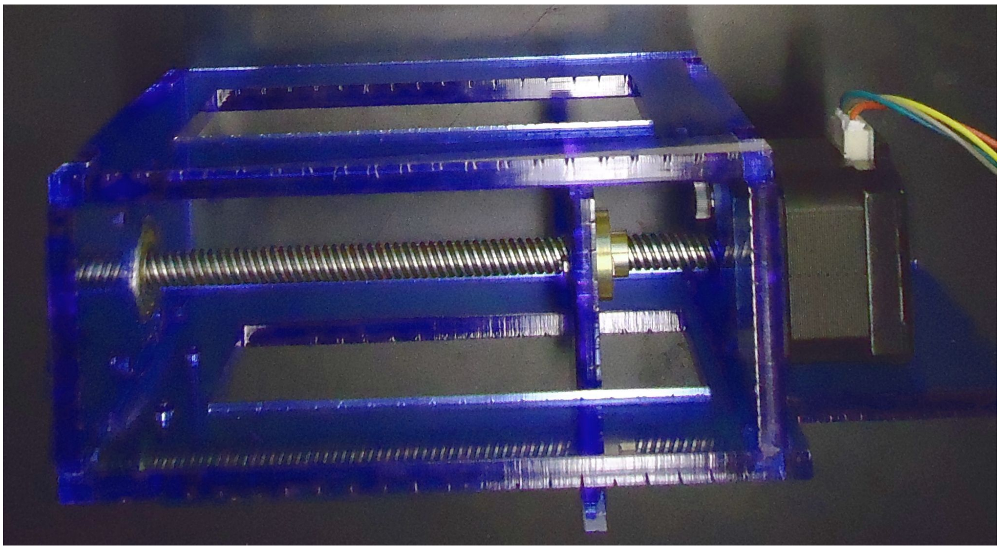

Figure 8: Positioning robot with stepper motor

Control and Data Logging Sub-System

The heart of the control and data-logging sub-system is a Digilent Nexys2 FPGA development board. FPGAs allow one to rapidly create high-speed digital logic systems that enable nano-second level of control. For each step of the positioning system, the output from each of the four remaining sensors is fed through a high-speed comparator and eventually logged by a computer for later analysis. The comparator acts as a 1-bit analog-to-digital converter – sufficient resolution to decode the 'modified frequency modulation' (MFM) technique used to encode the data. Each 'bit' flies under the magnetic sensor for approximately 103 nanoseconds (9.6 Megabits/second), so to ensure accuracy, our FPGA records a sample every 12.5 nanoseconds (~80 Megabits/second, or roughly 8X faster). The disk is nominally rotating at a speed of 3600 rotations-per-minute (RPM), so to capture one complete data track, we need to record data for 16.67 milliseconds. Continuing with our design-theme of including a healthy 'margin' in our sampling, the FPGA buffers 67 milliseconds of data (roughly 4 revolutions) at a time into an on-board SRAM chip before eventually sending it back to the control computer over a high-speed USB interface.

The FPGA is controlled via its USB interface from a driver written in C++ that is running on the data-logging computer. The FPGA also contains a small amount of logic to advance the stepper motor when directed by the computer.

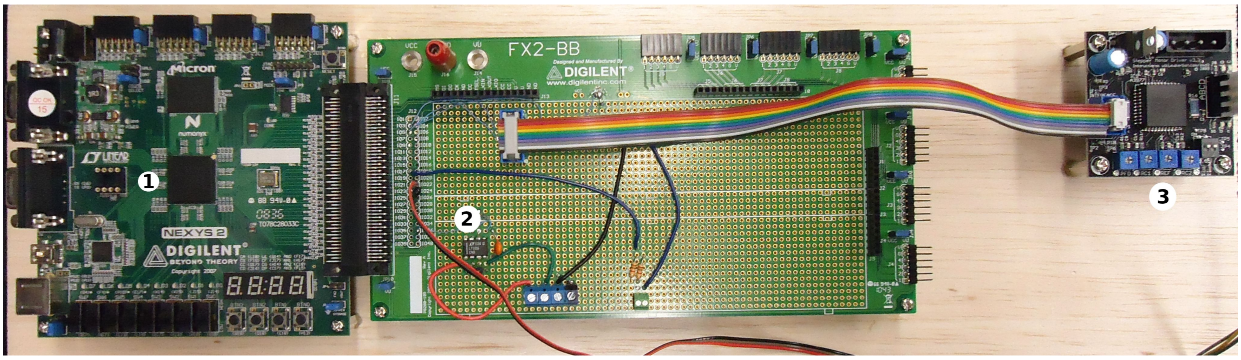

Figure 9: The FPGA (1), analog comparator (2) and stepper motor controller (3)

Putting It All Together

With the positioning system and control and recording electronics completed, the entire setup was mounted to the disk drive for testing.

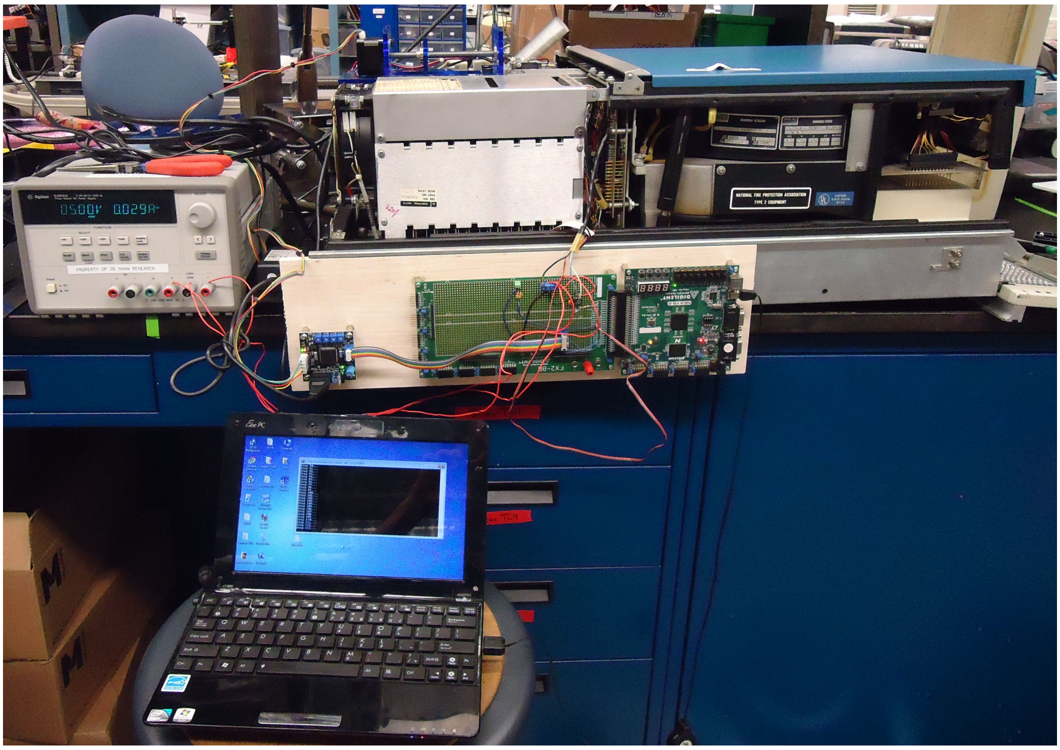

Figure 10: Final setup with electronics and positioning robot mounted

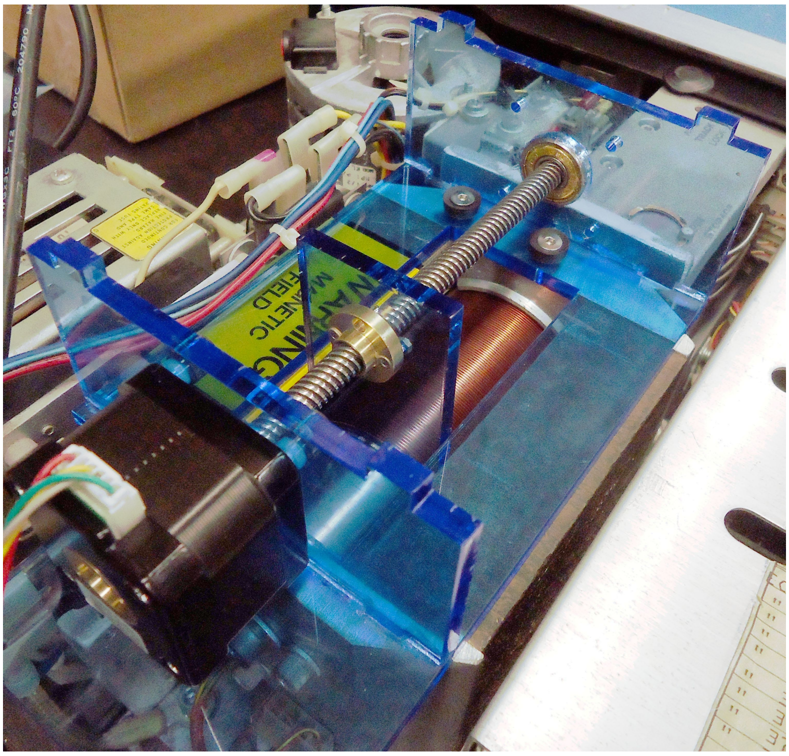

Figure 11: Positioning robot securely mounted behind voice coil

Figure 12: The moment of truth – the Cray-1 disk pack installed in the drive

An oscilloscope was used to verify that the analog data being read out from the disk was being appropriately converted to digital form by the comparator, and the data being sampled by the FPGA was tested and confirmed using a known data pattern.

Figure 13: Analog data (yellow) versus inverted comparator output (blue)

With everything tested and working as intended, the system was first used to record all four data surfaces of the Cray-1 disk pack accessible via the remaining read heads. At this point, the 5th read head, which had been removed from the drive (and carefully cleaned) following the earlier head crash, was re-installed in the drive. Typically, re-installing a read head is followed by a delicate re-alignment procedure needed to ensure that the sensor is in perfect vertical alignment with the servo head. Fortunately, our recording system ignores the servo data completely, conveniently allowing us to forgo the alignment procedure (which would have also required working drive electronics). With the now-clean read head reinstalled, the Cray-1 disk pack was re-installed, prayers were issued to the disk drive gods, and the head assembly was loaded. The cleaning procedure was apparently effective as the head loaded without incident, and the remaining surface of the Cray-1 pack was successfully scanned. With the Cray-1 disk pack scanned, the test disk pack was also scanned in a similarly uneventful manner (albeit at somewhat lower spatial resolution for the sake of timeliness) in order to provide a set of comparison data. All told, over 34 Gigabytes of data was recorded from the Cray-1 disk pack, and 8.75 Gigabytes of data was recorded from the test disk pack.

Future Work

With the target disk pack imaged with as high resolution as was practical, an enormous amount of data was generated. To actually recover the data will likely be every bit as challenging as getting the raw data off of the disk, and a great deal of work will need to be done in terms of signal processing and analysis. At a basic level, the following steps will need to be performed:

- For each 'sample,' a single revolution of the disk will need to be isolated from within the 40 mS snapshot (perhaps merging the data from all four revolutions to increase accuracy).
- All of the samples will need to be analyzed to determine which ones are properly 'centered' over data tracks, and which ones contain noise.
- Once a proper 'track' has been extracted, the track needs to be analyzed to determine the beginning and end of the track, as well as how many data 'sectors' each track contains.
- With each track divided into proper sectors, the binary data 'payload' can be extracted from the raw MFM-encoded data
- With the actual data extracted from each sector, work will need to be done to extract the underlying file system structure, as well as individual files.

Although the actual data analysis is beyond the scope of this paper, some very preliminary analysis shows somewhat promising results. As a simple experiment, a series of 39 samples (~3 data tracks) was extracted from roughly the middle of the surface recorded by head #0 (steps 5000-5038). Each sample was analyzed for long, contiguous streams of sampled 1's or 0's, under the assumption that valid data tracks might contain such features and noisier inter-track samples would be less likely to contain them.

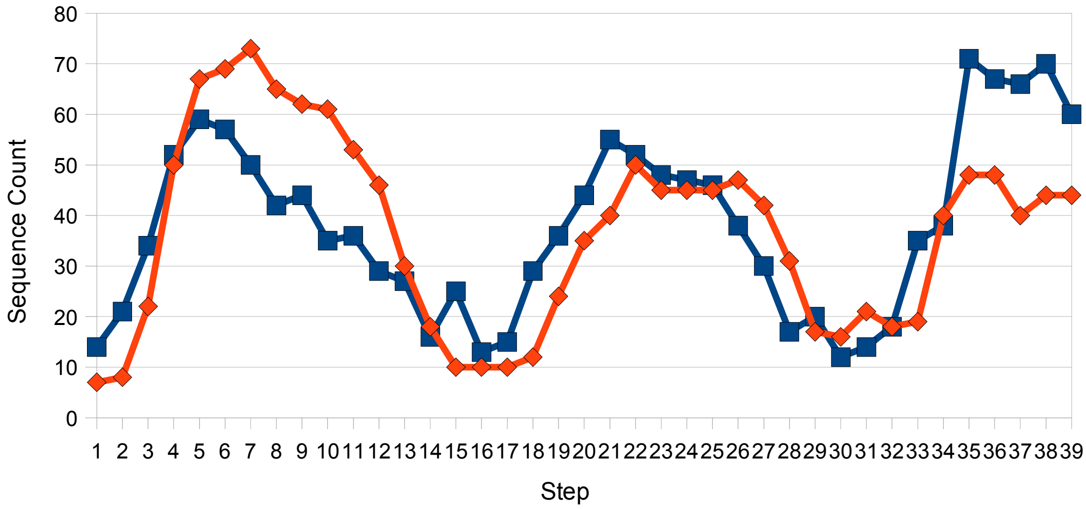

Figure 14: Occurrences of 24+ continuous 1's (blue) and 0's (red) vs distance

This data was captured with a theoretical spatial resolution of 13 samples per data track, so if the number of long sequences of 1's or 0's is correlated (negatively or positively) with the sensor being properly centered over the data track, we would expect to see a pattern recurring roughly ever 13 steps or so. Figure 13 clearly agrees with our expected result, implying that this might be a useful metric for identifying properly 'centered' samples.

Conclusions

This project has been an interesting and somewhat promising foray into the nascent world of digital archeology. The world is currently undergoing a rapid shift from easily-readable, long-lasting, low-density archival media such as paper or microfilm to hyper-dense digital storage mediums. As we hurdle towards an all-digital future, it is worth pausing for a moment to consider some of the challenges associated with maintaining long-term access to digital media. Within the past thirty five years, the CDC 9762 disk drive used for this project transitioned from cutting-edge storage technology to vanishingly rare antique. Fortunately, the same technological forces that have left this drive laughably obsolete have also given us the tools to allow a single engineer to potentially overcome these challenges. Digital archeology as a field, for both historical and forensics-related reasons, is likely to continue to grow in importance for the foreseeable future.

References

[1] C. Tse, C. Krafft, I.D. Mayergoyz, and D.I. Mircea, "System and Method for High-Speed Massive Magnetic Imaging on a Spin-Stand," US Patent 7,005,849 (2006).

[2] “TMR Magnetic Microsensor Probe.” 2011 MicroMagnetics, Inc. 28 Aug. 2011. <http://www.micromagnetics.com/product_page_stj030.html>

[3] “CDC Storage Module Drive – BK4XX / BK5XX Hardware Maintenance Manual.” 2011 Bitsavers.org 27 Aug. 2011. <http://bitsavers.org/pdf/cdc/discs/smd/>

[4] “Mud Dauber – Wikipedia, the free encyclopedia.” <http://en.wikipedia.org/wiki/Mud_dauber>

[5] “CDC Storage Module Drive – BK6XX / BK7XX General Description, Operation, Theory of Operation, Discrete Component Circuits.” 2011 Bitsavers.org. 27 Aug. 2011. <http://bitsavers.org/pdf/cdc/discs/smd/83322320H_BK6xx_BK7xx_GeneralDescription_Jul80.pdf>

[6] “MakerBot Thing-O-Matic 3D Printing Kit.” 2011 Makerbot Industries, LLC. 30 Aug. 2011. <http://store.makerbot.com/makerbot-thing-o-matic.html>