Oral History of Lloyd Thorndyke

Interviewed by:
Tom Burniece and Dag Spicer

Recorded:
August 13, 2010
Cisco Video Conferencing System

CHM Reference number: X5969.2011

Tom Burniece: Good afternoon, this is Tom Burniece, a volunteer with the Computer History Museum in California, and I happen to be in Bloomington, Minnesota today alongside of Lloyd Thorndyke. We're going to have a discussion with Lloyd about his history at Control Data Corporation, which spans almost 30 years, on both the peripheral side and on the computer side. As a result, we're going to have two interviewers: I'm going to work on the peripheral side and Dag Spicer, who is in a Cisco facility in San Jose, is going to discuss the computer side.

But before we do that I'd like to give some acknowledgements and thanks. In particular I'd like to thank Claudio Desanti, Landon Noll, and Craig Tay of Cisco for setting this up. They do this as a community service for organizations like the Computer History Museum and we really appreciate it. This is the second time that I've been involved in one of these remote interviews and the quality has been outstanding. I also want to note that I'm not going to be on camera, although I'm actually to the side of Lloyd, so I may be a bit of distraction for him, because he should continue to look at the camera, not look at me, as I ask questions from his side. Hopefully he'll keep his attention on the screen in front of him, which actually shows Dag Spicer.

So with that as a quick introduction, I just want to give one piece of qualification here, I did work for Control Data from 1969 to 1981 in the disk drive operation and worked under Lloyd for three of those years from 1971 – 1974. As a result, I am familiar with that part of the history but what I'm really interested in today is how Control Data got into the peripheral business in the first place, which started in the early sixties, when he joined the company. Thus, we're going to spend most of the time today talking about what happened in the sixties, with a little bit in the seventies, and then I'm going to turn it over to Dag to talk about what happened when Lloyd moved over to the computer side in 1974.

So Lloyd, why don't you just tell us a little bit about your family background; where you grew up; where you went to school, and anything else you want to tell us about your family in your early years?

Lloyd Thorndyke: I grew up in the southwest corner of Minnesota, in a place called Pipestone. It is 15 miles from Flanders, South Dakota where Mr. Amdahl of IBM grew up, so that area spawned a couple different people in the computer business. One of the other early people there was Ben Gurley, who was a designer at DEC. He grew up a couple blocks from me, but I didn't know him that well. I graduated from high school in 1945 and of course I had registered for the draft. So the next month, I took the pre-induction physical and was drafted into the Navy, which was unique. People have asked me how you get drafted in the Navy but the Navy had a qualification that they wanted me for. So I went to the Navy, which was fortunate because I was in the Navy Reserve, and when war was over, the Reserve started getting out. So I spent about a year in the Navy. The last six months in the Navy I spent at the Naval Research Lab in Washington, D.C. as a student. I was also a teacher, because after the war, the regular Navy people had to come back and get qualified for radar if they wanted promotions. Radar technicians were where the advancements were and that was where the need was. The problem was that radio operators had been on board ship for four years and probably never had to add two and two, while radar was all mathematics, with vacuum tubes and components. It was the math that they didn't understand, so at night we would sit in the cafeteria for three hours and I'd try to tell them what the instructors said that day. As a result, I was kind of an interpreter for them and was the last person discharged out of the Naval Research Lab, probably because I was trying to help the regular Navy guys. They even offered me a couple levels of promotion if I would sign up for two years to go to the Atom Bomb Test but I needed to go home. My older brother was shot down and killed in late March of ’45 and so the family ties were too important. Anyway that was where I went.

After being discharged, I went to Hamline University in Minnesota, which was predominantly turning out teachers and preachers in those days. But I got a BS in Physics with a Math Major and then started trying to find a job in 1950. The Korean War started but I was exempt from that because of the year of service, so I found a job working for a small company building neon sign materials and neon sign chemicals. They wanted to get into the electronic business, so we started making unique carbon composition resistors, and I worked for them until I got acquainted with the head personnel person at Univac.

He recruited me in 1954 to join Univac, where I got into a mechanical group that was uniquely interested in explosives. They had a series of contracts from the arsenals, including Picatinny Arsenal, who wanted to develop a high performance explosive using an electronic detonator, as against a mechanical detonator, since it would be faster. The real problem was the speed of triggering the explosive in shooting down a mach 3 aircraft, because they couldn't always depend on getting a direct hit. The radar system could lose the aircraft in a radar signal and when you did, the missile would fly an Archimedes spiral from the back of the aircraft up to the front and then throw out a series of bomblets and hope that the aircraft would fly into one of them and be destroyed. But the performance of the bomblet was one of the problems, so I did the testing on it. As such I came to the realization that the technology of caution of a normal mechanical unit was if the mechanical fuse didn’t work, the bomblet would stay in rebutment for 24 hours, to make sure the fuse doesn’t cook off. Well none of that existed for electronics, since it had a capacitor for energy storage, and after about four seconds that capacitor is dead from leakage. So when the power is zero, you can work with it all you want to without getting hurt, I thought. When you think about it afterward, however, you come to the realization that you're not going to have a long life <laughs> in that business, because there are no rules on how the electronic fuse will behave with various component failures.

So I got interested in designing peripherals at Univac, where they developed a system for the FAA or the CAA to automate the en-route traffic control unit, as pilots requested airspace across the United States. That required getting an automatic teletype from an area control unit to another area control unit that we’d request from Minneapolis to the East Coast. I designed some modules and installed the system in 5 centers. I then got an opportunity to transfer to NTDS, which is the Navy Tactical Data System. A book has been written about it, called *When Computers Went to Sea*. Univac had an exclusive love-in with the Navy, since it was formed [by combining Eckert-Machley Computer Corporation in Philadelphia with Engineering Research Associates (ERA) in Saint Paul]. ERA was mostly ex-Navy people who did some code-breaking work during the war and then were spun out as a company into Saint Paul, because there happened to be a war surplus plant there, and the Navy wanted the technical effort to continue.

So I started working on NTDS and it was a sole source project. I was transferred to NTDS, mainly because Control Data had been formed in the meantime and most of the people forming Control Data came from NTDS. So they took a few engineers and transferred us to NTDS, figuring that we were not recruitable <laughs>. I ended up having all the peripherals for the NTDS but some of the algorithms were so classified, that I was not allowed access. For example, how we digitized the radar return to display this on a scope was highly classified, as was how we could identify an aircraft. So I was not deemed to need to know that information. I could direct the people and didn’t have a problem with that, since I probably couldn't understand the algorithms anyhow. But there were areas where I had to tell these people that some of the Mil-Specs didn't apply, like the computer system had to take a 50G shock as required by MIL_SPEC E16400. Computers in the 1960's would not work after a 50G shock; you shock everything out of it, so that really didn't apply. But they certainly wanted to slam the computer at 50 g's. The logic was transistor logic with diodes in it for AND/OR diode combinations and we didn't understand how to ground cabinets very well. So when they got the first NTDS computer into Florida, they turned a search radar on it, hit the button, and we burned out about a million dollars worth of diodes. But we learned how to ground cabinets and we learned to do shielding, which was very critical in order to be immune to EMI / EMC. We used all of that learning later on at Control Data in our peripherals.

We never had power glitches take us off the air in our peripherals at CDC because we had line filters on every piece of peripheral equipment we designed. Of course, the cost reduction people felt that that was worthless but it was not. We had a system that would stay operating and we couldn't allow the system to turn off. Now the CDC 6600 and all the computers Seymour Cray designed used 400 cycle power; not 60 cycles, so we had the MG set isolation from the power lines. However, in the CDC peripheral group, I felt we had to be immune to power surges and power cycles. No one gave us a lot of credit for it, but since I was running the program, I put in what we needed. Later on we had an opportunity to write the design plans for follow-on peripherals, where I insisted you always put the line filters in. So that was just one of the aspects.

Anyway after Univac, I was interested in working at CDC. Because of NTDS, there was a gentleman's agreement with the Navy that CDC would not more recruit people from the Naval Technical Data Program. Part of the mechanical team I worked with at Univac, had gone over to CDC, so I decided that was an opportunity. I wanted to join, and hadn't been told that I was supposed to be frozen, so they hired me. That created a little hell, but they got over it. So now I'm over at Control Data.

Burniece: So what year was that?

Thorndyke: 1960.

Burniece: Okay and tell us a little bit about the culture and the company at that point.

Thorndyke: At CDC?

Burniece: At Control Data, yes.

Thorndyke: Well the culture in Univac at that time was different. The management tended to operate in somewhat of a vacuum, instead of telling us what was going on. There was not a culture of trying to keep the people together and talking with us and listening to us. If it was within the division you'd do it, but you hardly ever saw management, and when you did, you were probably in trouble. At Control Data, the interesting difference was that management was involved. Control Data management talked to us; they got us together in a big conference room monthly. The other thing was that there was a yearly performance review. At Univac I can't remember having a performance review in six years. You may get a raise now and then but the performance reviews were kind of "Well, you're doing a good job." It was not like Control Data, where you ended up writing a review and pointing out to people, “Well here’s what you need for your promotion. Here’s what you need to study or get ready for a promotion.”

Because Control Data was growing so fast, we tried to promote internally and point out that you needed some education for management one way or another. So Control Data was very high on internal education. Now later on, that of course got to be PLATO on a CDC delivery system, and you were expected to have a couple weeks of PLATO training every year. I don’t know if you know what PLATO was but it was a computer learning system out of the University of Illinois that we funded for years. I would say this about PLATO - it was the first attempt to use a super computer as a switcher and each system could handle 1,000 customers online synchronously. The supercomputer would do the calculating and you each had one millisecond worth of time. So you would send in the data and it would do the calculating and send back the results. When you look at the current systems, there is little difference. You had a computer in the terminal, doing the work. You were networked together and talking to different people, so it was a high performance network. PLATO was ahead of it’s time but CDC didn't spin the technology into its computers, because the champion of the PLATO was Mr. Norris [William C. Norris, co-Founder and CEO]. In Control Data, we had champions for different projects.

Seymour Cray was the champion of the computers. He reported to Norris, but he didn't take direction from Norris. Cray was the champion of it and Norris was in the background giving him the resources and keeping things on schedule one way or another, while not allowing interference. The same was true in peripherals, where our champion was really Bob Perkins. Bob Perkins was Norris’ first hire at ERA and he stuck with him the whole time. He was a staff man, almost from start to finish. We always kept saying he had “staff infection”, but he was the man who wrote the early design plans on what we’d do on a tape transport and on the first disk file. We worked very closely with him and he was out with us a lot. So we had champions. Any time CDC had a program that didn't have a champion who stayed with the program, the program was eventually terminated.

I learned one thing. If you lose a champion, you might as well get rid of the program, because the continuity is gone. CDC tried a couple of programs, where the champion changed and it was worthless, because the new champion didn't understand what we sold and they didn't try and read the specification of what the customer wanted. Every time they thought the customer needed something else, they thought that was a new design change. So it was a fluid area and you never went back to what you sold. As a result, I learned early on, that you own the program, and you stay with it from start to finish. I have always tried telling people, if you're going do the program, I don’t expect you to leave halfway through it. If you want to leave halfway through it because you're dissatisfied, start thinking about it. We’ve got a program that’s going to be three or four years, we want to stay with it. So the culture was, at that time, you stayed with it. You stayed with it because you were in advanced technology and you were doing new things all the time. Also the culture was that, other than a couple different examples, you avoided recognition.

Seymour got recognition, but he got recognition from marketing, trying to enhance CDC’s image by having this supercomputer designer. But you didn't find interviews with Seymour; you didn't find him doing interviews or anything else. The big joke was you couldn't even call Seymour on the telephone. The telephone bothered him. I remember he yanked the telephone wires out of the wall, because they’d call him too much, so everybody was afraid to call Seymour. But I talked with him off and on because we were developing peripherals for him and continually wanted to make sure that he agreed. So I’d bounce it off him – “well how does this sound?” Bob Perkins was also always interfacing to Seymour, as we developed peripherals for him.

But anyways, the culture was do your job and don’t brag about it. So Control Data didn't brag and, of course, we didn't write articles about what we did in peripherals, because it would take an engineer out of project work. We had the next program that was needed and we couldn't take the time to write anything, plus if you did write something, there were recruiters all over us. So I discouraged any writing because that would mean I would have a problem of recruiting. Our people were highly recruitable but there was no way that the recruiters could find their way in, because nothing was written down. So that was the culture of Control Data.

We also had a lot of money. When we decided to do the 808 disk file in 1963 - 1964, I only had about 25 people in engineering. We had a million dollars to finish the 606 tape unit and start the disk development. We were under a schedule driven by 6600 deliveries, so time to the product was essential. Well with 25 people, how do you spend the money? We decided we had to spend $4,000 dollars a day, being there were 250 working days, a year. So at $4,000 a day, how do you do that? We selected the technical approach we wanted to pursue, and then started funding the backup approaches. Like we wanted to paint the disks, but we felt if that failed, we’d need to do plating. So we farmed plating out to the research people. We ended up having other people who were funded as backup in other areas. We used hydraulic actuators in our first disk system, because of the mass we were moving and the performance we needed for Seymour. At the time we had the 606 tape unit done The 606 tape ‘family’ had 7 different products, from the slowest (CDC 601 at 37.5 ft/sec) to the fastest (CDC 607 at 150 ft/sec). However, the 601 and 602 were one machine and the other 5 were the same design. Only the tape speed differed. Unless you want to talk more about the tapes, we’ll go to disk.

Burniece: We should finish talking about the tapes, because I think you did the tapes first, right?

Thorndyke: We did the tapes first.

Burniece: What was it that actually caused CDC to decide to develop their own peripherals? What was the impetus of that?

Thorndyke: IBM would not sell peripherals OEM and our customers required that we incorporate tapes onto our computers. If we bought IBM tapes at full end-user prices, it was a losing proposition. It would cost us money even though we could make up some on the computer. In the meantime, the only tape unit available was the Ampex tape unit, the FR300 and later the TM2, and that unit ended up having “morning sickness”. When you turned on a computer in the morning, it’d be two hours before you had a reliable tape unit. The difference was IBM tape units had a tape cleaner that looked like a razor head and it would clean the nodules off the tape. As the tape units were used, they ended up having tape debris that was pressed in the tape surface when the tape was wound up. When the tape was unwound, the debris nodules were stuck to the oxide, so you had to scrape them off. Ampex never understood that and therefore they didn't do it. So what you had to do is you had to run it back and forth for a couple hours before you could use it in the system. So CDC had a hell of a time delivering computers, because the Ampex tape system was unreliable. As a result, CDC finally decided that they had enough money in 1960 to fund a tape development.

Now over at Univac, we had designed a tape unit for Univac computers, however, it never went any place; but it gave us some prior experience. Paul Bulver had developed a voice coil valve there that allowed us to start a tape unit in milliseconds, so we could start the tape moving faster than a pinch roller. The problem with a pinch roller in a tape unit is it burns the tape as it pinches down with a big crush, and accelerates the tape to cause local stretching. The tape is accelerated and it burns it with a polish, which is probably a better term. So by using a vacuum and sucking it down to a rotating wheel, we didn't damage the oxide surface of the tape, and controlled the acceleration. But we also incorporated in a tape cleaner and our tape cleaner was three carbide blades, so if the tape had a wrinkle in it, we cut it in half. That was one of the problems with a damaged tape, so if we couldn't read it then it was usually cut in half and that was a bad tape anyhow. But in any case, with the tape cleaners, we could achieve very reliable operation and we'd run 150 inches a second because we wanted a quarter to a third increase in performance over IBM. Any time you go against IBM, you want at least a third more performance or a third less price. Well it's easier to put the performance in than take the price out, because the price comes out of your gross profit, so that was our philosophy all the way through Control Data.

IBM would not sell OEM at all but Control Data had a very high level, reasonable working relationship with IBM. There was an IBM Vice President who the senior people could call, as an example, if we wanted a copy of their tape specs. We would ask if they would send us the specification for buying magnetic tape and that was not IBM classified or anything else. They would send us the same specifications they sent to the vendor of the tape they wanted manufactured as a second source. This gave us an insight so that we could begin to say that our tape can be compatible with IBM because it was to the same spec, with the same recording performance, especially in the high frequency pulses that come off the tape. That relationship continued way into the disk business, as well. It probably didn't stop, even though we sued IBM for antitrust, since we had a cross license agreement with IBM for patents. They had needed some of our supercomputer patents, so they put their stack of patents about 36" high on the desk and said, "we think you're violating some of those patents "<laughs> and we put our little stack on the desk and said, "yes, we would consider swapping if we had a "cross license" which was good for us and I presume was good for IBM too.

But anyway the tape transports that we started with gained CDC a reputation in the tape business, so AT&T came and wanted us to develop a special tape unit for their geo-ballistic computer that they were developing for the government. So I proposed to them that we do a two-inch tape unit that would have 36 tracks [which was really four 9-track ½" tape heads in a 2" frame], with the ability to intermix computer data and function data. They wanted to record the missile radar data, plus record all the other data on the same tape, so they could play it back as a simulation, as they started writing their software. Over time, they found out that they didn't need to shoot a missile, since their computer could generate what a missile track should be and record it on a tape in real time and then read it back as if it were the missile shot. So that's the first time I ran into a numerical simulation using tape. They wore the heads out on the tape unit we sent them after about eight years, so we built a new set of heads for them. I don't think it's operating any more but it was way in advance of anything that we'd ever done. We'd recorded it at 2700 bits/inch at 10 million bits/second off of the tape unit. That's the speed they wanted and that was good for us, but we never made a product of it.

Burniece: What year was that? Was that also very early sixties?

Thorndyke: That was probably about the mid sixties. Dale Strand is still around and is the person you can reference to find out anything more about that if you'd like. We had to buy a 3300 memory and use it as a disk buffer for the 36-track head, since shift registers weren't available at that time. We did wire that memory as a shift register and got all of the bits lined up.

Burniece: Let's go back just a little bit; the initial tape drive that you developed when the Ampex tape drive was not satisfactory for Seymour's work. I believe it was called the 606.

Thorndyke: Yes.

Burniece: And I believe you showed that in San Francisco at a computer conference somewhere around 1962.

Thorndyke: Yes

Burniece: And CDC ended up getting serious interest in supplying tape units to NCR, plus a couple of other people. So that became the beginning of CDC's OEM business, right?

Thorndyke: Yes, for peripherals but it is interesting that the CDC OEM business actually started out selling computers. In the early years, Seymour first developed a quarter-scale, 12-bit computer called 'Little Character', I think, and that proved his technology capability without building a big machine. The 1604 was the first CDC computer. However, the first OEM computer was the 160, productized and called 'Little Character'. NCR bought the 160 computer from Control Data and put it in their product line, so that was the first CDC OEM product. I don't think Control Data priced it to make a lot of money on it but the learning curve took us down where we could, so therefore we already had a love-in with NCR as an OEM customer.

The problem with NCR buying the 606 tape drive OEM in 1962 was it had a CDC interface and they said, "well gee, we can't buy the tape unit, because we would need a new controller and that it is going to cost too much time and everything else." I was sitting next to Tom Kamp, [General Manager of the Peripheral Division [8]], so I leaned over and said, "Let's have a private conference." So we went out and I said, "Tom, I can put an Ampex interface on our tape transport, and we will be plug compatible, so they won't know what they're talking to." He said "You can do that?" I said "Yes." He said, "How much is it going to cost?" I said, "Ah hell, it's going to take probably another 25-30 transistor cards. We don't have to change anything else and we've got to change the interface anyhow, because they've got to have a non-CDC interface." Even thought we sold OEM we never allowed the CDC interface to be sold OEM, so no one could buy one and plug us, and we never published the CDC interface. That is why on the OEM side of Control Data it was called a different name, not the 606. You couldn't buy a Control Data interface peripheral OEM. The 852 disk drive, the 6600 computer, and the 606 tape transport were never sold OEM. It was always the 9000 series with the customer's interface on it.

If we got a request to give an OEM the NCR interface, which was proprietary to NCR, we would not quote it. That way we could keep the OEMs independent of each other. Every OEM put on their own operating panel and they put on their own I/O connector which made it very unique. That first tape transport then became quite a standard. We sold a lot of them to the Navy for NTDS because their data centers couldn't use the Mil-Spec tape unit that was built for ship board operation, since it was too damn costly. It was also based upon a tape transport that was not that reliable, because of the necessity to have a Mil-Spec design.

So anyway I don't remember how many of the tape units were sold but we had 800 bits/inch and it was very reliable. Wally Edwards was a very instrumental engineer in that. Wally was deep into simulation and would have been a great QA guy, because after we built this 800 bit/inch tape unit, we were in production. Wally came in one day and always called me "Hoss". He said, "Come on over, Hoss, I want to show you something." So he took a CDC 607 and wrote a roll of tape. He rewound it, and we couldn't read it. He could read it after he had written it okay but when he turned it around and tried to read it again, he couldn't read it. I said, "What have you done?" He said, "Well, all I did is save a garbage dump. I just start writing on a tape and when I saw an error, I stored that data in the garbage dump and kept writing until I get a mess of errors. I did it two or three times." The problem was, our tape head was operating as a magnetic amplifier, so when you read certain bit patterns, you'd begin to generate flux through the head that caused unreadable tapes. So you could write a tape that was good when you wrote it, but it was bad after a read, because of that effect. Well that caused a little bit of heartache, as you would imagine. In fact, the problem was the edge effect in reading made the head core look like an infinite gap, as against a normal gap, so it was picking up infinite gap flux. We solved that by increasing the radius of the pole piece corner to reduce the flux created. But it taught me Christianity pretty fast too. Here I've got a tape that was perfect on everything, except for the fact that I got a guy who proved the customer could make their data unreadable on the tape due to the magnetic heads and a unique bit pattern on the tape. But anyway, we fixed that and we also did some other things.

A marketing guy came to me one day and said "Lloyd, you know we can increase our performance if we could read our CDC half-inch tapes backwards." He said, "You know, you've got to rewind to the front but it would be nice if I could read backwards, if I only want to go back so far." He said, "If you can do that, I can reverse the order of the data in memory a lot faster than if I wait to read it forward." So I ended up saying, "You know, I can read a tape backwards; I probably can read the IBM tape backwards." So I put in the extra electronics and Control Data then would read forward and read backwards rather than rewind. I remember the IBM people saying, "Damn it, you don't know the amount of headache you gave us by reading your tapes backwards, where we can't." <Laughs> I said, "Well that's too bad, you know, you've got to have a little fun in this world."

But then the other unique thing is CDC wanted to get into the seismic business and Texas Instruments had a 21-track tape that was too slow a reader for the 6600. So they asked if we could read a 21-track tape. I said, "Get me the head specification and we'll see what we can do." They got a head specification and it was for a one-inch tape, so we just added a half inch to everything in the front of the tape transport and developed a one inch, 21-track tape that was so successful that we then sold them to many of the seismic companies for supercomputers doing seismic work. That was very good for us, since it opened a new market to us.

All of this was possible because of the fact that we had such a flexible design. One of the other things in our design philosophy that I got accustomed to, not that I had initiated it, was when we started the design of a unit, we had 80 percent mechanical engineers and less than 20 percent electrical. The reason was that the electricals can't do anything, unless the mechanicals get it right, and then electrical engineers can come in and patch it together. Then they can come in and do power engineering, interface design, logic cards, connectors, etc. One of our requirements in packaging, was when you go to service a product, you can't take something apart that is working, in order to fix what's behind it. So therefore everything has to be modular. The power supply pulls out as a module and you can get behind it. The criteria of Bob Perkins was, if you disassemble something that works, and need to fix something behind it, you've got two parts you've got to make work again and that was a philosophy we carried through our peripherals. There was an extra cost but remember we were servicing the peripherals.

These products were not initially designed for the OEM market. Everything we designed was initially for our own computer people who were servicing it. Therefore the extra effort to make it easier to service, with less skill, was deemed desirable. Like the power supply, we had a unit on slides, so you could pull the whole power supply out, take the connectors off, and then put it back in again. And that held true for most of our peripheral designs.

Burniece: Let's shift over to the development of 808 disk file. Now early on you were doing the 606 tape and not too long after that you were requested to also do a disk drive. Tell us about that and why you needed to design a disk drive, what you had available at the time, and what you did.

Thorndyke: Well about that time, which was, probably about 1964, Seymour was designing the 6600, and he decided that he needed higher performance IO and reduced access times than tape as his first level of storage media. So he initially bought the Bryant disk file but it used large disks [and did not meet the 6600 performance or reliability requirements]. I think they were about 30 inches in diameter or more, with an angular positioner. So it was a heavy impulse moving all that mass into those disks. As such, it used to be a game for some people to program it so that they could make it walk. The disk file would vibrate and they could walk along the floor by hitting it at its natural frequency. Of course, a program was written to prevent that from happening. But when you impulse a large mechanical deflection into a structure, the structure rings. It vibrates and rings like a bell. So when positioning, the structure rings, and the heads are moved on and off track. Until that ringing dies out, you can't read or write. You can try it but it's not going to work. As such, that imposed a delay and Seymour required 10 million bits a second, through the 6600, which was significant. The 6600 was a CPU with eight peripheral processers that did the IO work and then locked it into memory so the central processer could work on it. As such, that was the requirement, so we had to get up to 10 million bits a second and disk was the only way you could really do it. Two-inch tape would do it but it was never considered for that. So we started trying to figure out how to build a disk file that solved this.

We came to a realization that the large force to move a head assembly weighing ten pounds required a lot of torque and the only way to get the torque we needed was to use hydraulics. So we ended up putting four hydraulic cylinders in a square with two pushing one way and two pushing the opposite way. And then we arranged these to access four stacks of eight 26" disks, mounted on two parallel double-ended motor shafts, so the head actuators would always position against each other. As such, we would have no reactive force in the frame because they were pushing the same amount of mass at the same time in opposite directions. Then we decided that since access time was important, we would accelerate until we got halfway to where we were going and then decelerate to a stop. So we accelerated at 10 Gs until we were halfway and then decelerated at 10 Gs. Because there was no impulse, we could hold 250 micro inches positioning accuracy for read/write with no ringing delay.

As a result we ended up with four stacks of fixed disks, two upper and two lower, on two double shafted motors. We also had two dual-opposing positioners, one upper and one lower. As a result, we could simultaneously use one positioner for reading and the other positioner for writing on the same motor. One of the problems back there in the early '60s was computers didn't have a lot of memory, since we didn't have semiconductor memories yet. It was still all core memories. Since the disks on a double shaft motor were synchronized, we didn't over-run buffers, when streaming data in and out and of the super computer processors at 10 million bits per second. The system could wiggle a few bits in the data and store it. Then you could bring it back and you just ping pong it back and forth continually, because of the nature of this design.

Nowadays you get synchronized spindles all the time and everybody thinks it's great. But we did that in 1965, not realizing that was a significant advantage for the 6600. So synchronizing the disks on a double shaft motor, and then transferring data at 10 million bits a second effectively gave Seymour what he needed for the 6600.

Burniece: Now Lloyd on that point, you got the high transfer rate by paralleling the heads right?

Thorndyke: Yes, we transferred data at 10 megabits a second by reading or writing on 12 tracks at once in parallel. Each head had six read / write channels so there were six upper surface tracks and six lower surface tracks per head arm. There were 32 heads per actuator and 32 discrete major positions. The problem was we couldn't crowd the head channels too close together, so we spaced them three tracks apart and had a reactive, proportional control servo system to incrementally move plus or minus one track spacing on either side of each major position for the other two tracks. As a result, you could read and write 1/32 of the total file with each major positioning move. [1,2]

It was interesting how much performance we got out of the 808, because some comments were made in hindsight that other companies failed to be competitive in the super computer business. They failed to be competitive because they did not have a disk file reading and writing at that sustained data rate. I put counters on the 808 and found that the 6600 was moving the heads about 19,000 moves an hour. We also formatted the disk, into sectors, not like IBM where you had to read the count key to find out what the data is. We had sectors, so therefore the software could look around and make a minor move to pick up another record, then move back and catch the one you've been working. You didn't have to wait for a full revolution and that was a significant advantage. We could to that because it didn't ring, so there was no settling time. In fact early on, we decided that we had to servo on the sectors, in order to position. So we effectively ended up in 1965 servoing on sectors, but then decided it wasn't worth the expense, because the positioner was accurate enough to position within 250 micro inches. Thermally the inside of our casting was thin and was also the air return. As a result, we could control the expansion of the aluminum casting, disks, and positioners, so that we did not have differential expansion problems as it warmed up.

Burniece: If you were servoing on sectors in 1965, you probably were the very first in the world ever to do that.

Thorndyke: Yes, but you didn't need to do it. We had headers on the sectors. You could switch the sectors, which were about 1k bytes, so there were a lot of sectors.

Burniece: Lloyd tell us about a couple of the key guys who were involved in that 808 disk file design. Who were the guys that made it happen?

Thorndyke: Well the mechanical engineer that really was unheralded was Dan Sullivan. Sullivan had a hydraulics background and had also been one of the head technology engineers on the tape transport. He had production experience in a hydraulics company and had a lot of casting experience, so we could go to sand castings, etc, when designing it. Then there were Dale Larson and Doug Hennefent, who were the two people on the head. We'll get into Hennefent later on but Dale Larson designed the heads back at Univac for the drum memories they used as their storage element. Now the drum memory was bit addressable, so effectively it was accessed like core memory. As such that required you to have a drum that was perfectly round, with no eccentricities in the drum, because the heads would be flying at about one to one and a half mils (0.0015 inch) from the drum surface. So you had to make damn sure that your thermal expansions between the rotating drum and the head shroud were about the same; otherwise the head would dig in.

I joined Univac in 1954, and heard the story that IBM came into Univac in probably 1952 - 1953 to buy a license to produce drums. Bill Norris, of course, was interested. He thought he was going to be able to sell IBM drums and I think they ordered a number plus a license to build. I wasn't there at the time, but the story goes, that IBM then come back and said, well we got this drum but how the hell do you address it? What is the electronics? Well the guys said, "Here is how you use a drum in a computer." So they spent some time to help IBM out, because they were buying a license on the drum. I think IBM introduced that as the 650 [Magnetic Drum Calculator], which I believe was their first computer. This sequence was repeated many times in the CDC OEM disk business.

So the comment from some of the early guys was "Yes, we taught the IBM guys how to address a drum and how to use it," which is good as far as they were concerned, because they were selling drums and making some money. We found out that in the OEM business, you had to get in bed with the customer and help him in the design of the controllers and software, or he was not going to buy it from you. Incidentally, we also found that in the OEM business you had to sell the product at half of the market price, so the OEM customer can make money on it. If you got a manufacturing cost like some of the people who sold the peripheral equipment, at three quarters of the market price, you would lose money on it. But if you can get the cost down to half the market price, you can make money. In our disk business, we were sometimes selling our disk products at 40 percent of CDC System's market price because of our learning curve. Well now the customer is out there starting to go crazy and you got a willing OEM customer. That is why the IBM peripherals never caught on with the OEMs, because they would not sell at a discount price. As a result, I believe that because CDC was selling disks the way we sold them, we were very responsible for a series of companies that successfully competed against IBM. Without the disk they could not compete against IBM and no one had developed an OEM disk product except Bryant at that time.

Burniece: Let's talk a little bit more about what happened next. But before we do that, the 606 tape transport did generate a lot of OEM business and actually you built a factory around it and everything else. The 808 disk file was mostly a Control Data Systems device but you did sell a few OEM right?

Thorndyke: Yes, we sold a few OEM.

Burniece: Who did you sell them to?

Thorndyke: We sold the 808 to ICL - English Electric, who had a unique name for it " Big Willy." Then Honeywell bought it for their computers, except they needed a bit serial version. They couldn't handle the 12-bit parallel, so we put a bit serial interface on it [and called it the 803/804 [3]]. They used it as a big storage device rather than for online memory swapping. That was about it, because we sold it OEM for over $200,000. As such you had to have a hell of a big computer system. We sold it for $400,000 in Control Data's product line. So unless you have a high performance computer, you couldn't afford it.

Burniece: There are a couple of interesting things with that 808, I'd like you to talk about a little bit. I know that cleaning it was a problem and determining when you had to clean it was also a problem. You solved that in a unique way. Tell us about that.

Thorndyke: Well part of that happened when I was moved from Normandale to CDC Research to do research on advanced peripherals, mainly for the Star Computer. The Star Computer was a unique computer that CDC was pioneering that used vectors. A vector in a computer is a stream of consecutive data points. So you have this stream of 1000 data points that you want to add to another set of 1000 data points and then store the answer. Instead of addressing each pair of data points individually, they wanted the vector processing unit to simply say, stream 1000 addresses of data from A to B, add them, and store in C. That way you could stream everything out and get a performance that was orders of magnitude faster, with just a few instructions. They also wanted to stream a double bit word and do two of them at once. I think Jack Walton of Los Alamos had originally proposed that you could get much higher performance if we would end up streaming the data in a peripheral in the vector processor. So as such, there was an advanced lab set up to develop streaming peripherals for the Star Computer.

We were over at Research for a short time, when a new general manager came into the Normandale manufacturing plant and he was having trouble making the 26 inch disks. He told engineering that they had to relax the flatness spec of the disks, because of manufacturing yield problems. The flatness required was 10 micro-inches per inch on the disks and with anything more than that the heads would not fly over it, because of the mass of the head. Otherwise it would be a washboard and you can't fly on a washboard, since you will hit the tops of the peaks. The heads have to follow the terrain and the only way you can follow it is to make sure that the waviness is so low, it is not near the natural frequency of the head. So they changed the flatness, and the next thing that happened was we got called over to Mr. Norris' office, where he reported that over half the disk files in the field on the computer system were off lease. Well, you know, that's a big pile of money. So Norris said, "You've got to go back to Normandale, find out what's the matter and straighten the problem out." He also said, "You have my signature to manage that program" which meant now, I have two levels of authority over the plant. We established that the relaxed flatness spec was a severe problem, and re-established the original engineering spec. That solved one of the problems.

The other was contamination on the disk and heads. We noticed that in manufacturing, certain technicians ended up producing good files. Other technicians' files always had trouble. So we went around and asked "What's going on?" and one of the guys said, well he checked them out and could tell when "Uh oh, I got a head that's not flying." We asked "and how do you know?" and he said "I can hear it sing." Well I'm pretty deaf that way, since I don't have any hearing over 2,000 cycles, due to the fact I did too much shooting when I was young. So I said "Well, you can hear it?" And he says, "Yes, and when it sings, I go ahead and clean it till it doesn't sing." So I got one of the engineers, Ferd Samuels, and I said, "bring the tape recorder and let's find out." So he brought the tape recorder in and I told the technician "when it's singing let us know." So we marked the tape, put it on a frequency analyzer and found out it was vibrating at 9,000 Hz. The head operated like a tuning fork at that frequency. So I said, "Ferd can you make a slot filter for me?" He said, "Sure I can." I said "Well make one for 8,000 to 10,000 Hz and let's put a microphone in and listen to it." Well he did that and a couple days later, put the microphone in and by God, a head was singing to beat hell. So I said "Clean it." When we finally got it cleaned and the singing stopped, the technician in the lab said for the first time "That will tell me if I'm doing it right." He said, "Up to now, we don't know if we're cleaning the disk and head well enough." So now we had now a device that would tell us that the heads were having trouble flying, due to the fact that the contamination built over time, and number two, we also had something that told us, when we had done the cleaning job well.

So we bought four microphones, put them in the four stacks and then put counters on all four, plus a detection system that determined if the noise extended for a period of time and would then inform the system that we were going to shut the file down and check heads. That allowed the technicians in the field to pace their work. Everybody wanted to call it Crash Detector and I said "We can't call it Crash Detector. My God, our disks don't crash. You want me to tell you I got a Crash Detector to prevent crashes?" I said, "No we'll call it the "Head Altitude Monitor", it's the HAM kit. So that went into the field and I don't think we had any more catastrophic head crash failures. I'm sure we had some head crashes, where one of the heads wouldn't fly. We also found out that if you went up to the shroud of the disk and simply tapped it with a coin, the tick detector would detect that frequency. So every morning a service tech goes around and he taps all four shrouds and says "our HAMs are working. I'm in great shape." We also put a mechanical counter in so that we could count the number of ticks we had each day and when that got up to a few a day, we knew that we had to do a cleaning because something was wrong. The other thing that we did was, we originally specified that after every week to week and a half, you were supposed to clean the disks. Well the data showed that just after you cleaned the disk, the ticks were most prevalent and then over time, they went down an exponential curve. So we said "let her run, don't touch it until the mechanical counter gets up." God, that was a savior because we were doing more damage to it every time we cleaned it. We also found that we had trouble with crashes at one of the customer sites and couldn't understand why the hell they're having crashes. So we went out there and sat around watching the disk. Well the janitor came in at night with his dust mop. He went over to the disk file, where we had four horse-power worth of fans for temperature control, and he used the fans to clean his dust mop.

<laugher>

Thorndyke: He says, "I just shake it here and clean it." Of course we had a dust storm go through the file, so we put up a notice that said "this is not to be used for dust mop delousing."

We also had a problem with water cooling at a Honeywell site. We cooled the hydraulics with water and if we detected we lost water pressure, we would shut the file down. The problem was we used a millisecond transition shift and we were on the end of the line that flushed the toilets. Well as soon as they flushed the toilets the water pressure went down, so the detector fired and we shut down.

<Laughter>

Thorndyke: About that time the pressure would come back up and we would be in the center of a retraction where it's shaking heads at 25Gs and having problems. So I said "Well, we can't redesign it to put in a thermal delay relay, so put a big capacitor across the relay, isolate with a diode to prevent discharge, so that once it picks, that relay is going to stay picked for at least a second." We got that all done and started sending the change to the field, when the QA people came in and pointed out we were using an unauthorized capacitor. I said, "We buy them by the gross for use all over CDC. You haven't had that one tested and you're going to lock us up with a stop ship?" So I said "Well that's interesting. Let's go to Mr. Norris and I'll point out that what I have is a life preserver that works and there is a guy out there drowning but you don't want me to throw it to him. Instead you want to use the swimming pool for a couple of months, to qualify the capacitor." I said, "We're going to send it to the field and if this solves a problem, then it's certified." He said, "I'm not going to allow it." And I said, "Well I don't particularly care if you're going to allow it. That's what we're going to do and why don't you go with me to Norris and I'll point out that this is going to save us about $150,000 rework every time, so we should keep it in until this capacitor is qualified. Otherwise, the computer is going off lease." I said, "How the hell you think your salary is going to get paid by having all these computers off lease? It's just a capacitor and if it doesn't work, we are no worse off than we are now. Otherwise they are going to crash. But if it does work, we're going to eliminate that problem." Well that was the end of that; you can get by with that and I did. But later on, a year and a half later, he had an opportunity to stop one of my programs for no good reason, other than the fact that he remembered that, so I had burned a bridge.

Burniece: Pay back time.

Thorndyke: I burned the bridge and remembered not to tell him, thank you. So anyway, that's the background of the 808 disk file. Then after going back in research, in about 1971 …

Burniece: Let's go back for a minute. We'll get to 1971 later. About the same time that you were developing the 808 disk file, IBM developed a removable pack disk drive called the 1311. You have an interesting story about how you decided you needed something like that for CDC. So tell us about how you ended up with the 852 disk drive.

Thorndyke: Yes. We were always alert to what was going on in the world. Within Control Data, the philosophy was you always built the biggest thing you could build and then start spinning it down into other markets. For example, after the 6600 was built, there was a 6500, 6400, 6300, 6200, and 6100 that came out of it. When they developed the first computer, [the 160 / 1604], they developed a whole series of logic cards for that computer. So when I did the 606 tape transport, rather than develop all electronics, I went up to Arden Hills and got all of their standard card types and used that in the tape transport. Well in a tape transport, what do you use to drive the magnetic head? You need a card that switches the core, so I found a standard part that had the power to drive the magnetic head and was about $5 a channel, so I didn't have to develop anything. That was one of the reasons we shortened up the 606 development time to less than 18 months.

So then we started looking at what was the follow on to the 808 disk file? We said, "Well you can't build one bigger and you certainly can't take the current big file and cut it in half or anything and save a hell of a lot of money." So IBM announced the 1311 and Dan Sullivan and some of the rest of us looked at that and said, "Why don't we buy an IBM machine to understand what that unit is and spin our technology down into it? We know how to build disks and we're building flying heads. Why don't we spin those down?" The question was "Where are we going to get the money?" I said, "Well you know, we're spending $4,000 a day and most of the escape technology is kind of coming to an end. We can terminate a few of those and all we've got to do is save $1,000 for 17 days, so we can buy one." So we pruned it down to $3,000 a day to buy an IBM 1311 or at least tried to buy one. At that time, it required a general manager's signature and Tom Kamp looked at it and said "Well, what do you want to do that for?" I said "Control Data has a very big product line of 3300, 3200, 3100 computers on the business side of the house. IBM announced the 1311 for their medium business computers, where the customers were not going to be able afford a big file. [The 1311 allows you to store data on a removable pack and even move it to a remote site, so it was also the replacement for tape]. It's the shelf store of disk, not tape, that is going to be the answer for short term storage - long term maybe tape, but not short term. So what we want to do is, to buy a unit, understand the characteristics and cost, and then design a pack compatible version for our computers. That way we can swap packs, at least on our computers." As it turned out we only swapped packs within CDC, due to the way CDC formatted it. We were never pack compatible with IBM.

Anyway, Kamp said, "I'll never sign that. There's no market for removable disks." I said, "Well the market isn't there yet." He said, "Well nevertheless you're not going to get my signature because there's no OEM product." I said "This is for Control Data and since Arden Hills is the way we are funded, the system business pays our R&D to develop peripherals. Therefore all I have to do is hook it on to a CDC computer and it will be funded." He said "I'm still not going to sign it." I said "Well then I'm going to Mr. Norris." He said "Okay." So I went over to Mr. Norris and walked in there and he knew me and said, "Well what are you here for Lloyd?" I said, "Well Mr. Norris, we got all this technology from this big disk file and I want to spin it down into a smaller disk drive for the 3300 because IBM has managed to place disks on all of their business computers, and we're going to have to do the same." And he said, "$17,000?" I said "Yes, but we feel that we can probably build it for less than $5,000." He said "A hell of an idea" and I got his WCN. So now I've got this approved and we bought it. But Mr. Perkins and Paul Bulver, the chief engineer, felt that the 808 disk file was at a critical stage in design, so we couldn't divert the resources to get into that removable device. As a result, they assigned it to another crew, which had no disk experience. This crew had many thousands of change orders as they learned how to design a high-volume disk.

The interesting part of the 808 was that we all had disk experience at Univac. In fact one of our understandings was that I wasn't going to work on disks, because we'd been through a couple of disk designs at Univac and I said, never again. So when Mr. Perkins came over and said, "Well we're going to do a disk design", I said "Remember we said we…", and he said "Yes, that's fine. You can sit over there and play with the tape unit. But that's going to go out of business. So if you want a job you're probably going to have to run the disk program. That was when I decided, by God, I can design a disk."

<Laughter>

Thorndyke: Besides that, the disk drive had the budget, so Sullivan and I decided yes, we could work on a disk too, as long as we needed jobs. So that was the start of the CDC disk group. We pulled together, all of the different people who had disk or drum experience at Univac, and that became our critical crew.

When Mr. Perkins brought us the draft of positioning issues that our servos would encounter, we determined that we would push weights against each other, in order to get rid of the mechanical ringing. That concept actually came from Jack Rabinow.

Jack was the head of the Bureau of Standards and had developed an automatic character recognition capability. CDC bought the technology and formed a company over time, called Rabinow Engineering. I would go out and visit Jack, who had a museum, with a lot of things that he had tried one way or another. One of them was for the Navy, where he addressed one of the problems they had with fire control radar, which dithered and bobbed back and forth, with an antenna that was massive. The problem was when it stopped, the whole structure was ringing so you had to wait until the ringing was done. So he put an equal mass behind it and shoved the dithering units against a mass. In doing that, there was no ringing, since you weren't deflecting the structure. So we looked at that and it became the background of our 808 file, in positioning two devices against each other reactively so the structure received no impulse at all. In fact the 808 produced 200 pounds of force, yet you could stand a nickel on a positioner all day long and it wouldn't fall over. It was tuned that well. So anyway, Rabinow Engineering was an interesting place. Jack had a lot of inventions. When I knew him, he had 250 inventions and was the only guy in Control Data who had his own patent attorney.

<Laughter>

Thorndyke: Joe Genovese was his patent attorney and Jack got every patent issued. He would say "Hey Joe, they missed this claim on a patent." So Joe would put another claim on the patent and a lot of his were additions to our existing patents. One of Jack's claims to fame was adjusting the clock timing in a car, which was a problem, because [automobile] technicians could increase the speed or decrease the speed but needed to go through three or four iterations before they got the timing reasonably close. He designed a 1000:1 gear ratio that tightened or loosened the timing spring . So as you tighten it, it would tighten up the timing wheel main spring at 1000 to one. So incrementally, as you adjusted the clock you finally got it perfect. Well he got a penny a piece for the royalty as an award and a penny a piece was a lot of money in those days. So Jack was very inventive but he was a model shop type designer. He didn't draw prints, just sketches, as he had some very, very skilled model makers. So he had them build the things and then document it.

As an example, he invested in a particle clutch or particle brake, where you had two surfaces. You put in a fluid with a lot of magnetic filings in it and when you put a force on it, you either had a brake or you had a clutch. We used that on our tape transport for a while, until we found out it had other problems and wouldn't release properly. So we finally had to put a friction brake in and then we put a shield on it, which sucked all of the dust back into the dust control system, in order to keep it from getting into the tape.

Jack was very good but he did not have a good understanding of electrical. I remember when he brought his people to Minnesota. They were highly educated and looked at that big file and Jack said "What did you invent?" I said "probably nothing" and he said "Well then why did you do it?" I said "What do you mean why did we do it? We needed the storage for the 6600." He said "Yeah but why did you do the file? You didn't invent anything."

<Laughter>.

Thorndyke: I got quite a charge out of that because they didn't understand that you need to make money in this business. They invented a page reader that would read five different typewriter fonts at 100 pages a minute. The page reader system was the size of a modern computer system, and the optical system had technology that used analog information or analog reading, in that as we could take a character and blow it up. Then, where the pattern was black you got a one bit and white you got a zero bit. One of the problems was an E and an F were quite similar. The E has a line at the bottom and the F doesn't. What it read was that every time it's black you got a unit of current and where it was white, you got a negative unit of current. Jack set the probability of a bit at about 1000, so if he got a word that is about 1000 bits long, he would take this 1000 bits and produce an analog voltage of 12 volts as a perfect read. Anything less than 12 volts down to maybe for a 10 volts was pretty positive. So you could say that 10 volts was really an E on this typewriter. But different typewriter fonts had different characteristics of E.

They had transients and couldn't always read it, since there was all sorts of electrical noise. I went there and looked at the reader and it had a bundle of 22 gauge wire about as thick as your arm. I think there were something like 100,000 wires in it and they were shoving 100 amp pulses down the cable. It created so much cross-talk noise they didn't know what the hell was going on. I looked at the ground line and it was jumping three and a half volts with a five-volt detection threshold.

<Laughter>

Thorndyke: So I went out and got a welding cable and clamped the two cabinets together with the welding cable and they thought I was out of my mind. But for the first time they had a good ground reference and it started working. They just didn't have the practical experience in the electrical area. They had never suffered through this sort of grounding problem.

As a sidelight, one of the problems I had at Univac was we had a card-to-tape reader or a card-to-tape converter that we sold with the Univac One to Park Davis back in probably about 1955. Every Monday morning at about 9:45, the tape unit would run away. They would have to re-do all of their card-to-tape work on Monday. So every Sunday night I would go out to Park Davis and be out there on Monday. After I got done, I usually went down to the coffee shop and had coffee with the nurses. I said "What in heck goes on here in the place at about 9:45 in the morning? What is going on in here?" One of the nurses said, "Oh that's when we turn the x-ray machines on and start x-raying people." The computer room was in back of the medical facilities near the X-ray machines. So I went out and bought a big isolation transformer and hooked it on to our card-to-tape system and then wrapped aluminum foil on the walls. That was the last time I went to Park Davis but it taught me that EMI comes through the line and also through the air. Later all of our computer rooms were lined with aluminum foil, because the search rays of a radar at the airport (about 1 ½ miles away) started erasing our magnetic drums. It would come through the air and hit every wire, which was an antenna at the right frequency, erasing the drum track. But anyway that was a good lesson.

Thorndyke: Going back to the 852, this became the first CDC removable disk pack device [4,5]. After we got it into production, the OEMs came to the realization that without a removable disk, they were not going to sell against IBM and there was no other removable disk drive around. Now there were two or three other companies that were quoting removable disk development that would take $3 million and two or three years. So Control Data was approached on whether we would sell the 852 OEM. And the answer was, “sure” but the Arden Hills people were all bent out of shape, since they paid for the 852 development. They were not very anxious for us to start selling OEM. But Norris was a driving force behind selling OEM. His point was, “You got this developed. Why don't you sell to other people and reduce the cost?” He said “The problem years ago, was the printer people would not sell a high performance printer, so they ended up having people get in competition with them and they finally ended up not selling at all, because the competition developed a better printer.” So he believed that you should sell to your system competition. As such we called each a different number in the 9000 series, with the requirement that every customer had to specify a non-CDC interface that became unique to them.

So we had customers that needed the 852, and we were already in production. I used to joke about the marketing people who would lose sales if they were not manning the telephone over lunch. We had customers call us up and say, “I need that disk. Can you get me one to work with so that I can design my interface?” So we had production scheduled with spare disk drives available that we could send out to the OEMs, so they could develop their controller. About four months after the controller was developed, we started getting orders. We ended up upgrading the 852 to the 853, [with a higher transfer rate, and the 854 was also double the track density, so the capacity went up to 6.1 megabytes [5]]. Anyway, that became the real OEM seller and went against the IBM 2311. It was IBM pack compatible but was never IBM interface compatible [4,5]

About that time, Tom Kamp appointed me as General Manager of the Disk Drive Division and we built a plant in the northwest side of Minneapolis for the sole purpose of building rotating memories in an environmentally controlled, contamination-free environment, with individual air-conditioning in each room with filters. We never moved into it, however, because Control Data started suffering its first downturn in earnings in 1964. The corporation thought they made three and a half dollars a share but by the time they eliminated interdivisional transfers, it was losing about a buck and a half a share. We had transferred peripherals to the computer division at a cost higher than the standard cost, but it was a fictitious markup. Once you eliminated the interdivisional transfers and just looked at the cost, Control Data lost a buck and a half a share. So we closed plants and other things. But it also taught us that you don't transfer around a corporation at a profit and make money. If you remember, a company sold some systems to their end-user division at a profit and were telling the stockholders they made money but it never left the corporation. Of course, the IRS corrected them. So we learned early <Laughs> that you can't do that. Anyway the 854 became the big seller and, as I recall, we had 87 different customers.

Burniece: Eighty-seven different OEM companies?

Thorndyke: Yes, 87 different OEM customers as I remember. The thing that we did was we created competition for IBM, because without our disks IBM would have dominated that market and no one could have touched them. The fact was we had the disk drives in production and had already paid for the engineering, when we sold OEM. We never priced them against the cost for the development because that was already paid and we had no sense of recovering it. So what we did with OEMs is we priced it on a learning curve at about double the cost and as the cost went down, when more OEMs bought, we’d go back to the learning curve and lower the price. As a result, once they got on the learning curve, they had to come back for the next order, because they got a decrease in price. They were going down the learning curve a ways with all the total volume and, as such, they had to stay with us. In the OEM market, we had many OEM customers for 20 years, as their sole source. It turned out that we had the dollars to develop the new technology, which was always based upon the super computer needs first, and we were developing it for that and then finding the OEM market.

Burniece: So in the middle to late ‘60s, you were selling the 852-854 family to lots of OEMs [6]. In your opinion is that when the OEM business really took off for CDC? Was that the impetus of CDC’s OEM disk business?

Thorndyke: Yes, I think so. In 1968-1969, we had so many customers that the manufacturing VP, who was kind of running everything, had the engineering people stop all advanced development and had the engineers service the OEMs. Well that was fine, except the 854 ran out of gas and there were more advanced OEM disk products out there. All of a sudden we didn’t have an answer because there were no engineers working on it. So in about 1970 Normandale was losing money and we were losing customers. Our product was obsolete and there were better products out there. I was over in the CDC Advanced Projects Lab and John Titsworth had come in from Lear as the general manager of the disk drive division. John was continually trying to get me to come back to Normandale and I said, “I won’t come back as long as the VP of manufacturing is there, because he’s the guy that created the problem and I’ll be damned if I will.” He was the man that created the problem that the 808 had and now he’s created the problem that there were no engineers working on advanced development. As such, I told John that I wouldn’t go unless he was removed and he was John’s boss. So finally I got called over to Norris and he said, “Normandale is losing some money and I’m told you are the guy who can fix it.” I said “I have a problem Mr. Norris.” He said, “I know about it. I removed that man. He reports to me. Now get the hell over there and get the place straightened out.”

So it was a series of unique errors in that the new manufacturing VP seemed to feel that none of the engineers could be trusted. If you wanted a pencil, you had to send the requisition to him to get something done, which is not the way to treat engineers. I moved in immediately, and found that there were some engineers who came late and left early. We had programs that were critical and they didn't open up the print shop until 9:00, so there was no sense getting there early, and they would close at about 3:00.” I said “why the hell wouldn’t you work till 5:00?” They said “Well the manufacturing VP did that.” So I said “Well, the print room will be opened at 7 o’clock and it will stay open till about 5:30.” The manufacturing VP came up and said, “No you can't do that.” I said “I don't care. It’s going to stay open and the same thing on the doors.” Five o’clock they wanted to lock the doors. I said “The doors with the guards will stay open till 6:00/6:30.” He said “You can't do that.” I said “I can do it. I got a key to this place.” So it was a test of strength on whether engineering was going to manage engineering or a manufacturing person was going to manage engineering. Well, we got engineering back under control and made it convenient to work.

Then I came to the realization that the engineers were ordering parts delivered to their home and bringing it in to work, while trying to submit a payment request for a device that never came through the proper area. I said, “Why are you doing that?” They said “Well if you send it to manufacturing they’ll only work on our parts when they have nothing else to do. So it could sit in manufacturing receiving and inspection for two weeks.” I said “Well it will sit in receiving and inspection for one day from now on, because they don't need to inspect. All they need is to send it to us. There is no inspection required. Just send it to us and we’ll send you a note, if it was not the right unit.” Then we got our own purchasing agent and we started operating like an independent business. Of course, the engineers I brought from research were disk people, so we started looking at the product line and came to the realization that the product line they were developing didn't make any sense in a lot of areas. So we reviewed the programs and changed several of them and that is where the SMD came out.

Burniece: All right, so let's get into that. I'd like to have you focus on SMD for ten minutes or so and then we're going to get Dag to take over on the computer side.

Thorndyke: Okay.

Burniece: It's 1971. You've come back to Normandale and I was there now. You start straightening the place out and now you're trying to figure out what to do next. There was a program called the Memory Module that was being developed in the backroom and tell us what happened next.

Thorndyke: Well, the Memory Module was a device that was supposed to be [like the IBM Winchester]. But IBM had decided to put everything on fixed media and our belief was that the OEM market didn't see a need for it and didn't like the idea, and of course, the computer division had no interest in the product. They still liked the idea of low cost, removable media. The IBM 3330 had outgrown part of our OEM market, which was still centered in the 10 to 40 megabyte range, especially for mini computers. So there was a big argument going on and I think it was Dan Sullivan and possibly Dale Larson who came to the conclusion that the IBM type Memory Module was not the right OEM product. We also had Tom Dugdale, who was an electrical engineer that went into marketing. Tom was actually opposed to doing a Winchester type device. So we started looking at it and he said "Well, the problem is the IBM 3330 pack is too damn much data for the minicomputer market." So we peeled it down to 5 disks instead of 11, reducing the capacity to 40 Megabytes, instead of 100 megabytes. As a result, it was a miniature 3330 IBM removable pack. The argument was we had already developed an IBM 3330 equivalent drive [the HPD], so we don't have to try to figure out the electronics and everything else. But at the same time, we didn't need to be IBM compatible and the head people, Harold Beecroft and Doug Hennenfent, had developed a head that flew at an angle rather than flying flat. By flying at an angle, there was a very small area that was in contact with the disk and we could get down to a flying height in the order of 20 micro inches, instead of using a big pad where any piece of debris that came through could tip the pad. As such, it could effectively move the debris over to the side and that allowed us to get closer for a greater bit density. So that was the combination that we put into the SMD. We also realized the HPD technology, combined with the SMD technology, could be used to obsolete the 808 disk. To that end we configured a 20-disk pack that was equivalent to 4 SMD technology packs. This was sold at around $75,000 against $400,000 for the 808 and space-wise, was the size of an HPD disk drive. We called it the CDC 819 and enjoyed a big market with Cray, Inc.

Also, we had previously, about three or four years prior to that, been criticized by the corporate officers of producing devices that were noisy. Their comments were that we were producing noisy devices and it was a health hazard, so we needed to quiet them up. Well our comments were we didn't know how the hell to make them quiet. But we had more money than the ability to spend, at that time, so I sent two engineers to the university to learn about acoustic noise and how to suppress it. Well they spent about nine months over there and came back and started working on our noisy tape transport. The result was you almost could not hear it run, with no loss of performance or anything. What they did was they added an orifice to release the vacuum slowly. When you break a vacuum you get a sonic noise out of it but with an orifice releasing the vacuum, as the tape rotated, you couldn't hear it. Then they suppressed the fan noise by putting reflectors in to reflect the fan noise back into the fan. You suppress the noise by feeding it back at a half wave length delay. You remember Yeager, when he flew around the world in his aircraft, there was so much aircraft noise that he couldn't talk? What they did was they put in a microphone and delayed it a half wave length and fed it into his earphones, so it was quiet.

By applying this approach to our products, we ended up with a disk drive that was so quiet you could put it in an operating room in a hospital; it was that quiet. That created a hell of a lot of problems in the field for other people because they had noisy disk drives and the Control Data disk drives were so quiet you could hardly hear them. That was part of the effort that we did with the SMD. We put a lot of technology in it and it really was a very successful product. It had a short stack IBM 3330-type removable pack for shelf store but we increased the bit density for more capacity, so it had to fly closer. We didn't have to be compatible with anybody but ourselves. The fixed disk project was initially called a Memory Module but my comment was memory belonged to computers and this was storage, so it became the Storage Module Drive. There were a lot of objections to it but, since I was running the place, I could call it a Storage Module Drive (SMD) and get away with it. So that is what it was.

Burniece: Thank you very much Lloyd. We actually had a long interview with Tom Murnan and Dick Berreth last year on this and got a lot of the SMD history in that Computer History Museum oral history interview [7]. So let's wrap up the peripheral side of this interview with just a couple of additional questions. What accomplishments are you most proud of in your time in the CDC Peripherals?

Thorndyke: Well you know, my greatest satisfaction was designing products and going out on the production floor and seeing a multitude of people putting the product together, making a living at building a product that is sellable. There is no fun in developing a product that you can't sell. I used to have a rule of thumb that the total development cost of the product had to be on the order of one percent of the total sales value of the product. If you started looking at it and said "I can't get that", maybe we were working on the wrong product, because we've got hundred's of people out there depending upon building products that are going to have volume. So we didn't design products that we couldn't sell in volume. We were not a model shop design house, although a couple were done that way because of unique customers like AT&T. But basically our job was to design products that we could build in volume. We ended up increasing the technology incrementally, so that the people building it didn't have to go through a revolutionary change. They could gradually increase their skills. We didn't change everything and force them to start over

As for the products themselves, they were like your kids. You loved them all. The 852 decision was monumental from Control Data's standpoint. The fact that we were able to force through the development of our own version of the 852, put us ahead of the others in the OEM business forever, in my mind. Now somebody else has it, I can't talk for that, but up until the time that CDC sold the disk drive operation we were the dominant OEM disk drive player

Burniece: Are you talking about selling the disk drive operation to Seagate?

Thorndyke: Yeah, Seagate.

Burniece: that was 1989.

Thorndyke: Wasn't it Shugart?

Burniece: Yes, Al Shugart [was founder and CEO of Seagate and bought the CDC disk drive operation, which was called Imprimis, from CDC in 1989]

Thorndyke: In my estimation being a guy who started with disks at IBM, Shugart was in the wrong part of the product market (possibly due to his IBM employee contract). He was down in the very low end, not in high technology. All we were selling was the high technology disk drives out of Normandale. Once Seagate got ahold of Normandale, they were no longer just the low end, and they leveraged the CDC technology to become a dominant high-end player.

Burniece: They still are by the way. The Seagate operation in Minneapolis was the old Normandale group, until most of them retired, but they are now in Shakopee and it is still their premier development facility. Just one more question for you and then I'm going to hand you over to Dag--

Thorndyke: Well, first let me finish up.

Burniece: Okay.

Thorndyke: July 1, 2010 was the 50th year anniversary of CDC forming the peripherals group, because it was getting into the tape business and then later the disk business. Now the early people like Bulver, Perkins and Sullivan, plus quite a few others, are deceased but there are a few of us left who can still talk. When you're of that age you don't have a truth squad come along and follow you up.

<Laughter>

Thorndyke: But there are a lot of the guys, like Dale Strand, Dale Larson, Doug Hennefent, Larry Matthews, and others, like Wally Edwards, who are still around and can give a lot of insight into some of the detailed products we went through. It may not be all that interesting to a lot of people, I would imagine. As I was saying earlier, the IEEE foundation has been after me to do this type of discussion with them for quite a while and hopefully they will have access to this effort. The head of the IEEE foundation is Dr. Gowan. He was on the board of directors over at ETA and is a very close friend. I'll visit him next week.

Burniece: Well we really appreciate your spending the time with us on the peripheral side today and we definitely are going to try and get to some of the other still living members of that gang. Larry Matthews, Wally Edwards and a number of others that you mentioned still live here. When did you leave CDC Peripherals and go over to the computer side? I am going to turn the rest of the interview over to Dag

Thorndyke: I was pulled over by John Titsworth when he took over the Systems Division in 1973-1974. He pulled me in as kind of a technologist, basically to try to get some discipline in Arden Hills. They all thought they were Seymour Cray class designers and could do their own thing. John pointed out that Seymour did do his own thing but Arden Hills needed more discipline, from what was designed and how it was designed, plus the characteristics. The problem with Arden Hills was every computer was implemented differently. If you had five different computers, there were five different versions of everything. No one used what the other guy used. Now when Seymour settled on the original 6600 design, Jim Thornton, who was one of the designers, came to Arden Hills and spun out the 6400. Then there was a 6500, 6300 and 6200, so all of a sudden Seymour's designs were spun all the way down, not reinvented. They didn't change everything. Later Arden Hills lost that discipline and all of a sudden every computer was different, across the board. Thus, with near every computer being different, that meant software, hardware, experience, spares, training, everything was also different.

So when I went to Arden Hills we started trying to get some discipline. It was hard but since I was talking for Titsworth, who was the manager of all of it, I could get by with it, although I ended up irritating a lot of people. For example, the Star Computer was four years late and was so complicated, that when you fixed one problem, you sometimes created a few other problems. This realization occurred under Neil Lincoln and Tony Vacca long before I arrived. We came to the realization that unless you started using simulation for a design, you couldn't do it. There was no longer a model shop assembly approach. You had to use simulation tools.

So they came to the realization that we had to do the research to develop the simulation ability for our computers but that went on for about five to six years before it began to come to fruition. About that time, I was pulled by Norris again and he said, "I want you to take over [super computer research]." Seymour had moved on and we went through about three general managers in the super computer area. They had tried to build a super computer out of Star, using large scale integration, but that was impossible.

Incidentally, Seymour failed at his Colorado-based Cray Computer because he did not believe in simulation. He did not believe in tools that allowed someone to build his modules. He would build a module but you could probably never duplicate it. His module was about two to three inches square and they were stacked vertically. As such, when I visited him, he had a barrel that was three-quarters full of modules that didn't work, because the components all had different characteristics, so there was no way you could use them. You can't build a computer that way.

So Norris said, "I want you to go over there and either get me in the super computer business or get me out. The way I'm going, I'm putting money in but I'm not getting anything out." I said "I don't have any experience designing super computers." He said "I don't want you to design them. I want you to either get one built or get me out."

Well I went home and started thinking about it. I said, "Well if I get him in, I got a job. If I get him out, I don't have a job. So I'd better figure out how to stay in rather than get out." So I took over as the general manager of it but there was a person between me and the rest of the people. Very clearly that wasn't going to work. I had to get down into understanding what this technology was and why we were having trouble, plus what could be done about it. So that person gracefully left for another job in Control Data and I began to come to the realization that we were promising a performance that could not be met by semiconductors for the next two or three years. Therefore, rather than trying to go at a speed of 13 nanoseconds, if we changed that to 20 nanoseconds, it would be compatible with the old technology of Star. Then we could do just the scalar unit and make it twice the performance of the existing unit and call it the 203, allowing us to prove the technology.

Prior to that, the research people were too conscientious in listening to Arden Hills, whose argument had always been, "We are the only people who build computers. You research people are not allowed to build computers." Well, the super computer was in the research area. Now CDC Research developed large scale semi-conductor integration back in about 1973 - 1974. Thus, the 6400 was a big computer that had big bays and everything else. What research did is they built a foot-square version of a 6400 processor using large scale, custom, semi-conductor integration that had five times the performance. However, in doing that, they listened to Arden Hills too much and left one instruction out, since Arden Hills insisted that this instruction could not be put in because then it would then be a super computer. So all of a sudden, they had a 6400 that could go on a desktop, in a foot square with five times the performance at hellaciously less cost, that we could quickly check out. The 6400 was such a complicated computer, it took six months to check out, because you had to tune the wires for delay. The fact that the LSI 6400 computer would not run standard software was excused as a failure because it was not a computer in Arden Hills mind.

So instead of taking that and modifying a couple of chips to make it a true 6400, we could have sold them by the gross because they were so darn fast, low cost, and reliable. But Arden Hills took an axe to it and chopped it up, so no one could use it. They didn't learn from it but it was a vehicle that taught us the fact that you can have large scale integration and [greatly reduce size and cost]. So that was the technology that we put together, which could prove that the 203 star computer could be built on large scale integration and simulation would tell us that it would work. All the instructions were executed in memory and the unit then was built. After we sold the 203, we did the large scale integration of the rest of the STAR into the 205, and that was also sold. We began to come to a realization that the next advance in large scale integration would have 64 gates on a chip. Then the next integration of technology would get a lot more gates than 64. As I remember, it was 20,000 gates on a die, due to the fact that lithography strength had come along. Well we built the 205 out of ECL (Emitter-Coupled Logic), a current-coupled approach. As such, we come to the realization that the computers were getting so big and so hot that we had customers saying, "I'd like to buy another computer, but I don't have any room for it or I don't have the cooling, I don't have the capability."

Tony Vacca, a CDC computer semi-conductor technologist, had been working with CMOS (Complementary Metal Oxide Semiconductor) technology, which is slower than ECL but scales better and was focused on memory. He had come to the realization that CMOS was the way to build a computer, because of the drastically lower power. He realized that every time you do a 2:1 lithography reduction, you get about a four-time speed up. So his thought was the real key is to get into the memory side of semiconductor technology, since memory is 80 percent of all semiconductor use, or used to be. As such, that was where the volume was and where all of the cost reduction and effort was. So therefore, we should start building super computers using CMOS, because that was where the cost reduction was.

We started going down the road to designing a CMOS-based super computer. After talking about it however, we came to the realization that Control Data could not afford another super computer development. So it was decided by a small group, that what we needed to do was spin out of Control

Data. So I went to talk to Control Data to see whether we could form an independent company, owned by CDC, to develop super computers. Being the semi-smart kid I thought I was, I didn't go to Norris, because if I did and he said, "Hell no," I would be dead. So I went to Norm Berg, his right hand man. Norm Berg had Norris' ear, so when I went to him and suggested this, he let it fester and then said "Yes, talk to Norris." So I said to Mr. Norris, "You know, we need to spin it out, both from the standpoint that there's too much competition for R&D money, and with spinning the super computer out, your R&D budget will be in better shape. So let us take the super computer and spin it out." Finally Bob Price called me in and said, "We're going to spin the super computer out into a separate company, do you want to run it?" I said, "Yes," and he said, "Okay, start putting it together."

So I put together the staff we needed and all the time that was going on, no one in the computer division was allowed to know, so it was tough. Here I'm working for Bob Duncan, who was running the computer division and he doesn't know what's going on. He was never told. So I formed the unit out from under Bob Duncan, and we existed for about two months, going through all this organization before it ever came out in the open that it was being pulled out. Of course, it got some people semi-irritated by the fact that we formed the company, and they didn't know about it, but it wasn't my job to tell them because the attitude was "the fewer people who know, the better off we are." From a corporation standpoint it could mean a big plus or minus in the stock market. So there were no leaks, because we had everybody sign a piece of paper that said, "If there's a leak that's traceable to you, you are responsible for the defense of the lawsuit, not Control Data." <laughs> Yes, I thought that was a way to keep it secret.

Burniece: Excuse me, gentlemen. We have only 25 minutes left on the telepresence time, so Dag needs to have his questions asked, so go ahead, Dag.

Spicer: Okay. Thanks for the great answers. You've anticipated probably at least half of my questions.

Thorndyke: Okay.

Spicer: And answered them already, so that's great.

Thorndyke: Excuse me, let me say one thing about CMOS. One of the advantages of CMOS is that the way it was designed, we ended up having delay lines - semiconductor type delay lines, in order to tune the timing of the computer. Where you used to cut wires, we now had a piece of semiconductor that could change a count. Therefore, the design was you stayed on the chip for a clock cycle; you went off the chip for a clock cycle. That was not the way we used to design. You ended up with a clock cycle that was a combination, so instead of buying the very narrow range of performance of the ECL chips, we bought the entire bell-shaped curve of production for CMOS. We used a slow one in a slow computer and a fast one in the fast computer, so we had a 5 to 1 performance range with the same design. The slow one eventually became Piper and came into the market at about a million dollars. The big fast one, called ETA-10, was $25 million and used liquid nitrogen that lowered the resistance in the wires and sped up the semi-conductors and processer. The Piper and ETA-10 computer boards, and all the semi-conductor chips were identical, except for speed.

If there were problems, all you needed to do was slow the clock down. When you did, it started working again. So in the field, if one of the semi-conductors began to get slow, you slowed the clock down. Since it was all DC-coupled in the semi-conductor die, there was only one semi-conductor chip that determined the speed. The rest of them were finished in about a half clock cycle and were just waiting for the slowest one. Anyway, that was a critical piece and why we ended with a bell-shaped curve in getting such a low-cost computer.

Spicer: Okay. I wanted to ask you about the 6600. How big a project was this for CDC?

Thorndyke: Well, Seymour did the 6600 over in Chippewa Falls. He ended up, I think Watson of IBM commented, with 35 designers including the janitors and secretaries. Seymour was such a designer that every other design failed. He tried to take the 1604 and upgrade it into the next performance, and that ended up failing. Then he finally found out what he had to pay attention to and that became the 6600. Now, in the 6600 he changed the logic to a resistor transistor logic, as against the 1604, which was a complicated circuit, I think a Darlington circuit that had two transistors that were coupled together for gain. Previously at NTDS, he'd tried to design a computer using a single transistor for gain, and as such, it was impossible for that unit to be manufactured. He used an SB100 transistor, a surface barrier transistor, but it required a gain of 15. If you found one of those at 15, you hoarded it, because there were certain places where, if it didn't have a gain of 15, the circuit would fail.

Spicer: Right.

Thorndyke: You couldn't send that to the Navy. After Seymour left Univac, I was asked to head the NDTS management presentation of the design and redesign of the [ANS]Q17 computer to the [ANS]Q20. It used a transistor called the 1500.This effort was in the circuit development effort of NTDS. We used an OR inverter circuit, and the problem of the Q17 was Seymour had used two diode drops in his computer, so there was no margin. This design used one diode drop and that really made UNIVAC successful in controlling the cost, so it was a very successful computer. Then, of course, they miniaturized it. But with the CDC 6600, the secretaries were laying out the taping of the transistor cards. So Seymour used all the people and directed everybody. If you started inventing yourself, you got "crazed". <laughs> You were on his list.

Spicer: What was Seymour Cray like to work with?

Thorndyke: He would talk technology but he did not have time for chit-chat crap. So if you got in the meeting and, all of a sudden the meeting got dull, he would be gone.

Spicer: <laughs>

Thorndyke: He would just walk off. You also didn't smoke around him or he'd just walk out. He was friendly to talk with but he was a person who focused on the job. Like all types of development projects that I have been in, when you've solved the problem, that's when you stop and go to the facilities or go eat. But if you're in a problem area and someone suggests lunch, you responded, "I don't have time to eat." The only time you eat is when you're bored. You can't stop and say, "Well, I've got this problem about half isolated, so I can't stop", so therefore you skip lunch. I used to lose about 15 pounds in every development. You can't afford to take the time to eat. Then, of course, afterwards you go out on the marketing show and you pick up 25 pounds. <laughs>

Spicer: <laughs> That's right. What was the intended market for the 6600?

Thorndyke: Well, interestingly, all super computers opened up a new use, however, all super computers that I'm kind of aware of were sold as a cost reduction over the existing machines. You could run existing programs and they would run faster and at a lower cost. So the new machines were mapped into old applications because they had a lot more speed. They were usually about four times faster, when you ran existing programs. Then you would start writing new programs and two years after you got the new super computer, you would have none of the old programs on it. They were back in the old machine again.

As an example, in aerodynamics, you might do a two-dimension analysis on a wing but when you get a new super computer, you go to a third dimension that's eight times the data. As such, the super computer is a new tool, and most of the customers will figure out how to put new programs on the new tools. In the early 50's, the argument was that the Eckert Mauchly machine (the UNIVAC 1) was going to saturate the market with, as I recall, nine machines. That was as big as the market would be. I would like to point out to people that "back there in the 1950s, engineering was a lot of square root work and the Friedan mechanical calculator was in a data center, so you signed up for time on that data center in half hour blocks. You would come in at midnight to get a half hour time to start doing square roots on it." Well, maybe you would get a square root every half second, maybe one every second. Well, if you take that and you said, "Okay, let's assume a second effectively utilized. How many machines does Frieden sell?" So now that defines how many square roots are going to be done. Then, you go to the super computer and it does 1,000 or 10,000 square roots a second. So you divide that into the number and it says "Nine machines will do all the square root work that the data centers are doing," so that's the size of the market. Well, that wasn't the size of the market at all, but at least that was a number.

The problem was the super computer opens up a new market every time, doing different things, and especially when you've got the performance. The wind tunnel was used to define the aerodynamics of flight. Until you got a fast machine, you could not afford the simulation. But if you do have a fast machine, you can design a wing, fly it, and within two hours, know the performance of the wing. Now, every now and then, you take one of them and you go to the wind tunnel and verify it. But the problem with the wind tunnel was you got bad data from it, because in an aircraft there was no constraint on the wind flowing off the wing. In a wind tunnel, you've got the sides of the wind tunnel reflecting it back, so you get extraneous data. So you could put an aircraft in the wind tunnel and get different results than the models predicted.

I remember we had a contract with NASA Ames to build the digital wind tunnel. It was about a $250 million effort to start doing simulation of aircraft design. We got part way down the road and they decided that they couldn't afford it, so they bought a Cray machine to do the first analysis. When we got the ETA-10 running, we ended up with people who had different algorithms for aircraft work. One of them that we were researching with took a delta wing aircraft that they were going to put in production and ran simulation on it. Our simulation said that when you go to land it the left wing was going to stall and you were going to have a circular flow pattern on the wing, so it is going to go over on you. But we did that on our own, just trying to see how our ETA-10 would run. It predicted the fact that the aircraft was going to have an airflow due to the delta wing, and I guess you probably had to cant the nose a little bit to deflect the air on the other side. The simulation indicated a circular flow in the center of the wing and, when the plane was flown, it stalled, so it flipped over. Simulation is where the new computers find their use. It's numerical simulation, where you need to predict the weather and everything else. Numerical simulation is where the super computers do the work.

Spicer: Right. I'd like to ask you about software and the role of software in selling super computers. In the early days, and correct me if I'm wrong, customers got very little software. They were sort of left to fend for themselves. And then later in the '80s, more application packages started to appear.

Thorndyke: That's correct. Early on, there was a relationship with the manufacturers and the AEC, where essentially, we had a partnership. They said, "Send the machine and we'll start developing the software, because your machine's got the right characteristics and it's so fast. So we'll start developing our applications in parallel with you." As such, Livermore was a partner and would say, "If it runs these loops this fast, we'll buy the machine." Yes, we would give them some software with it, but they would come in and start developing their own. We developed the compliers and others, but they wrote their own applications. So that was quite a relationship, because they had come in many times and said, "Well, we need a faster computer." The 7600 did these Livermore loops, which at that time were 24 programs that were representative of the type of work they did. If it ran that just so fast, you sold a computer. So that is what you tuned, running those particular loops, and they were tough. They did a lot of scalar work, as well as a lot of vector work, and so it was tough.

Early on, the work on the software was disastrous in that the people on the west coast were supposed to write software, but they got interested in something else, so the software lagged and didn’t mature. Finally, Seymour took over some of the software.. As a result, Cray got the application piece going well before the west coast CDC did and we suffered from that all the way along. At ETA, the low cost Piper machine would allow people to start developing applications because it was cheap enough. I think we sold it for a million and half dollars. The big machine was designed to run at five nanoseconds, while the Piper would run between 12 and 13 nanoseconds on the clock. We could take the big machine processor, put it in Piper and slow it down, because you could always slow it down, but you just couldn't speed it up.

Spicer: My final question, because we're running out of time, is you wrote a great article in "Frontiers of Super Computing" [9] on the demise of ETA, and you gave some of the reasons for that. Some of them were you started with too many people, failed to consolidate Cyber 200 marketing and expected government support, which never came. Are those still valid reasons and would you add any new reasons?

Thorndyke: Well, part of those are sanitized reasons, obviously. If I were to write it today, it would probably criticize CDC much, much more. We never really got spun out. Control Data still had access to all the technology, but never tried to use it. We had an automated plating line and produced our own printed circuit boards. As an example, the Cyber 200 had about $2 million worth of coaxial cable and connectors to wire it.

Spicer: Wow. <laughs>

Thorndyke: Well, there was lots of wire in it.

Spicer: That's amazing.

Thorndyke: And the coaxial connectors were expensive. They would print a circuit board for an ETA-10, at a manufacturing cost of $15,000. We ended up manufacturing 65 computer circuit boards and could use 96 percent of the boards we made.

Spicer: That's amazing, yes.

Thorndyke: The Printed Circuit board production line was all automated, in that we had a pick-and-place robot, because we ended up having an impedance-controlled board that required you controlled the plating temperature, current plating density and time, as well as continually sampling the solutions and dripping in the replenishing solution as you were plating. We also had developed the ability to plate the printed circuit board at about five times the rate as anybody else plated. We ended up having a very high-performance plating system, and even cooled the plating fluids, because with the current density, it heated the solutions. No one at Arden Hills, or other CDC computer division, was interested in doing anything. All they had to do was come over and use the software, map it in and production would come all the way through. The attitude was "Well, Seymour gets to invent everything, and so we do, too." But Seymour was unique, in that he could handle building a computer in a model shop. Arden Hills people had no ability to build a computer in a model shop and remember all the things necessary.

Spicer: Yes.

Thorndyke: Computers today are getting more and more complicated. As a result, the only people who can afford to design computers today are semiconductor people. We thought we had a big deal having four layers of wire on a semiconductor chip [circa 1985]. My understanding right now, is they have 11 layers of wiring on a semiconductor chip.

Spicer: Wow.

Thorndyke: But look at it. We had 20,000 gates, which is probably about 100,000 to 150,000 transistors. Currently, the technology is 3 billion transistors on a die. The problem is, what the hell do you put on it?

Spicer: Right.

Thorndyke: Because you're connector-limited, you're pin-limited.

Spicer: And there's the parallel programming problem, as well.

Thorndyke: Pardon?

Spicer: There's the parallel programming problem, as well, with multiple cores.

Thorndyke: Oh, yes. And in my estimation, and I don't necessarily talk to a lot of people about it, the real invention today is the cross bar.

Spicer: Oh, yes.

Thorndyke: And nobody talks about their crossbar design.

Spicer: Right.

Thorndyke: But there are two super-computers being funded, one at IBM and one at Cray that have a goal of ten to the sixteenth instructions per second

Spicer: <laughs>

Thorndyke: Now, if you are a floating point person, so you understand that number.

Spicer: Yes.

Thorndyke: Most people don't. And I believe it will have about ten to the fifteenth of rotary store and ten to the twelfth of semiconductor store, with 10,000 processors.

Spicer: <laughs> Yes, amazing.

Spicer: Yes, they said for me to keep asking questions until apparently we just get cut off and that's the end.

Thorndyke: Yes, well, that's okay. We can still talk.

Spicer: <laughs> So let me see, if you looked at the last 40 years, in which you've been involved in computing, what are some of the major trends in super computing that you've noticed changed? How has the market changed, for example?

Thorndyke: Well, early on super computers were extra fast in existing problems. The real change is the numerical simulation, where they're so hellaciously fast today, because a multiplicity of computers can be connected together with a massive memory, and each work within their own domain. Like the first example, I think, is weather. To do the weather prediction, you can divide the earth at the equator. Everything on the north and south of the equator is independent, so you crank away for a while and then you stitch it together, and say, "That's fine. Well, now we take that upper half, let's cut that into quarters and calculate that." Well, that's quarters, but why don't we make it 16 slices, rather than four, and put a computer on each slice, and then stop and stitch it together." So you can begin to see how certain problems can probably be broken down relatively easy that way. There are other problems that don't break down worth a darn that way, so there is a numerical simulation.

When we did the ETA-10, we had it in the memory of a Cyber 205. I had a dedicated Cyber 205. Well, all the instructions worked to the minute or to the second, if you wish, which was the times. They built a next generation of ETA-10 with about 150,000 to 200,000 gates on a die and put it on a printed circuit board, back to back. The point is that they had the chips. They soldered them on the board and the comment was made, "well what do you want to do?" Tony Vacca said "load the Livermore loops." So they loaded the Livermore loops and they executed and the guy said "I've never seen a computer that executed." Well that is what simulation does for you. It reduces a risk. Tony Vacca ended up taking off later and went to Cray.

Spicer: Right.

Thorndyke: So, it is a drastic reduction in the risk to do things.

Spicer: The simulation is useful to design computers, as well as being the principle application of them.

Thorndyke: It's essential to design.

Spicer: Yes.

Thorndyke: Because you make one error on a semiconductor chip, that's $1 million worth of masks.

Spicer: Yes.

Thorndyke: Probably $2 million worth of masks now, even at the cost of Intel doing it. So you have to have everything perfect and, of course, you then have to have a database. You say, "I've got to know the propagation speed of a wire. I've got to know a propagation speed to turn a corner where you've got more mass, which means more capacity." With all that data, you are effectively going to databases that are firm. In fact, we developed and got a printed circuit board that had 90,000 holes in it, 16 by 20 inches. As such, how do you check it out? Well, what we found is that we could touch every wire and we would get a capacitance on that wire. So we got back to the simulation that told us what the capacitance should be. Now, if it was close, you could finally figure out that, well, this is five percent high, so you would normalize. All of a sudden, I got one that's way out, so you stick it over in a file. Later on, you're going to find another one that's connected to it, that is in fact the same number. So now, those two are connected together, and then you go to the file and find out where it's connected or how it's connected.

Spicer: Yes.

Thorndyke: So, all of our printed circuit boards were checked for accuracy and one of the binds in putting the circuit board layers together is you've got glue in between, since you're gluing all these layers together. Well, when you put it in a press, there are always three high spots. Any time you touch something, there are only three high spots. That's why you always have a three-legged milk stool, because it doesn't rock on you. So when you start putting pressure on, the hydraulic force starts squeezing fluid out from the high spots, because it has such pressure, the copper wires at the high spots would start to move and short out. So all of our printed circuit boards were put together with a water-filled bladder between the PC board and the hydraulic press, which evened the hydraulic pressure, but didn't end up squeezing the units together.

So the simulation is a critical piece, and it certainly is as good or bad as the software, but you couldn't have the computer chips today without it.

Spicer: Right.

Thorndyke: Now, that's the area where other people have a problem because they don't have that capability yet. And well, China just announced that they've got a super computer.

Spicer: Right.

Thorndyke: I don't know, have you read that?

Spicer: Yes. The sun, something about the sun, I think.

Thorndyke: Yes. Well, my understanding is, they've got a super computer and it's running one of the programs at super computer speeds, but just one.

Spicer: Right.

Thorndyke: The older boys tell me that.

Spicer: Can you give us a quick view of the 7600 and what need that was meant to address?

Thorndyke: Well, Seymour did the 6800. The 6800 was promised as a 7600. The problem was that with the 6800, he came to the realization he could not do it with the size of transistors that were used in the 6600. So for the 7600, he came out with a transistor that was the size of a small match head. It was very tiny with three wires on it, and he ended up having about four different units. There was a red dot, a green dot, a blue dot, and a yellow dot transistor. Therefore, he ended up having much higher packing density. The 1604 was two circuits; the 7600 was eight circuits. So he got eight circuits in about an inch and a quarter.

Spicer: Right.

Thorndyke: But in the same physical square size, using Freon cooling and as a chill plate.

Spicer: In the 6600, how hard were the cordwood modules to build?

Thorndyke: A bitch.

Spicer: <laughs> Can you tell us a bit more?

Thorndyke: The worse of it was repairing it.

Spicer: Okay.

Thorndyke: You've got to cut through all the resistors to get at the transistor inside. So you've got to start cutting everything loose. Now, some designs gradually put the transistor on the outside, so you could pull them out and put a new transistor in. But any transistor on the inside, you're dead.

Spicer: Okay. That's all the time we have today, Lloyd. I'd like to thank you again for spending the time with us. You gave some really great answers.

Thorndyke: Well, if I didn’t answer the other question properly on the problems that ETA encountered, you've got to ask Bob Price.

Spicer: Okay. Thank you.

Burniece: With that, I think we'll call it a day. I'm pretty sure that Dag and his partner who couldn't be here today, Gardner Hendrie, are going to probably have a ton more questions for you over time and they will probably also talk to Tony Vacca. Hopefully they are going to do that.

Thorndyke: Yeah, Tony Vacca has written quite a discourse. Tony was the person who convinced me that if we put it in liquid nitrogen, we would get twice our performance. We wouldn't dope it, because we would dope it for air and liquid nitrogen operation. But if we had doped the transistor to operate only on liquid nitrogen, we could have probably gotten five times the performance. But you can't test it.

Spicer: Right.

Thorndyke: Also in the ETA-10, Tony put in an onboard computer in each LSI chip that would test the unit. So you could load in about a 250-bit operand, run it for a second and a half, and get out a result that simulation says what the answer should be. It was about 98 percent or more reliable. So the first thing you did with the ETA-10 is turn on the power. You load a program to test all the chips. The next thing is you load a program that would send all the data to the other chips, and you go over there and you see if they got it under clocked periods. And if those two said, "Okay, I've got a chip that works, and propagation," and all the connections are alright, it's not shorted or anything. Then what you do is you load the master clear start up, and that was called clock cycles to get the condition of the chip proper, in order to turn on the go button. We could call from St. Paul to computer units in the USA or overseas, run a complete diagnostic on the computer, and check it all out. So you didn't end up having two years worth of training like we did in the Cyber 205.

Spicer: Yes.

Thorndyke: And you don't need to service it. However, the spare was another computer.

Spicer: <laughs>

Burniece: We're going to have to say goodbye now.

<crew talk>

Thorndyke: Okay, bye.

Spicer: Well, thanks once again to Cisco for hosting us. You guys have been marvelous and I'm looking forward to seeing a tape.

Burniece: Thank you very much, Lloyd, for being here today.

Thorndyke: Yeah, that's great.

Spicer: Thank you. Thank you.

Thorndyke: All right.

END OF INTERVIEW

References:

1. Control Data Corporation "Real Time" Newsletter, Vol 2, #14, September, 1965

2. Control Data 807/808 Disk Files, Components of the Control Data 6607/6608 Disk File System, Brochure

3. Control Data 803/804 Disk Files Brochure, 40820900

4. Control Data 852, 853, 854 Disk Storage Drives Brochure, 40818000

5. Control Data 852 / 853 / 854 Disk Storage Drives, your move to better man / machine communications, Brochure, 100.217 - 6/66

6. Control Data Corporation Midyear Report for Six Months Ended December 31, 1966

7. Computer History Museum Oral History Interview of Richard Berreth and Tom Murnan, June 23, 2009, http://archive.computerhistory.org/resources/access/text/Oral_History/102702131.05.01.acc.pdf

8. Computer History Museum Oral History Interview of Thomas Kamp, June 12, 2010, http://archive.computerhistory.org/resources/access/text/2011/10/102702337-05-01-acc.pdf

9. "The Demise of ETA Systems", Lloyd Thorndyke, *Frontiers in Supercomputing II- A National Reassessment*, University of California Press, 1994, Ames and Brenner, editors, pp 490-496

Editing notes:

1. Initial Edit by Tom Burniece, November 15, 2011
2. Gardner Hendrie edits, January 16, 2012
3. Tom Burniece added references, January 16, 2012
4. Lloyd Thorndyke edits, August 13, 2012
5. Tom Burniece final edits, October 23, 2012

