Oral History of Dal Allan

Interviewed by:
Tom Burniece

Recorded: March 12, 2013
Mountain View, California

CHM Reference number: X6787.2013

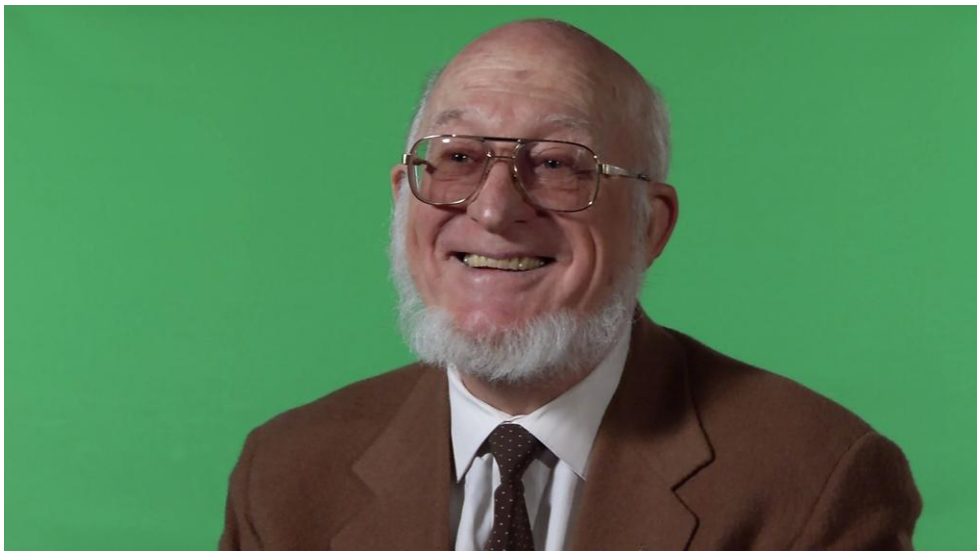

Dal Allan

Tom Burniece: Hello, this is Tom Burniece. I'm a volunteer at the Computer History Museum and we happen to be in the museum's Orientation Theater today. I'm going to do an interview of Dal Allan. Dal has been in the storage business almost since they invented disk drives and has a long history at IBM, ISS and several other places. But a good part of the story we're going to talk about today is his involvement in industry standards for disk drives, which he did over a period of about 30 years. So in order to get started, Dal why don't you introduce yourself and tell us a little bit about your family background, where you grew up, where you went to school and so forth.

Dal Allan: Basically my dad went over to Europe during the war, in Scotland, and met a girl. I was already around at the time, because she had been previously married. We wound up going to Australia after the war. I don't remember the war years much but Australia was great. We lived in the country; had all the space in the world. Then Mom decided to go back to Scotland to visit family and that was a pretty miserable period going from wide-open spaces to a tenement in the middle of Edinburgh.

What I learnt there was something very interesting. When you are in a class, where you cannot understand the language, you have to listen very closely - they spoke with such strange accents. When we got back to Australia, I went into high school because of the age gap. With the half years being different, I was moved forward a year in Scotland and moved forward a year in Australia. I wound up going to high school with a few other people who'd been moved forward and we all graduated early, from Armidale High School. We entered university a year younger. We were kind of awkward; we weren't at the right stage in our lives to be at University.

I worked for BHP, the steel company of Australia, because they were the only company that came to high schools to recruit. The pitch was, "Hey if you get a high pass in the Matriculation Certificate [statewide

high school examination] we'll give you a scholarship to university." That sounded good because the school guidance counseling had given me the advice that I should be either an accountant or a teacher. Dad was a teacher and when I got home and told him that he said, "Don't be bloody silly."

That was the end of my future in teaching so I took Economics, because that's what BHP wanted me to do. I was full time for two years at Sydney University, which was great. We had income during the holidays, because we had employment at BHP; I spent some time in the blast furnace, the rolling mill, and the sinter plant. The sinter plant was so dirty that you would have guys come in in the morning to start their shift at seven a.m. and, when the smoko [coffe break] came at 10:30, they'd go to the showers and leave. It was such a dirty job. The exposure was to a huge range of nationalities, people and personalities - very, very different to my upbringing.

When I became weary of student-poor I went to work for BHP full-time. BHP paid on seniority and I discovered I was the lowest paid in the office. That's how I got to IBM. I heard at class [university] one night that "They don't worry about things like that." When I applied, the guy in Administration said, "Oh I don't think we could pay you as much as you're getting at BHP but we pay the sales people a lot." He set up an interview with a branch manager and I was hired. So that's a snapshot of how I got into the computer business. It was because I was too young to get a decent salary at BHP.

Burniece: So how old were you when you actually started with IBM?

Allan: Nineteen.

Burniece: Wow—as a sales guy?

Allan: No. IBM didn't hire you in as a salesman in those days. What they had was hiring characteristics. You had to be a graduate and been in the workplace for three years. IBM would hire you and decide what you would be. Would you be a systems guy or would you be a sales guy? That was a personality thing. They were looking for characteristics to begin with and the branch manager, Frank Barr-David, interviewed me, gave me the aptitude test and offered me a job. I owe Frank a lot, because he kicked over the tracers on that one. He gave somebody who wasn't eligible a job and I finished my degree, while working for IBM.

Burniece: So what was the training that IBM gave you then? Did they give you substantial training?

Allan: Oh yes. In those days IBM didn't have a cadre of people to choose from. First of all, nobody knew what a computer was. Nobody even knew what card punch equipment was and IBM was completely unknown in the country. You were immediately sent to school for six weeks, to give you education on the

equipment and also to give you the feeling of whether you wanted to work there. It was a real weeding out process, because some guys would come in and say, “I don’t need to put up with this” and they just quit halfway through their training.

Then you’d go into the branch office and be assigned roles. Typically you’d be given an account. This was unit record days, so everything was punched cards. Trainees would be given either to a large site, to learn the ropes there, or be given a new site and told, “Just get it running.” It was pretty hard to get a unit record system running, because every step has to be programmed, every patch panel has to be designed, every piece of equipment in the flow has to be in order, so the operators can handle it. It was an incredibly detailed process and during this time the branch manager will take you places with him or he’ll send you with a senior salesman. Johnny Bagshaw and I had joined the same day and Frank walked up to the two of us and said, “Hey let’s make a call together.” “Sure.” And on the way out the door he said, “Where are your hats?” We looked at each other and said, “Hats?” He said, “You can’t make a call without a hat.” You have to remember this was the day of white shirt, dark blue suit...

Burniece: What year are we talking about? Just to get calibrated.

Allan: This would’ve been 1960.

Burniece: 1960.

Allan: Yeah. Big signs all over the place - little leather notepads, with "Think" on them and so on. It was still kind of the rah-rah environment — in fact people used to sing early in the morning. Start the day with an IBM song to rev you up, much like the Kirby vacuum cleaner.

Burniece: Were you singing the IBM Fight Song or what?

Allan: There was a guy by the name of David Greatorex, a salesman who wrote his own. When Frank was out of town you had to sing Dave’s and I was very, very lucky. I was hired during a period of examinations, so I didn’t have to come in and do the singing on Monday morning. I always had an exam at 10 a.m. and I had to be at that.

Burniece: Coincidence, huh?

Allan: Oh yes, it was sheer coincidence. As a result, the only time I sat in a singing session was when Frank arrived back from a United States trip and I remember sitting there in the conference room with the

sheets handed out for a song Dave had written. It started "Welcome back to FBD from the good old USA" and that's never left me; it's stayed in my mind forever. It's like "what did I get myself into?"

I'll tell you the difference between a salesman and a non-salesman. John looked at Frank and said, "Do we have time to stop at the store on the way?" He bought a hat.

Burniece: This is for the hats?

Allan: Just for the hat. I didn't. I was systems from then on. That was it.

Burniece: So the lack of a hat made you a systems guy.

Allan: The lack of willingness…

Burniece: To get a hat.

Allan: …to get a hat. But the fact that John volunteered "Do we have time to get a hat?" He went on the call with Frank and I stayed in the office.

Burniece: All right. Now just again for another calibration point, 1960: What was the kind of equipment you were working for with IBM, what model numbers and so forth?

Allan: Calculator was a 602A; accounting machine was 421; collator 088; sorter 082; card punches 024 and 026.

Burniece: You just happen to remember all these?

Allan: You don't forget that kind of early stuff. The only removable element in these machines was a patch panel, which was basically an engineering panel, and you used a variation on the banana plug. In this corner would be accumulators and over here would be different functions that you had to program the machines to do what you wanted. Every step in the process needed a different patch panel. For example, on a billing system the women operators had to keypunch the meter reading into a card that was pre-punched with the old reading and an identity number. You would have to calculate and punch the difference and behind that card you'd have to collate the name and address for the printing.

Each of these steps had to be put together and they were all different patch panels. You had dozens of patch panels to wire to get the things rolling, as well as understanding what had to be done at each step. It was tremendous training for the nitty-gritty of architecture; if you didn't get a step right the system didn't work. The patch panel for that step had to be fixed and probably the patch panels behind the failed point, as well. It was pretty intensive structuring but it was good training.

Burniece: How much programing were you doing in those days and what kind of programming was it?

Allan: The Service Bureau had a 650 and all of us, even the sales trainees, had to write programs for the Service Bureau. Every salesman had a quota to sell X amount of Service Bureau. Everybody who was not a salesman had to fulfill that quota by writing the programs for it. There were maybe a couple of other 650s. Certainly the AMP insurance company had one; MLC had one. My first computer experience was writing a job for the service bureau on the 650 and we did it early on in machine language. There wasn't a compiler.

Next came a compiler called SOAP (System Optimizing Application Program). The 650 memory was a drum and we'd been taught to write the programs in class as Housekeeping, Mainline and Subroutines. You'd have three pads of paper and you'd write the programs accordingly. Nobody, and I didn't even realize it at the time myself, knew why my programs seemed to run faster. I compiled mine as Subroutines, Mainline, Housekeeping, because the early commands were optimized.

An instruction took X number of cycles, which was so many positions on a drum, so the first commands would be optimized. This command would finish and the next one would be right there on the drum; it was very fast. Towards the end it got pretty bad, because you had to wait partial rotations to get to the next available memory location. My subroutines did run fast, and that's why I earned a very unusual reputation as a great programmer. It had nothing to do with my programming skills; it was the position on the drum of the next command.

Burniece: So that was a little insight you gained early about the value of understanding the storage system and how it works, latency and all that.

Allan: Yeah. It didn't seem to be a breakthrough. It just seemed natural that, if the next position was the right place to be, the things you wanted to run fastest should be the commands that ran most often. It just seemed very normal, because latency was part of the 650. It was a programming factor - by putting subroutines first you had the fastest running programs.

Burniece: Tell me a little bit about that drum. Was that something that IBM designed or something…

Allan: Yeah.

Burniece: …they picked up from Univac or where did they get it?

Allan: I don't know. It just turned up in the field. The 650 was a bi-quinary machine, you looked at the panel and you had a light and two columns of four lights; off/on was your 0 or 5 and then you had 1, 2, 3, 4. I'm assuming the drum came out because that was the only technology available at the time for memory. It also had a very, very expensive magnetic core memory. It was an add on to make programs run faster, because it wasn't dependent on latency.

Burniece: Now of course there was no hard disk at that point. But before I ask that, do you remember the capacity of that drum?

Allan: I'd have to look it up. I'm sure the answer's here in the museum somewhere. But it wasn't huge.

Burniece: Yeah it was probably about a megabyte range or something like that, I would guess. Right?

Allan: It was only used for programming; it wasn't used for storage. The input/output device was a tape drive or a card reader. And so every job would be set up with a card reader to read the cards in. You also had a patch panel on the 650. It was so low in compute power that every job you ran you actually had to use a patch panel to move the data to the correct position on the drum, so you could process off the punch card stream.

Burniece: Explain a little bit more about that because I'm not sure I understand that. When you say you used a patch panel to move it, you mean you're actually connecting directly to certain tracks on the drum, with a patch panel?

Allan: No. No you had to move data into the memory space, which was on the drum, and to get it there you had to read it from the card reader, so it had to be moved in there. You didn't have read and write commands to anything.

Burniece: Oh all right. So you literally are doing this manually.

Allan: So you could read cards. Yes.

Burniece: Reading and writing.

Allan: The printers, which were attached, had patch panels; the card readers had patch panels. And so the 650 had a patch panel for the card readers to move data to certain places and then you could manipulate it in memory with multipliers and so on.

Burniece: Now RAMAC had already come out four years earlier. Were any of those in Australia at that point?

Allan: Oh no, RAMAC came out well after the 650.

Burniece: RAMAC came out in ’56.

Allan: Yes but the RAMAC was attached to different equipment than the 650s. As I remember the first RAMAC came in to go on something else.

Burniece: I don’t remember the model number but I believe it was a 350 system.

Allan: In fact—

Burniece: But that did come out in ’56.

Allan: Yes but we had a 650, which was an earlier machine. The 350 went into Colgate-Palmolive and that was a mindblower for the whole branch. I didn’t work on the Colgate-Palmolive account. I’d been there and visited. There was a room full of punch cards and these were in bins. It was staffed with women who, when an order came in, ran around plucking cards from the bins. If a wholesaler wanted three cases of Persil, which was a soap powder, they would pull out three cards. The women ran through plucking from the bins and then the accumulation of those went to the unit record equipment and eventually you’d get an invoice out.

The RAMAC came in to solve that and it had to have assigned addresses, which were unique; synonyms slowed everything down. Everyone was involved, because the operation needed a simple random allocation algorithm, and the whole branch was put onto the job to help the engineers on the account. We all had to learn about the RAMAC because we had a branch problem; they couldn’t come up with enough unique addresses.

Burniece: So again what about timeframe was that when that happened in Australia?

Allan: It would've been '61 maybe.

Burniece: Okay. So that was your first experience with a disk drive. But your initial experience was actually working with the drum.

Allan: Yes.

Burniece: And patching in the data directly for read and writes from the card deck.

Allan: Yes. And it was kind of funny because, years later when I was at ISS and we were acquired by Univac, guess what? They had an operating system based on a drum. When they talked about one-word writes, I knew exactly where they were coming from because that's what we used to do. Of course the idea of a one-word write on a disk drive was anathema but that's what the operating system was based on. It kind of like came full circle there.

Burniece: Well tell us a couple of stories about those early days with IBM and some of the things you worked on, plus maybe some of the people you remember or some of the problems you had to solve.

Allan: The biggest problem in those days was me. I was just too young. I owe a huge debt to a few people in that timeframe because they taught be how to behave, let's say, in society. It was kind of rough and tumble growing up in the country; there were no such things as tact. We didn't know what tact was. We just told each other blunt truths.

That didn't go over too well for some. I remember on the first account - it was North Shore Gas - and the accountant came down and he asked me how long it would take to get finished. And I said, "It'll never happen; there's not enough equipment here; we need this, this and this." Next day David Greatorex has me in and he's chewing me up, down and center, because I'm not supposed to tell customers that. [He was] supposed to go in and talk the guy into more equipment. Here was I blurting out that IBM underbid the job.

I needed a lot of help and we had an assistant branch manager by the name of Art Langlands. Art was an interesting story. He was in Colorado Springs and the Air Force had been one of his accounts. Remington Rand won a bid and IBM lost their equipment in there. For all of the equipment that came out, Art had to pay some huge penalty, like 250 percent of the points he would have won. The money he would have earned for selling that equipment was now a negative on his balance in IBM. He had to repay IBM that loss and could only pay it out of commission at 10 percent of all future commission income.

Art won the account back the following year but he couldn't do a lump sum return, so he was committed to being on commission at IBM for the rest of his life, or he could leave and the debt was wiped out. IBM didn't fire [for losing an account]; they simply gave you a reason not to stay but Art stayed anyway. He was a blue guy, right? He could be in Australia on assignment and be on commission, as an assistant branch manager. Art took me under his wing and taught me how to function.

He also taught me the psychology of inspiring people. I may have told you this story before but we were having a down year and Ken Hudson had been our trainer, before coming into the branch as a salesman. He was one of the most popular guys you'd ever meet. We came to work one day and Ken's desk is empty - no sign of Ken. No word, just Ken's not around. Art said, "Oh yeah Ken won't be here anymore." It was like "Wow, they fired Ken!'

There were worse salesmen than Ken, so why did they get rid of Ken? Art asked "Well who do you think should've gone?" and I said, "Oh Bob or Fred"; whoever it was. He replied, "Yes, he is the weakest. What would you do if he was fired?" and I said, "Everybody expects it." That's when Art pointed out, "Nothing would change. I want to be in the Club. I'm not going to get into the Club at the rate these guys are selling because my quota is 175 percent of their total quotas. They have to do better to get me in--take a look at so-and-so. I know he's got X number of points accumulating for next year but he doesn't want to bring them in this year, because he's comfortable. And take so-and-so." Art went down the list and then said, "These guys need to be motivated. If I got rid of . . . they wouldn't be motivated. You watch, I'll be in the club in September."

Sure enough, when he went down the list, he knew every guy working for him like "Bill's wife, she's learnt to live in the style she's accustomed to; she's not going to let him get fired." That was it; you don't motivate people by firing the worst performer. You pick somebody who's up near the top, who's very popular, and scare everybody. That was probably my strongest life lesson, the kind of things that he told me. He had a habit; every day he went to the same place for lunch, ate the same meal, and he took a different person out of the branch every day just to chat. I learned a lot.

Burniece: Wow. How long did you work for this guy?

Allan: Art? He was the guy who sent me over here. He called me in and said, "We're sending you to the States. You got to be in New York." I said, "When?" He said, "Next month." I said, "What about my fiancée?" "We don't pay for girlfriends". I picked my fiancée up after she'd come off shift that night - she was a nurse - and said, "I'm going to be in New York this time next month. Do you think we could move the wedding earlier?"

Burniece: If you want to go along we need to get married.

Allan: Ros graduated three weeks later, married me one week later and got on the plane the next day. That covers the '62, '63, '64 timeframe.

Burniece: You still married to this lady?

Allan: Oh yes, it's worked out well; it's been a good run for both of us.

Burniece: So about '62, '63 you moved to …

Allan: It was actually April of '64 that we came over here.

Burniece: You moved to New York.

Allan: April 19th. I can tell you the date exactly.

Burniece: And that was up to headquarters someplace?

Allan: No, I went to Endicott. When we were in the branch, I had been to the first 1401 class with John [Bagshaw] and the two of us didn't impress the instructor at all. We thought it was so easy - you had Read commands to read cards in. It was so much better than the 650 that we couldn't believe how easy this machine was. What we didn't know was that the instructor was one page ahead of us on the teaching guide. We were pressing for more information, because we didn't need five hours to learn this stuff. It was understood in five minutes. The instructor was irritated by that and when we returned to the branch, I was assigned to install the first 1401. I went to work for a guy by the name of Paul Gyjax.

Burniece: This is literally the first 1401 ever?

Allan: In the country (Australia).

Burniece: And where was that?

Allan: Coming to the Sydney Service Bureau, that would've been around about '61. Paul was a Swiss National who had worked everywhere in the world and he taught me about cultures. He gave me some really interesting advice, which stuck with me for a long time. "When you go back into the branch after this 1401's in, you're going to be the center of attention for a lot of people. Don't let it go to your head." He said, "Worse than that is you're not going to treat them correctly. If they have to come to you 10 times

in a row, they're going to start to resent you. You have to prevent that happening; find out something about them that they can give you advice on." "It doesn't matter what it is - your car, whatever - so that they feel they're giving you back something in exchange. Then they'll come visit you as a friend. They don't come to you for advice; they're just coming to get some feedback. You have to get into their head, not be somebody that they have to keep begging information from." That's a pretty invaluable insight when you think about it.

Burniece: That was a fascinating piece of insight.

Allan: You know it is and I needed more than a little of his advice at the time. He probably realized I would turn into a pretty difficult person because I was just too young.

Burniece: I got a little bit mixed up on the dates. I thought you left for Endicott in '64.

Allan: Right. I just flipped back to '61, when you were asking about the branch experience.

Burniece: Okay, all right - the story goes back, right?

Allan: Yes.

Burniece: But the 1401 that you installed, when was that, the first 1401?

Allan: That was '61.

Burniece: It was '61. Okay all right.

Allan: The 1401 came out of Endicott. Before it arrived and the reason I was sent upstairs to work for Paul was that IBM Endicott sent all these listings. Remember the old green and white sheets on the wide paper? There were huge books of these sent out to the countries. We had to sit there and find the bugs without the equipment, by reading through the listings, and surprisingly we did. That's how bad the software was and that's what we did until the hardware came in. We had to do support and everybody was raw, really raw. The only people who knew the 1401 were the CEs who'd been to the U.S. and the software wasn't all that great. We kept finding things wrong and would keep submitting APARs, basically a bug report.

This was to turn up a few years later because, when I arrived in Endicott, it turned out they knew my name because of all these APARs from Australia citing problems. We were pushing the hardware much further than most of the U.S. users, who bought a larger piece of equipment to do the same job that Australia bought a small piece of equipment to do. We were pushing it harder and, as a result of that, we found problems earlier. Turns out there was a whole list of things that we reported on that Endicott didn't think was a problem, until it turned up as a problem later on.

Burniece: So you were a little infamous in Endicott about these bug reports that they didn't feel they needed to hear about.

Allan: For the wrong reasons and another thing is that, because I was in that first install, I wound up being assigned to AMP, the biggest account in Australia. They had intimate knowledge of the equipment and they had a very low opinion of IBM people.

Burniece: What does AMP stand for?

Allan: Australian Mutual Provident, a big insurance company. They knew their equipment better than the IBMers sent to "help them". AMP figured it was part of the cost of doing business with IBM; they had to put up with these people. AMP gave me this job to go in a corner, just to get me out of the way, and said, "Make this fit the 1401. It's a multi-run, multi-pass actuarial job on the 650; make it work on the 1401." They had decided it wasn't possible but I didn't know that. I went off in a corner and figured it out and they were surprised, so they didn't get rid of me at the end of the 60 - 90 day period, whatever they normally did.

AMP bought a 1410 and about a month later one of the guys mentioned to me that the Sort was blowing up. I asked, "Why didn't you tell me when it started happening?" and the reply was "I didn't think you could do anything." That evening I stayed up all night trying to find this problem in the Sort and went home at dawn to get breakfast and some sleep. I left a note on the senior programmer's desk saying: "The bug's on this page; I can't find it." About four o'clock I came back in and the place was running. He found it immediately, twigged where it was in a shot, but he gave me full credit. They'd been looking for a month and hadn't found where it was happening. He found it in five minutes, once he knew where to look.

It was interesting, we both told the same story. "He found it". That's how I learned that when you give away the credit for something, it reflects on everybody involved. Another lesson learned. I thought I'd failed 'cause I didn't find it. He thought I'd succeeded because I'd narrowed it down. Pretty much a quid pro quo, so I became part of the AMP team. I was asked to do something to improve the loading of software, because everything was done with the programs on punch cards. Programs were loaded in the card reader, so one card out of order or lost or crinkled and the job couldn't load. The operator had to

figure out what it was. I came up with a 160-character resident routine; the operator would put in a header card and we pulled the program off tape. I did a lot of odd things like that for AMP, which was pretty good learning.

Burniece: Before we go back to New York, any other fun stories from your days in Australia, at IBM?

Allan: Probably the day I should've been fired.

Burniece: That sounds interesting. Tell us about that one.

Allan: Honeywell had turned up at the door of AGL, the Australian Gas Light Company, with a huge computer that they had coming into the country. Honeywell had sold one and they needed a second one as backup, so they were searching for a home and approached AGL with this great deal. IBM had a competitive library and you could go in there and study all the equipment. The Honeywell machine was in the 1410 class and they were selling it at a 1401 price. I used to be pulled in on most of the competitive situations, because I was one of the few guys who'd sit there in the library and try to understand what the opposition was.

Come decision day, Alex somebody [the data processing manager] said, "Pretend you're sitting in my chair and I'm going up to the Board of Directors at nine o'clock." It was a quarter of nine and he asked, "What would you do in my seat?" "I don't know Alex, you're getting this machine at 30 cents on the dollar, and IBM can't come up with a number that'll match that." "Okay, thanks." Alex went upstairs and I went back to work and about 11 o'clock the branch breaks into great glee. They had just received the phone message that he'd chosen IBM. Everybody knew that I'd been in to see him just before the decision process and assumed I'd talked him into it, when in fact I'd actually done the very reverse. I'd told him that Honeywell was a better deal. Nobody asked me what we talked about; they just assumed that I'd done the right thing.

Burniece: You'd done the right thing.

Allan: If anybody had ever said, "What did you tell him?", I'd have been in trouble because I would've blurted out the truth.

Burniece: So how did IBM get the sale? Do you know?

Allan: I think Alex looked at the deal and said, "This is too good; they're going to screw me somewhere." That's my guess.

Burniece: Okay.

Allan: Alex had a good head. IBM had bid a 1460 and Alex did not hire programmers. He brought in all of the people in the billing/accounting system and said, “Turn them into programmers.” That was very wise. At the time I thought he made a very foolish mistake when he insisted on 'everything' being in the system, because IBM had bid it based upon the exceptions. If you are leaving town, you call in for a meter reading and it is not part of 'your' day billing cycle. The month was broken into processing cycles. So, if you were out of cycle, it was a special, which would be handled on the side as an exception.

Alex wanted everything in the day, so that meant every night we had to run the entire number of tapes, There weren’t enough hours in the night, because the software driver would read into memory, move to a work area, move from the work area to an output buffer and write to tape. The tape speed was the same as memory speed. So you had 4X time for over 90% of the input. I wrote a driver that read into memory, checked for the need to process, and if not then wrote from the same memory. That’s what I meant by Australia pushing the envelope. We would never have finished in one night, if the tape database had to be moved through memory four times. Those were some interesting times in the field.

Burniece: Fill in just a couple more pieces for me before we move on to Endicott. You said you started at IBM before you finished your degree.

Allan: Yes.

Burniece: Did you switch degrees at that point and do something other than Economics?

Allan: No I stayed in Economics.

Burniece: So you got an Economics Bachelor’s Degree…

Allan: Yes.

Burniece: After you started at IBM and were learning how to program and …

Allan: Yeah.

Burniece: And become a systems engineer. Wow.

Allan: When you said systems engineer it reminded me. When AMP didn't want me for anything. I spent all my time with the customer engineers to fill in my time. They taught me about hardware, even taught me how to read logics. I don't know if you remember it, but the equipment came with trolleys of books - over in the corner was always a huge stack of logics.

Burniece: Sure.

Allan: They were on wheeled carts and they explained to me how things worked - walked me through the logics. Some commands were pretty complex and they'd say, "Oh this is the way it works." I think they enjoyed somebody showing interest, because they were hanging around waiting for a call. They were on site and we had two CEs on call all the time. That's where I got my first engineering exposure, learning from them.

Burniece: That's fascinating. So you actually got an Economics degree while you were learning to be an engineer, essentially learning engineering.

Allan: Not qualified as an engineer.

Burniece: On the fly. At least a computer scientist anyway, as they now call it.

Allan: Yeah, whatever.

Burniece: Wow. Anything else on the Australia side?

Allan: Culture?

Burniece: Yes, do you want to talk about it a little bit before we move on?

Allan: Culture was very different. You lasted as long as you were performing well on 'this' job. There was no credit for past successes or anything like that. It was the quick and the dead. You were rewarded as well; motivation was everything. I remember one month we had a competitive win and it was against NCR. That's when I got all the CRAM exposure. IBM had nothing like it. It was a fantastic product but we eventually won the account and I was given a raise the next month that was made retroactive a month. I didn't think it was very good compared to what success we'd had. Then the next month I received another raise, which would take place the following month. Personnel would not allow too big a raise in the same month so the branch managers finagled the system to give me one retroactive and one going forward.

Not only was motivation part of it, because I can tell you that motivated me. The branch managers also manipulated the system to keep their people pumped and make things happen the way they wanted them to happen. Managers didn't let guidelines get in the way. They didn't overtly break them but they bent them and moved around them or circumvented them. There was a lot of entrepreneurial stuff going on in the branch at that time, so it was a great learning experience. We were 'starting' the IBM company in Australia.

Burniece: It does sound like a very interesting learning experience. Tell us a little bit about the NCR CRAM thing and, just to fill in for those that aren't aware of this, you just wrote an article not long ago on the NCR CRAM for the Storage Special Interest Group (SIG) site [1]. Talk about that a little bit and what that experience was all about and what you thought of that machine.

Allan: It was impressive. NCR had a good track record in the banking industry and they used cards for everything. They were head to head with Burroughs and I can remember—

Burniece: When you say cards you're talking about …

Allan: I'm talking about magnetic cards. NCR had some extremely complex accounting machines for banks and they translated that across into an accounting computer that was very much biased towards the industry they served. NCR came up with the idea of using the same kind of strip that was used in their computators, or some name like that, in the branch offices. The operators would pick a card from the bin of customer accounts, pop it in the computator and update it

Burniece: And these are now magnetic strips, not 80-column punched cards?

Allan: Yes, not punch cards.

Burniece: Okay.

Allan: The computator did the daily updates in the branch and NCR took that principle and automated it by wrapping strips around a drum. The system could read and write on that drum and you had all the output facilities. CRAM had a reasonably good capacity - 256 cards in a box (0 to 255), which were notched, so the selected strip would drop in. CRAM wiggled the pins; and the only one that matched would drop out. Pretty much like a key tumbler. I only saw one once and it was kind of a horrifying experience. It was noisy. I have no idea how loud; it ran on compressed air. There were three or four of these in a room and it was noisy. I wasn't there when they had a crash but apparently that was like the wail of a Banshee - the card would drop in and concertina when it hit the drum that was rotating at speed, so the drum would slow down and the magnetic card would crinkle. It was apparently ear splitting …

Burniece: Literally be a screech.

Allan: Yes, Hound of the Baskervilles routine. CRAM was impressive. IBM had nothing like it at the time [2321 was introduced two years later]. The capacity was 255 cards in a box and I think 16 boxes to a cabinet. You could have all the boxes you wanted stored around the room. You'd just go over and pick one up. What was really dynamite is that a customer could walk in and say, "I'd like my account statement as of today." NCR could give it to you at a time when most branches had listings that were run once a month. They would ask, "What checks have you written?", total them up and say, "Well this is what we think your balance is." Banks with CRAM would have an operator [at the head office] go over to the wall, pick up the numbered/alphabetic box, pop it into the machine and the answer would come out on a terminal at the teller. There was only one or two terminals in the office and you'd be there chatting about the weather, or whatever, and a minute or so later you'd get your answer. There was nothing like that from IBM. It was truly an immediate experience.

Burniece: Real-time account information in what? About 1962 or something like that?

Allan: Yes.

Burniece: Wow.

Allan: That was a mindblower. Everything IBM did was batch and here was something happening now. It was like - how do you sell against a machine that can do that? IBM wasn't the leader in the game and we [IBM] were behind in a lot of technology.

Burniece: Part of what they were trying to do with the RAMAC was to actually be able to automate that kind of stuff and be able to get fast answers.

Allan: But the RAMAC wasn't big enough. That was the beauty …

Burniece: From a capacity standpoint, okay.

Allan: That was the thing. On the RAMAC, like Colgate-Palmolive, they took one function, basically inventory stocking or fulfilling orders, and put that on there. They couldn't put anything else on, because it was too small. NCR solved that by making the data removable, what you might call accessible.

Burniece: Right.

Allan: It wasn't immediate. If you ever saw the arm on RAMAC moving up and down. It was incredibly fast.

Burniece: We got one right behind us.

Allan: That's something isn't it? NCR's CRAM card solution was unit record, individual cards, grouped, removable, and accessible. It was the first near line storage and it was pretty terrific for its day.

Burniece: This is a little bit of an interesting segue here. You went back to Endicott in '64 and that's right about the time that they announced the 360.

Allan: The 360. I can fill you in on the background there.

Burniece: System 360 was '64, right?

Allan: From what I heard, in 1960 or thereabout, IBM examined what it was selling. It had commercial computers like the 1401, 1410, scientific computers like the 7090 and variations between the two for different applications. The senior executives sat down and decided on having one machine for everybody. About the time I joined IBM in Sydney, back in Poughkeepsie and Endicott they were coming up with the idea of a single machine.

Burniece: At least a single architecture because they had ….

Allan: Single architecture, yes.

Burniece: They had multiple models, but a single architecture. That was the 360 architecture, right?

Allan: That's right. IBM put a tiger team together from people around their different laboratories to help with the architecture and then they faced having to implement it. They had the engineers but they didn't have the programmers. And this was the first [general purpose] machine that really, really needed an operating system to run. You had one on the 7090, but it wasn't very sophisticated, and the 360 was a stepping stone beyond. IBM Labs told World Trade, we want programmers from everywhere and every country had to contribute one person. I was the token Australian.

Burniece: So, that was the reason they sent you to Endicott, to be the token programmer on 360?

Allan: Yes.

Burniece: Right, Let's walk in this area a little bit because the guy who ran that program, I can't remember his name, wrote the famous book, "The Mythical Man-Month."

Allan: Yes. Fred Brooks.

Burniece: He found out that you cannot just throw more people at the problem and get it done faster. The more you throw at it, the slower it gets.

Allan: That's right.

Burniece: It takes longer, because of additional communications and so forth.

Allan: That's what happened in Poughkeepsie.

Burniece: Right, but that original idea was we'll just get somebody from every place, and we'll get this done in a small amount of time by having lots of people working on it.

Allan: Poughkeepsie didn't want the World Traders. I was originally assigned to Poughkeepsie and halfway over learned I was being sent to Endicott instead. Endicott wound up with all of us - South Africans, Swedes, you name it. A real melting pot of foreigners dumped in the middle of this provincial town in upstate New York but the Glendale Lab welcomed us. Endicott is where the model 30 was developed, and where DOS (Disk OS) and TOS (Tape OS) were developed. Four of us did the design for [BOS (Basic OS), which was the predecessor to] DOS, Frank Sebastian and Lou Stabile were—

Burniece: This [was] the [original] disk operating system?

Allan: Yep, they were the guys already there. Ray Lorie from Belgium and myself were assigned to Frank and Lou to do the operating system, the supervisor. Three of us didn't know what an interrupt was; the 1401 didn't have them. Ray had experience on the 7090 so we used two offices. The three of us would be in one room brainstorming on the chalkboard and, as we thought of ideas, we'd call them out to Ray in the next cubicle. He's there programming them live, like the first thing we figured out was we needed an interrupt handler. "Can you do that, Ray?" "Sure." His English wasn't so good, so he loved that part of the task, programming our concepts off the chalkboard next door. That's how primitive it was when I arrived in April '64 and we didn't have a machine. We had to do all this on SUPPAK, a 7090 program that

simulated a 360. Lou did Job Control. I was given the Linkage Editor, and Frank kept the supervisor with Ray.

Burniece: I'm a little fascinated by the timeframe here because I graduated from college in '64 and interviewed at IBM when the 360 was being announced. I know that was in the spring of 1964. You joined them in April? And they didn't have the disk operating system, yet?

Allan: No.

Burniece: So, how long did it take to do this?

Allan: The 360 was a paper tiger.

Burniece: This is a paper announcement, the 360?

Allan: Yes, it took four years to plan it, put the structure together. Now, I don't want to make it sound that bad; there were huge architecture books. A whole team of people developed the architecture for the 360, and that was in these binders; they physically filled a room of this size. Every shelf from floor to ceiling was filled with architecture books because they covered every aspect. An Englishman had come up with the specification for the Linkage Editor and, of course, they sent him back to England three weeks before I arrived, so there was nobody to talk to. That was the same for all of us. IBM brought in a wave of architects to do the design and then sent them back to the labs in England, Europe and stateside. We were working from these books, and the day the System /360 was announced it really was paper - a good solid idea with an architecture but there was little or no hardware and there was no software.

Burniece: How long did it take until they actually had the first prototype 360 hardware / software combination running?

Allan: We probably got a working Model 30 sometime in late '64, early '65.

Burniece: So, later that year.

Allan: Yes, and I can tell you that we didn't know if the manuals were right or the hardware was right. The Load Address command was supposed to zero out the top byte of a 32-bit word because the address was only 24 bits long. What we'd written on SUPPAK was working but, when we ran it on the first 360, it was blowing up. I tracked down what it was. There were bits being left in the top byte of the Load Address word, so I went through my code and in front of every Load Address command I added a Subtract

Register to zero the contents. We assumed the hardware was right but as things turned out it was a bug in the Model 30. I wrote software for the next two years with Subtract Register in front of every Load Address, before finding that out.

Burniece: So, you actually almost invented the idea of a clear before you loaded the register.

Allan: It was supposed to clear anyway - just that the first machine didn't do it. I thought the manual was wrong and the machine was right. Back then you compensated; that's how primitive things were. It was pretty raw.

Burniece: So, that was the first real [disk-based] operating system IBM had done?

Allan: I believe so.

Burniece: And that was in '64 timeframe?

Allan: Yes.

Burniece: Before then, they'd had the RAMAC and the 1301. Now they we're coming out with the …

Allan: 1311.

Burniece: The 1311 – the first removable drive? Was that also coming out before they had a disk operating system to run with it?

Allan: The 1311 had been out there before 360. AMP had a couple of removable disk packs on the 1410. As I remember, at the time none of us understood what a cylinder was. The 2311 was the same machine revved for the 360.

<Laughter>

Allan: The cylinder was kind of a wild thing. The 1301 came before the 1311 and IBM had taught us about this big box that used things called cylinders. We tried to figure out how that worked but, since none of us had experience on it, we did not grasp the concept. We heard about it in the field and tried to figure out/guess what to do with it; then we had to write programs for the 1311. It was a real exercise because, first of all, the hardware was touch-feely. The whole I/O structure was very much touchy-feely.

One interesting thing is that because I had the Linkage Editor, nobody could run, unless the Linkage Editor loaded their software, because it was a dynamic mechanism. I had to complete real quickly and, after it was running, I had the time to rewrite, rewrite, and rewrite to improve its performance and functionality. I replaced all the individual I/Os with a single routine, which encompassed everything and ignored all provided parameters by requestors. I used the parameters provided by the Open to decide what I/O request [parameters] to execute. Every now and then, this routine would post a read error on a write command, which 'couldn't happen', 'shouldn't happen' and the manual said 'can't happen'. That was a 2311 bug that kept popping up through my life over the next few years.

Burniece: That's a great story. So, go ahead and keep going. Tell us about that bug.

Allan: Okay, I'll stick with that story. After being in development for eighteen months, I was supposed to go back to Australia and I was moved to Product Test. All of my complaints about this read error on write commands had been ignored but now I was in Product Test and you can't ignore a Product Test report. Product Test was the gateway to the outside world; machines and software had to pass all hurdles. My first day I wrote up about thirty or forty Product Test reports and Frank was my best friend and manager of the Supervisor group in Development.

Burniece: Frank's last name?

Allan: Frank Sebastian - and he went berserk. Until the day before, I'd been working for him and he went to the Lab Director and accused me of using inside information. These were things that I'd said before that we needed to fix and Frank had always said "we don't have time to fix that; it is good enough." Here I was, aiming at this litany of things and one of them was the 2311. The Lab Director shot Frank down and told Frank to fix his bugs but San Jose refused. We were stonewalled and everything was questioned.

Burniece: Did they know what caused the bug?

Allan: I don't know, but …

Burniece: Did you know what caused the bug?

Allan: I knew I was getting a read error on a write request and figured out what was causing it. This was the CKD (Count Key Data) architecture and to do a write you had to search R0 to write R1. You had to find the previous record to write the next one and the search command was failing. If you got an error on search R5 it would break the chain with a read error and you would never write R6. Sort used full track writes for efficiency. I was back in Australia working in Headquarters support, the last place to call before sending a problem back to the U.S. General Motors Holden had an industrial action. Thirteen guys were

paid twice, though nobody knew that at the time. All they knew was that thirteen guys didn't get paid. The issue goes straight to Detroit.

Burniece: They thought thirteen guys had not been paid and they actually had been paid?

Allan: They knew thirteen guys had not been paid; that was what the industrial action was about. In a General Motors plant, thirteen guys didn't get paid, they rolled up to get their envelope, and it wasn't there. The first thing that happens on a strike pending is an SOS goes out to Detroit. It is industrial action they've got to deal with, because the unions are screaming. The rocket goes up inside IBM too, both from the U.S. and back to IBM Australia, and then from the branch back to headquarters. "Get down there, and fix it!" I asked if thirteen guys had been paid twice because if Sort was searching R0 to write a full track R1 and got an error then the old R1 stayed there. Those were the thirteen guys who were paid twice; the guys who should have been written on that track were the guys who did not get paid.

Burniece: Oh, so it was a different group of thirteen guys, okay.

Allan: Yep.

Burniece: Now, when you say the Search was not working, was that a failure in the mechanics, or …?

Allan: San Jose fixed the problem on the 2314. The 2311 used a single Address Mark. If it was missed, then you had a Search error that was reported as a Read error. The 2314 solution was to use a pair of Address Marks. It was solved on the 2311s at GM, because San Jose sent out the fixes and a six million dollar bill.

Burniece: A six million dollar bill?

Allan: Yes, to figure out what the problem was and to fix it. I had the outstanding trouble report from Product Test Endicott mailed back to us in Australia and the report was given to WTC who took it back to San Jose and said "No. We're not going to pay this. You knew about it already. You just didn't bother to fix it." Years later I'm interviewing Marty Halfill at ISS for software called Environ One, a database response based system, using telecommunications. Marty was the VP of development and he asked, "Why do we want you? You're a programmer; we're a hardware company. I've never been too keen on this software stuff. "What do you know about disk drives?" When I replied, "Well, about the only thing I know is that on the 2311 …". It turns out Marty was the program manager.

<Laughter.

Allan: So, that bug—

Burniece: He didn't know up until that point, you're the guy that found the bug?

Allan: That's right. He didn't have a clue.

Burniece: Was he happy to learn this, or not?

Allan: No, he just said, "you're that guy".

Burniece: Darn, the Australian that found it.

Allan: Yeah.

<Laughter>

Allan: That was the end of my interview with Marty and the following week I was working for Irv Tjomsland in Environ One. Then Environ One was sold. It wasn't a good fit for ISS. It had been one of those decisions like "We've got a lot of money. What can we do to complement our hardware? Sell something that uses a lot of disk drives." It had been very appealing to the "Dirty Dozen" to come up with this but it wasn't a good fit, so it was sold to Cincom. It was a rancorous deal. Nobody was happy about the situation and Russ Brunner put me in charge. Russ didn't trust Tom Nies of Cincom because ISS remained responsible for the contracts in place, even if Cincom failed. All the tapes had to go in escrow and he told me "Find a way that I have control of this." The contract did not mention cards, so I punched all the programs off tape and boxed them up in order and stuffed them in the back of the wagon. I kept them at home for a couple of years.

<Laughter>

Burniece: Let's hold that thought and go back. I want to fill in a couple of the cracks. Let's go back again into that late Australian transition to Endicott. Talk a little bit about the storage side of this because I think you've touched on a couple pieces that are going to be somewhat confusing until people sort this out. Why don't you give a quick description of the difference between a RAMAC, a 1301, a 1311, and a 2311, just so people understand what the differences were?

Allan: Well, RAMAC was just a gigantic machine, twenty-four inch …

Burniece: Fifty platters.

Allan: Fifty platters, I think it might have been fifty mega characters of capacity [Editor's note – it was five mega characters]. The interesting thing about that was it used Bowden cables; it could move from the inside top to the inside bottom at lightning speed. It was fantastic to watch. It had a huge arm that used compressed air to push the head away from the surface.

Burniece: But it only had a pair of heads.

Allan: Yes, and so—

Burniece: It could only access two disk surfaces at once.

Allan: Very limited access, but compared to picking punch cards, it was dynamite. Then came the 1301, which introduced the concept of the cylinder, one head per surface, and it was fixed storage too. It was pretty big as well. I only ever saw photographs. I was never in the room with one. You could tell it was pretty big just from reading the literature. The mind blower was [the] portability of the 1311, a relatively small fourteen-inch pack that you could take on and off. I think it might have been five mega characters, something like that, and you could have two or more machines. Kind of like the CRAM, you could store everything over there on the shelf and bring it over to use.

Burniece: That was actually one of the reasons I wanted to come back to this. Do you believe that IBM actually understood what CRAM had done with removability and had to now reproduce that in the disk drive rather than a magnetic strip system?

Allan: No, I think these things were going in parallel. I think they were looking for a way to lower the cost.

Burniece: Okay.

Allan: If you look at the way to lower cost —it's like a taxi, which transports an awful lot of people with the same vehicle that you use and re-use. That's removability; it was that simple principle of "we put tape reels on and off, why can't we take disk reels off and on?"

Burniece: All right, okay.

Allan: It was the search for more availability online - more stuff to do during a day. It was still batch processing; we weren't talking about interactive. You could schedule your activity during the day and that's what was done. The real value for most shops was getting rid of the punch cards as the programming entity; you had your programs on the disk pack. You'd load up all these reels of tape, pop on the right disk pack, and it could load all of the programs in sequence that were needed because everything [data] was still heavily on tape. I don't remember the 1311 as a dynamite data device but it was a dynamite-controlling device to prevent lost cards and dropped cards. It was a real automation boon for the programmers and operators because it was right here. Nobody could get it wrong, as long as it stayed, and it was safe. The same program executed the same way all the time.

Burniece: So, that's an interesting piece of insight. Expand a little bit more on that. You're saying that basically the removable disk pack, in those days, was more of a program load device. And the tape was the data storage device.

Allan: Yeah, that's right because—

Burniece: So, this disk operating system you were writing had to figure out how to marry those two pieces.

Allan: Yes, you didn't have much support for these products when they went to the field. Like on the 305, there wasn't a huge amount of software support. The field branch office, the guys in charge of the installation, had to figure out how to use this newfangled thing. And that's why, as I said, the whole branch was used to find out a good random number. We wound up, by the way, using one from Colgate Palmolive's inventory numbers. It turned out to be a better random allocator than all the algorithms in the world. Use the last few digits; reverse a couple in the middle, and you had a random number. The disk drive was not a huge data device in those early days [unless the application was specific]. It was simply an expedient way to control the elements that were difficult for a machine room to manage.

Burniece: So, how long did it take to get that all sorted out and working well? How long did that disk operating system for the 360 take before it was solid, basically bulletproof, and that system was really working well? Did that take many years?

Allan: No. Let's see, I'm trying to think back here. It [BOS] was probably running by around about September, October, on the first machines. I spent a lot of 1965 on planes visiting other centers [for DOS]. They were writing compilers and the Linkage Editor was loading their output. If they had any problems, I was the Endicott guy that got sent because one, I was finished, and two, I usually got the problem solved in some way. DOS was definitely running in '65. It wasn't a big operating system. First of all, the minimum machine was something like 4k, or 16k, very small memory. The software had to be written to fit the smallest machine. The 360 was a base register architecture and you had banks of

memory to address by loading a register with the base address of a bank. If a machine was large enough, the Linkage Editor brought everything into memory at load time. The overlays weren't on disk, they were now in memory, and I'd change the base address to use the next overlay. Everything ran out of a large machine's memory and that got me audited.

NOTE: BOS was a project to cover the low end because OS/360 could not support the small hardware configurations. BOS morphed into DOS and TOS when it was realized that OS/360 could not make its delivery dates. BOS was an Endicott-only effort, and DOS/TOS had to include APIs to integrate COBOL, FORTRAN, and utility software which was being developed in the other laboratories.

Poughkeepsie reckoned I was leaving stuff out of the Linkage Editor because it couldn't run that fast, unless I had. Their program was chunk, chunk, chunk to process cards and mine was going brrt at full speed. Poughkeepsie had written code for the minimum machine, so they would constantly be accessing disk to get the next overlay, versus my overlays being all in memory. What put me in all of these places where compilers were having problems was that programs were failing, because they were thrashing the disk. A disk error was pretty normal and it wasn't always detected, so the program would fail. On a couple of the compilers we went through and redid the compiler to be resident in memory and the problems went away. Issues associated with disk errors from beating the disk pretty hard were a problem. The error detection was good but not tremendous.

Burniece: Now, taking that just a little further. The 2311 problem that you found where you had a read error that shouldn't have been a read error because they couldn't search correctly, was that also in the 1311? The 2311 was basically a bigger version of a 1311, right?

Allan: I didn't have a 1311. I never had one in the field and 1311 was fixed sector, like the 305.

Burniece: But whatever it was - it was solved by having a double Address Mark.

Allan: Yes, it was a pretty simple fix.

Burniece: Many years later in a 2314, after it had been in the field for a while.

Allan: Yes.

Burniece: So, that probably was an embedded problem all the way through those early disk pack drives?

Allan: Yes, it was. Looking back with knowledge, it was a problem, which only happened under heavy load activity, a Sort for example.

Burniece: Sure. Sure.

Allan: Or loading software and thrashing software. Interesting days, those.

Burniece: Tell us, before we go back to Australia again, about the people in Endicott. Who were the people that you worked with, who you really enjoyed working with, really admired.

Allan: Frank Sebastian probably had the most influence on me. He was my first exposure to a development personality. How do you structure something for x number of people to do, which is bigger than x number of people, knowing that you're going to get new people who don't know anything coming behind you? How do you divide that up? How do you apportion the responsibilities to people? How do you convince so many different personalities, who want to do things their way, [to] agree on doing it a single way? And when one of them is failing, how do you cover for him for the project, and knowing that he's a good guy, still not destroy him?

Frank handled all of that because, as I said, there were four of us and then it just started to go like this [hand gesture straight up]. As more people came in, as they were hiring people, we had programming schools. One of the best guys that I ever worked with had started as a tech on the manufacturing line and he'd moved his way up and eventually worked into a programming position. Bob Simonik was one of those guys you could depend on. Was he brilliant? No. Did he produce code that worked? Yes. Was it brilliant code? No. Was it a hundred percent? Yes. On the other hand, my cubicle mate was a brilliant programmer, ninety-eight percent solution. He worked quickly and then somewhere you'd come across a problem that was inherent and couldn't be fixed without a total rewrite. Give me a Bob any day. So, you had these issues - Frank really handled all of that, with no prior experience.

Burniece: What was his background, Frank's background?

Allan: Frank? Irish, been to college, came to IBM and worked on the 1401, that was it. He was—

Burniece: As a programmer, as a software guy?

Allan: As a programmer, and this was his first managerial job. He was given this.

Burniece: Now, some of these insights that Frank was developing on how to structure a project, how to assign pieces of code to people that could be put together and worked and be picked up by subsequent people etc. This sounds like the kind of things that eventually got built into the “Mythical Man-Month” book. Was he an influence on that?

Allan: No. Poughkeepsie and Endicott were not close. In fact, the antipathy between the two labs was blatant. 1401 was the cash cow. It was seen by Endicott as having funded all of the 'disasters' in Poughkeepsie. Poughkeepsie viewed itself as the brains and those were idiots at Endicott. I really think the reason I wound up with the job of full liaison on the architecture team between Endicott and Poughkeepsie is that it didn’t bother me. I was a stranger to both places.

The meetings would always start with the Poughkeepsie guys dumping on Endicott, because I was from Endicott, and I'd agree, "it’s got lousy weather". Then we’d start a meeting without having a fight first. The 1401 was the success of Endicott, but the 1410 follow on was done in Poughkeepsie because the lab had more political clout. You had all of these simmering enmities that were personal, departmental and lab wide, and I was in the middle. We’re supposed to be collaborating on the new 360 but though the architecture books were good, they didn’t cover everything. We had to agree on how to approach certain implementations. I kept coming back with agreements to disagree because the goals for [OS/360] and DOS were different in terms of the size of the operating system. Fred Brook’s book was based on [OS/360], where things were just out of control. There was so much to do because they wanted to do everything in the first release. [BOS/DOS] was a different philosophy. Frank started with "what can we do?"

Burniece: Okay.

Allan: What can we get done? And when it’s done, we can improve on it. The entire process on [BOS] was to get it running. Pretty primitive, but it ran. We could load cards and get them into memory with the Linkage Editor and execute. Lou Stabile got his Job Control language going and that grew and expanded little by little. We started with a running base and added to that [for DOS], whereas Poughkeepsie wanted everything in one, grand, glorious, first release so they never got anything to run. It was all in turmoil. And that was really the difference between the two.

Burniece: So, it sounds like Brooks didn’t necessarily learn from Frank. But Frank actually was out ahead of him in terms of how to structure projects like this.

Allan: Fred Brooks was not an implementer. He was a thinker; he wasn’t in the trenches doing the work. The problem was that the Poughkeepsie structure had the goal of having the pre-eminent operating system and wanted everything to come together at once. Poughkeepsie never really got the nucleus going until everything was working. Frank’s approach was "I don’t know what the hell we’re doing here,

but let's get this going first, okay?" The three of us, with Ray's help, worked that way and as we added people we could gild this part of the lily here, because people are having trouble with that. That's how the Job Control language developed; we kept adding more functions but in an easy to understand manner, while the old stuff was still running. DOS grew from a very simplistic base. It wasn't very sophisticated but it was the only thing people had to run on 360s. Long before [OS/360] was out, every 360 machine in the world was running DOS or TOS (we had it running on tape, too).

Burniece: So, the bottom line is that you managed to get this DOS done in less than a year.

Allan: Oh yeah.

Burniece: With a team of four guys—kind of a step at a time approach.

Allan: Well, no four plus. It was—

Burniece: And, in the meantime, the [full OS/360] operating system itself was still way behind schedule and not working well.

Allan: Yes.

Burniece: And not a very well managed project at that point.

Allan: That's why I was audited; they were looking for scapegoats. I'm sure somebody was beating them up saying "they're running why aren't you?" Poughkeepsie faced huge management pressures and it was a time of turmoil. It really was, inside the company. In fact, IBM ran out of money. Did you know that?

Burniece: No, that I wasn't aware of. They literally ran out of money?

Allan: Yes, IBM was building hardware and faced a huge working capital requirement, plus was hiring programmers up the kazoo. The cost of development was zooming out of sight and customers were not paying for machines that could not run [OS/360]. For a while there, Prudential was paying us.

Burniece: Now, that's interesting piece of information I didn't know. Tell me just a little bit more about this.

Allan: I don't know a whole lot more.

<Laughter>

Burniece: Was that public knowledge at the time? IBM's a public company. It must have been.

Allan: I don't know what happened. All I know is that at one point in time, somebody came around with paychecks and said courtesy of Prudential I went blank. What do you mean? "We ran out of money and had to go to Prudential for some more."

<Laughter>

Burniece: Wow.

Allan: We didn't know a whole lot more than that at the time, but we knew that for several weeks, if not months, Prudential was bankrolling our jobs. I'm assuming it was a corporate loan or borrowing against future revenue. Remember, the lease base was still there producing revenue but if you sold a 360 with an operating system, and you can only install a 360 with DOS, you weren't meeting your full commitments to the customer. The lease revenue on the new 360s was down. You were building these things and getting them out there, using more money than the early lease revenue. Those were very difficult times for IBM.

Burniece: how long did it take to get that resolved? Was that resolved by '66, something like that, or earlier?

Allan: I think the cash probably started to roll in in the '68 timeframe.

Burniece: Okay. So, how long were you actually in Endicott during that tenure? Were you there three, four years, or less?

Allan: Let's see, April '64, eighteen months, extended six months, so we got back to Australia in mid '66. I became headquarters support for the Australian New Zealand region - the General Motors Holden problem and a few other crises. One time I was sent to Singapore for a problem that reached a critical stage. They were part of the Japan region and did not want to take the problem to the U.S. It was a problem I'd seen already but I had to stay for two weeks. Henry Oh told me that he could not have me return to Australia immediately or the Japanese would lose face. For the next two weeks, I was in Singapore with the salesman for Shell and the two systems guys. That really paid off big later on because they introduced me to the Asian way of thinking.

I'd looked at Henry and asked "Why not, I've got a wife that's pregnant. I want to get home." He talked me through it, gave me to his guys and said "educate this guy, we might get him back here again some time." Amongst other things I went to a birthday party for the grandmother of one of them and it was eye opening. I had no concept of Asia or Asian thinking and styles.

Burniece: So, you're saying you were basically kept for an extra couple weeks, so it didn't look like you solved it too quickly, because otherwise they'd lose face.

Allan: It was scheduled for two weeks because that's what they thought it would take. Shell had been screaming for months and the Japan region had been in there without fixing it. It was an odd problem; it was another Sort problem and, in fact, it turned out to be AMP's Sort problem on different hardware. It must have been the same programmer who wrote both. If the tape reel was full and some circumstance arose at end of tape, it would blow up. If there wasn't a full reel, Sort worked fine. At my first visit I suggested they size the reels in half and try it again, and it worked. Bingo! Pure luck.

I wasn't allowed to go home and Henry did explain it to me. He told his staff "take this basically ignorant White man out and educate him in the ways of the world. And they did. The American club had been closed to those of color up until two years before. It had been colonial 'white' Singapore and everything was opening up. The place for everybody to be was the American club, so prestige in Singapore revolved around being a member of the American club. This turmoil was happening in Singapore at the time I was there and the guys walked me through what it had been like before and what was changing. I was like a sponge. It was all foreign and new.

There was one Canadian who was thicker than a post. He wore long sleeves and long pants and in Singapore short sleeves and shorts is de rigueur. If you're in long pants or long sleeves, you're a visitor and the moneylenders descend on you to exchange the money. This guy would walk down the street besieged by Indians who were in the exchange game every day of his life in Singapore. He was not willing to give up that vestige of Canada and adapt to the local environment.

Burniece: So, how long a period did you then stay back in the Asian area.

Allan: I was only there for two weeks, but I actually …

Burniece: I mean that whole tenure.

Allan: Oh, the whole tenure? I went back to Australia in mid '66 and around late '67; I received a telex with an offer to come back to Endicott.

Burniece: Okay, so you were there about a year or so.

Allan: A year and half, yeah.

Burniece: And why were you called back to Endicott the second time?

Allan: IBM had decided to do a multiprocessor system. MTBF on 360 was pitiful and on a multiprocessor system, if one failed, the other could carry on. The hardware was being done in Hursley [England] as the 40 IMS; the model 360/40 had come from Hursley. The software to run the multiprocessor was at Endicott. I'd fly over to Hursley and we'd walk through what could be done. If someone had given us an 1108 manual, we would have had no trouble doing it, but we didn't know anybody had done a multiprocessor, so we started from scratch. It had been Frank's project and, when he'd been promoted, he called me in to take over.

A big part of that job turned out to be hardware assistance for the software. We'd come up with what we wanted and I'd go over to Hursley and sit down with the engineers. The first day I'd go through all the things we wanted and they'd come back two days later and tell me what I could have. I had to be walked through why they couldn't do what I wanted and in machine terms explain what they could do. I'd come back to Endicott and we'd implement with what they'd given us. It was pretty terrific; one of the guys that worked for me asked for something from Hursley, which cut the swap time from milliseconds to microseconds. It was a brilliant step forward in terms of shifting activity from machine to machine.

That's probably what killed the project, because we exceeded our performance specs by fifteen to one. The Intermediate Multiprocessing System outperformed a model 65/65/50 combination, which was what they had for larger systems, at a fraction of the cost. Running several times faster at a fraction of the cost was not a viable solution.

Burniece: So, you're saying they buried the program because you outshone the high-end system?

Allan: I don't think it was that simple, although it was probably a factor. The other factor was the System 370. IBM had figured out that with higher integration, they could get the MTBFs they wanted on a single machine. It was a combination of things but it was a very bitter pill for all of us, because we'd reached the point of final integration and everybody on my team was going over to England for six weeks to get our software up and running. I'd overlap the first and last week with each team to manage the transitions and we were at the peak of success. Suddenly, we got a cancellation letter that says IMS is dead and we took it very badly. We were not told that it was cheaper to build a single machine with higher reliability, by using more integration, but it would not have made the pill easier to swallow anyway. That put me on the 370 architecture team.

Burniece: So, at that point, they moved you to 370. This is when?

Allan: This is 1968.

Burniece: Okay.

Allan: And that was a disaster. I got along well with the engineers. Software would ask for something, and Engineering would come back with "we can't do that, but we can do this". The votes were going the engineering way, 6:4, 6:4, 6:4... and I was called into explain why I was voting with Engineering when they had not given what Software asked for. My "well, they gave us something we could use" answer did not set well, because Engineering was a different political empire. Software was there to get what Software wanted, not what Engineering could give us.

IBM had transitioned to the point, where there was really nothing much for everybody to do, so the political structure became more important. Instead of not having any politics because there hadn't been enough of us in the beginning, politics became the modus operandii. There was no chance for advancement by succeeding at something because everything was already done. It became "make work projects" and "lobbying for position". It was a new empire and a new set of rules, and it was kind of frustrating. That's when I transferred to San Jose and I lasted no time at all there.

Burniece: You say you transferred. Did you ask for a transfer? Or somebody pulled you to San Jose?

Allan: I wanted out. I'd been to San Jose in '64 and '65 when they were doing compilers in Building 26, which is where I first met Irv Tjomsland. He was back east for something and I asked, "You got any jobs out there?" He said, "Yeah" and got me the transfer. It was Irv who took me into ISS years later. I've always had this thing in life with various incidents. People popped up now and then that would just come and say, "Hey, let's move over here. Let's do this."

Burniece: That's kind of the way it works. So when did you leave for San Jose? What year?

Allan: I left for San Jose in early '69, and left IBM later in the year. The goal there was AM1 (Access Method One) to replace everything that did I/O to disk with this grand and glorious, wonderful, single thrust. It was constantly being gilded and gilded and gilded, and it could never make it because of the daunting level of complexity. I was given the job of porting the existing random access support to run under AM1. We did it outside the umbrella of AM1. We put every function into a task table and used that to generate the required AM1 commands. We interpreted the request by the old software and called on the AM1 equivalent. Very simple and easy but my managers didn't want anything outside of AM1; the function had to be integrated inside AM1.

IBM managers were not supposed to interfere with technical people, when it came to implementing a project. I'd been in second-line management at Endicott after the 40/IMS, and I'm not a good priest. I really had a tough time dealing with employee problems when I was thinking, "How can you have this problem? You're smart. How can you be overdrawn with all this income you've got?" I'm not good at that sympathy part. <laughter> My wife's best friend's husband is a parson and if I'd a wanted to be a priest, I would have put my collar on backwards like he does. The transfer to San Jose achieved two things. It got me to the West Coast, which my wife really, really wanted, and it took me back into a staff position.

Burniece: Just, to finish that part of the story, this Access Method 1, was that a fundamentally flawed idea or they just couldn't execute the idea?

Allan: It was fundamentally flawed from the beginning - "let's replace what we've got, with one singular thrust that does everything and let's improve it beyond that". Everything became gilded before it was implemented. It wasn't a case of let's come up with a more efficient Sequential, knowing that you also have to do Random, so leave the hooks there. It was one of these glorious empires that died of its own weight, because you can't anticipate everything in advance. When you keep changing the specs to cover all of these things you haven't done yet, you never get anything done.

Burniece: So it went away?

Allan: It went away. Sometime after I left, it went away.

Burniece: And they went back to the DOS system you had before or something else?

Allan: No, no. IBM kept the old drivers and came up with new individual driver improvements, so they dropped AM1. AM1 was San Jose software's pitch in the contest for projects and money and funding, etc. Someone had come up with this idea of a universal driver and it appealed enough to get budgeted. IBM eventually went back to just improving what they already had.

Burniece: So, take us to the end of your tenure at IBM. How much longer did you stay there and when was it that you actually ended employment?

Allan: It was pretty difficult working for the AM1 management, because they did not want things to be done in any other manner but theirs. Things were being done for appearances, rather than effectiveness, so when a start-up approached me with the offer of a job, I said, "Yeah, sounds good."

Burniece: Start-up being ISS?

Allan: No. Computer Synectics.

Burniece: Oh, another one. Okay.

Allan: The 360 had problems and IBM San Jose had a team of guys led by Dudley Warner, to develop a monitoring device, basically an oscilloscope with counters. You could hang probes on the system, get the numbers back, data reduce them, and figure out what the problem was. IBM chose not to take it any further and Dudley said, "Can I have it? I'd like to start a company." IBM said, "Sure", so Dudley left with IBM's blessing to pursue the SUM (System Utilization Monitor). At the time I joined, it seemed to be a real growing concern. I wound up doing a lot of consulting with SUM because it was a difficult hardware sale. Buyers had to open up the covers of their machines, hang the probes to accumulate the data, and then they had to interpret the data.

Consulting became the income for the company to keep it rolling. Phillipe Yaconelli came in after the, hardly the politically correct term, vulture capitalists for EM&M (Electronic Memories & Magnetics) / Caelus relieved him as CEO. Many years later, EM&M sold Caelus to ISS, after it had been acquired by Univac. Phillipe came in with a completely different set of ideas but it was pretty clear the company wasn't going to go very far. That's when I got the call from Irv Tjomsland about Environ 1 and why I wound up at ISS.

Burniece: Okay. So your time in that startup was relatively short?

Allan: I don't know exactly how long, but it wasn't that long.

Burniece: So when did you start at ISS?

Allan: Probably in the '70s, '71 time frame.

Burniece: And how far along was that company by then?

Allan: Oh, it was booming. Telex was selling everything ISS could build. The 714 was a huge success. ISS was the only PCM with a disk controller for a long time. The other PCMs were device compatible and ISS was channel compatible, probably because some of the "Dirty Dozen" were controller people.

Burniece: Let's just fill in a little history, for the record here. Talk a little bit about the, quote, "Dirty Dozen" and what that was all about, how ISS got formed and where they came from.

Allan: This is all secondhand. The hydraulic actuators on the 2314 were a real pain and a new project had started inside the lab to do a voice coil. Some of the "Dirty Dozen" were involved in this. IBM decided to stick with the hydraulics and moved in the direction of what became the 3330. This left a few bitter engineers and my understanding is that this was the catalyst, which led to so many leaving together. It was a pretty well-rounded team - you had disk drive guys and controller guys. Frank Sordello was pretty bitter, because his last patent was not pursued by IBM after he left. He wanted that one.

Burniece: That's a good story, too. Tell that story in a little bit detail. What are you talking about? What number did he have to reach to do this?

Allan: I think it was fifty.

Burniece: So fifty patents and you're in a special club of IBM.

Allan: Frank's fiftieth patent was in process when he left, and he didn't get it. ISS's biggest failing in engineering was that excellence was judged by the number of patents you had on a project. We never had a power supply that worked well because every one was different than the last generation because some new wrinkle was invented to get a patent on it.

Burniece: So the culture there was invent as much as possible and then patent it?

Allan: Yes, invention-driven. Patent it. Although the products were great, if there hadn't been that focus on originality as well, I think they would have been even better.

Burniece: Focus on what they already had and extend it. From an evolution standpoint, they basically reinvented it all the time.

Allan: They evolved but with a twist. The twists were what got us into trouble, because the twists were always to save a few gates, save a few of this, save a few of that, and it would have been cheaper in the long run to have used those extra gates and parts. This is from the outside looking in and it may not be a valid judgment but I sure felt like that, when we had the same problems popping up generation after generation.

Burniece: What year did “the Dirty Dozen" leave IBM and start ISS?

Allan: That was probably late '60s.

Burniece: Right. So the company was a couple of years old when you got there? But it was booming at that time.

Allan: Oh, yes. Telex was selling everything they had. ISS never sold direct. They always sold through somebody, and Telex was their sales arm.

Burniece: And the secret was they had a true-plug compatible controller?

Allan: Yes. The disk drive was not plug compatible but there was a good controller. You could have ISS singles. IBM sold you nine in a box. Only eight were useful; the ninth was a spare. It wasn't left alone, because if you left it alone, it might not be working when it was needed, so it was constantly running diagnostics. IBM developed a diagnostic culture, because of that ninth drive, and that knowledge was also in the 7830. That paid off because ISS's 7330 was the first disk drive to have a microcomputer and a display panel (a telephone pad with some extra lights) and ran its own diagnostics. The engineer no longer had to use an oscilloscope. The diagnostics knowledge that came out of the IBM 2314 Controller was translated to the ISS controller for the 714 and the double density 715, before being put into the disk drive itself. A stream of knowledge that wouldn't have happened in the disk drive, if the controller people hadn't been there.

Burniece: That was a number of years later, though. You're talking about more like the late '70s, early '80s when that happened, right?

Allan: No, younger than that. When did the 3330 come out?

Burniece: The 3330 came out in 1971.

Allan: I'm talking '73, about the time we had it in. The idea to put diagnostics in the drive would have come out at the time the 3330 was announced and ISS had started developing it before then, because we had inside information. Guys were leaving IBM and hiring in. The fact that the ninth machine had to be working when it was used as a spare, translated downstream, and eventually wound up in every disk drive at ISS. That was a straight line.

Burniece: So what was your role at ISS? What were you doing there?

Allan: After Environ 1 was sold, it was kind of like, why don't you stay? I did planning work for disk drives and controllers and figuring out where things were going, doing industry intelligence, tracking down what IBM was doing. Employees would come in and I would debrief them to find out what they knew. I'd

sit in on some of the discussions between the engineers and when Itel acquired ISS, that side of things evolved further. The disk planning guys were let go but a lot of my intelligence gathering was being done at industry events, like GUIDE and SHARE.

One of the benefits in IBM Development was that you'd get R&R once a year, possibly twice a year. You would be rewarded by being sent to one of the IBM conferences, either as a speaker or attending, and you were always lionized by the customers. If you admire what an engineer's doing, he will tell you more, because for starters, nobody ever asks them. I'd meet an engineer talking to a bunch of customers at the social events by simply joining the group. I'd find out what he was doing and if I was not interested, move on, but then you'd strike gold. You'd find a guy from San Jose, and you'd steer the conversation in a group of people about this and that, and if you hit the right question, he'd just waffle on. The customers would start drifting away, because now he's talking techie stuff they did not understand; then I'd wind up going to dinner with him and listened.

Burniece: So this is the beginning of your networking career that eventually got you into a full time career with the interface work.

Allan: Yes...

Burniece: That you interfaced with.

Allan: I extended that to visiting OEMs and companies, which built disk drives not competitive with us. When Saburo Adachi came to town to create a business for Fujitsu, he and I would get together once a month. He was learning the game and I was finding out about the disk drive business in Japan, where they were hoping to go. That was the beginning of finding interesting people and letting them talk to me, about doing things differently at ISS. Could we learn from them? Sometimes we could.

Burniece: So during your tenure at ISS, which was how long?

Allan: Well, it was a bounce-around. It was ISS; then it was Itel.

Burniece: Yeah, I know it got bought by multiple people.

Allan: Yeah. Oh, gosh - some where probably from '70, until the time CDC bought it in the early '80s.

Burniece: All right. So basically the whole '70s you were at ISS.

Allan: Not quite. I went back to Australia for eighteen months as a branch manager in Canberra.

Burniece: Working for who?

Allan: Univac.

Burniece: Oh, so when Univac bought ISS you joined the Univac sales force or was this before Univac bought ISS?

Allan: No. Univac bought ISS but we were still in the IBM PCM business and we were now in the Univac business as well, so we had visitors coming from all over the world. The Director of Australia Sales came in on one trip with a customer and we did the dog and pony routine for them. That evening he asked, "Do you want to come back home?" I said, "Oh, my wife would love to." "Are you serious?" "My wife's serious." He came back with an offer. Univac had just made some big sales in Canberra and they were starting a new office, so that was the job offer.

Burniece: So you left ISS and joined Univac.

Allan: Univac Australia in 1975. My timing was terrible. There was a constitutional crisis, the Prime Minister was fired and the opposition leader appointed as the new one, so a whole change of government. The previous administration had been spending money hand over fist on automation and it became a case of "No more money." I flew out with a 1976 quota of several million dollars, a small chunk from a huge pot, and now it was several million above zero. My territory was re-aligned and I was given TAA and a couple of other companies in Melbourne. TAA was the government owned airline but it was outside some of the restrictions. I'd fly down to Melbourne on Tuesday and come back on Friday morning so my territory did have promise. TAA was good and I did make quota but it didn't look optimistic for the next few years. So, when ISS offered me my job back, we returned.

Burniece: So you went back to ISS. Okay.

Allan: I came back to ISS, and that was hard. I had no idea how much technology could change in eighteen months. I was at sea for weeks, if not months, just trying to get back on that treadmill of constant change. I'd been totally out of it in the antipodes and I was scrambling to catch up. It was a very sobering experience. You can’t step out of this game and step back in because it's ….

Burniece: Now, what time frame was that?

Allan: That would have been - let's see - I came back in late '76, early '77.

Burniece: Let's go right up to the point of when you decided to go work the standards thing and then let's take a break.

Allan: Okay.

Burniece: How much longer was it until you started moving towards the standards arena?

Allan: NBS and GSA really pushed to get the IBM BMC (Block Multiplexer Channel) as an American National Standard. They failed.

Burniece: What time frame was that?

Allan: This was early '70s. So, we're talking '74—

Burniece: Oh, it was in the middle '70s. Okay.

Allan: Middle '70s. The non-IBM manufacturers didn't want IBM BMCs, so it didn't get to be an industry standard, but NBS did make it a FIPS (Federal Information Processing Standard) standard. GSA was not happy with that because nobody had to comply, unless it was a FIPS compliant sale, i.e. an IBM machine. GSA decided to go for disk drives and Del Shoemaker of the GSA called an X3T9.3 meeting.

Burniece: Was that the first of those meetings?

Allan: The announcement was to investigate disk interfaces and the first meeting was being held in Irvine. Nobody from Univac wanted to fly out so I was a surrogate sent down. I remember the date very well; it was my wedding anniversary, so let's just say I wasn't very popular. <laughter>

Burniece: At least with the person you'd married**.**

Allan: I'm sitting in Irvine instead of celebrating my wedding day. It was a two-day meeting and Del threw out this huge pair of books and said, "We've had this study done by DataPro and the conclusion from this study is that we're going to make the 3330 interface a standard." Del went on to talk about all the reasons why. There were two interfaces at this time on the 3330, the original CUDI (Control Unit Disk Interface) and the DDC (Direct Data Control), which was a B-box or head of string interface. The analog

CUDI didn't have such great performance on long cables, so they put the head of string next to the disk drives and the data-separator DDC cable over to the controller.

I've got two columns, right? And I'm writing questions in each column, because I didn't know what Del's comments applied to. Some of them seemed to be CUDI and others were DDC. After a soliloquy of about two hours, we're coming up to lunchtime and Del says, "Do we have any questions?" I asked, "Which interface are we talking about here?" "Where have you been? We've been talking about the 3330 interface all morning." "Yes, but there's two interfaces. Which one?" "There's two?" <laughter> Oh, boy.

Burniece: Really, he didn't quite get it at that point?

Allan: No. He didn't know there were two. He never read the study. He knew what he wanted. The study was just …

Burniece: So he had not studied it. He basically just brought the books and then just said we're going to standardize this.

Allan: Yes. He didn't have a clue what he was talking about. The study had been done to provide justification. It was fairy floss and that's when SMD became the first project because ...

Burniece: We're going to table that and come back to that later.

Allan: But that's why.

Burniece: All right. Before we take a break, though, anything else prior to that that you want to talk about - good stories or good people or whatever in the whole trek through IBM to ISS back to Univac etc? Any other great stories you want to fill in there?

Allan: I could probably sit here and come up with a lot of stuff but no, I think - when we look at this later, you probably can say, "Hey, fill in this gap." But, yeah, you've heard some war stories....

Burniece: Let's take a break.

Burniece: We are back and we're going to pick up where we left off. Dal, you were just talking about how Del Shoemaker of the GSA said we're going to standardize the IBM 3330 interface and you asked

which one. He looked at you like what are you talking about, which one? How did that revelation that there were two suddenly morph into the CDC SMD interface becoming the one you picked?

Allan: You can't stop anything but you can redirect it. ISS was not compatible with IBM disk drives, so I had to protect our entire business. Dick Whitcomb from Itel was there. Itel was still around; they had sold ISS to Univac but were still in the leasing game. Dick thought it was a great idea, until I took him off in the corner. I had worked for him when I was at Itel and told him "this will kill us, because we're not disk drive compatible, we're channel compatible." Dick made it very clear, "Okay. We've got to stop this thing." "We can't kill it, we've got to do something else." During lunch, I talked to everybody who was at that meeting. I went from table to table and just chatted to find out where they stood.

Burniece: Now, roughly, how many people and companies would that represent?

Allan: There were probably only twenty people but it was a diverse group. We had a guy from Livermore Labs who later turned up at Quantum. It was a mixed bag. Memorex was there, of course. They thought it was great, because they were compatible at the hardware level.

Burniece: Was IBM there?

Allan: No. I don't remember IBM at all.

Burniece: It was their drive being talked about to standardize but they weren't there....

Allan: IBM did not encourage anybody to make their products a standard. They did not like that.

Burniece: Was CDC there?

Allan: Yes. CDC was there.

Burniece: And who would that have been?

Allan: That would have been Gene.

Burniece: Gene?

Allan: I'm pretty sure it was Gene. Because I'd been running around industry to meet as many people as I could to find out stuff. I knew that the CDC interface was the defacto industry interface. If you walked away from IBM, people used SMD, and by the way, if you were looking for another choice, you had SMD. SMD was the only product out there. It owned the market in every field except the one GSA was pursuing, which was mainframes. That became my candidate, so I led the lobby for one and a half days until I wore Del down to the point where we all voted n:1. Actually it was n:2, because NBS voted with Del.

Burniece: So Del and the guy from NBS still wanted the 3330?

Allan: Oh, yes. They didn't care what we wanted. They just wanted what they wanted. They weren't listening. They came in with a mission. They knew what they wanted and they left without it. <laughter>

Burniece: I love it. So, at the point that you had lobbied your way out of the 3330 interface, in part because it would have damaged your employer …

Allan: Yes.

Burniece: ... and lobbied for the SMD interface, which at the time was the only other really well-defined and well-used interface, after being developed by CDC…

Allan: And manufactured, by the way.

Burniece: Yes and manufactured by CDC. Did Del and the guy from NBS agree at that point, or did they …

Allan: No. Remember, this is a volunteer committee.

Burniece: After this loss of vote, did they go along with it?

Allan: It's a democracy. Yes. They had no choice, right? <laughter> Ampex was definitely at that meeting because the Ampex guy volunteered the specification to work off.

Burniece: CDC refused to give you the spec at one point, right?

Allan: Right, but the Ampex guy said, "That's okay. We'll give you our spec." Now, 'our spec', by the way, was a redacted copy with black blocks everywhere “CDC” was written. <laughter>

Burniece: So you were truly plug compatible across the industry. There was only one manual. I'm forgetting Gene's last name.

Allan: Milligan.

Burniece: Yeah, Gene Milligan. What was Gene's reaction to all of this?

Allan: He was not too happy, because CDC didn't want a standard either, but he didn't want **...**

Burniece: He voted for it?

Allan: Well, he wanted the 3330 interface even less.

Burniece: All right. So he had voted for it, despite the fact that his company didn't particularly want it either.

Allan: Right. Gene was a clever man and he was there to observe, not to decide anything.

Burniece: Talk about him a little bit. I knew him also but just give his background, because he was a very fascinating guy.

Allan: Gene was one of those characters who has the ability to read a specification, see something on page 7, log it away. When it pops up again as an implication on page 93, realize that it's different. So …

Burniece: So he almost had a photographic memory of details?

Allan: No. It only lasted the length of time that he read the specification. His comments were terrific, because he had to remind himself of what he found. His comments were always very specific, referenced both places, and said why there was a difference. He didn't try to resolve them. That wasn't his style. He either didn't care or it wasn't his choice. Gene understood how the political structures worked and he never wanted to vote for anything negative. If he wanted a negative result, he put forward a positive motion and voted against it. Great guy - two glasses of red wine with dinner.

Burniece: Say that again, he didn't want a negative result, so he'd propose a positive and then vote against it.

Allan: Yes.

Burniece: So the negative result would be …

Allan: A down would be his win.

Burniece: The de facto?

Allan: I don't know whether it was a Robert's Rules of Order or a Gene thing or what it was, but that still lives on today. You will not see anything in T10 or T11 that is a negative motion.

Burniece: It's always a positive motion.

Allan: We will—we should do this. No, it failed. All right. That was …

Burniece: That was Gene.

Allan: Gene really enjoyed life. He was a heavy smoker and had at least two glasses of red wine with every dinner but didn't care what it was. He just liked red wine. Very principled, very strong-minded and we butted heads regularly. One of the great things is, we could butt heads all day and then go to dinner, and that was really nice. We could disagree philosophically and still find things to enjoy chatting about. Gene was very family oriented and when he had his heart attack, followed by open-heart surgery, we thought wow, we're going to lose Gene. He was back in like five weeks. The doctors had him up and running around the ward, he told me, the next day. His life had changed, with the same determination he showed elsewhere. He went for a power walk at lunchtime every day. It was like a two-mile sprint to keep up with the man. That's what the doctors had told him to do and he pursued it. He gave up all the things he should have, took the regimen to heart, and lived a few more years because of it. Gene was a good guy to work with.

Burniece: Yeah, he was a great guy and I remember him well. I think he was CTO or Chief of Engineering at NCR's disk drive operation, when CDC picked them up in the early 70s.

Allan: Yes. He had been acquired.

Burniece: And he ended up in Oklahoma City when CDC acquired Honeywell's disk drive operation there and formed MPI in 1975. He was there for the rest of his career. He became the MPI standards guy after that first meeting, which goes from a vote for IBM's to CDC's interface. Then he's not going to give the spec but Ampex gives them the CDC spec. That is quite a story …

Allan: Yep.

Burniece: At that point, is he on the bandwagon, or is he fighting it?

Allan: I don't remember Gene being anything in those early periods; he was laying back in the weeds. I don't remember any of us knowing where to go. We lost the guy from NBS almost immediately. He had cancer, so Gary Robinson, who had been in the standards game for Honeywell on tapes and the floppy disk, became the leader. Gary was the chairman, and Gene had known him from previous escapades. We were ticking along on the SMD and I checked in for a meeting somewhere in Colorado. There was a message at the desk, which said, "I won't be there, you're the vice-chairman." Gary had run into an employment issue, wasn't getting travel permission, and that's how I became a vice-chairman, a field command. <laughter> I opened the meeting the next day, but it was also the beginning of the first …

Burniece: Was the chairman still Del Shoemaker at that point?

Allan: Del was chair of X3T9. As the master committee, he had several committees under him of which most had petered out. They were doing pretty much useless things and they'd gone nowhere. X3T9.3 was a new birth, so to speak, and he stayed out of it, which left Gary as the one in charge. We worked quite well together. I did the technical side and he did the administrative side. We succeeded on the SMD and Gene did come on board. Somewhere along the line, it was decided by CDC that, hey, it's not such a bad thing.

Burniece: It turned out to be a blessing in disguise, actually.

Allan: It did. And it did something else, too. It said that there is benefit to being a standard and the standard will weed out the differences that can occur between implementations. Customers liked the fact it was a little more rigid, plus the things that were ambiguous or obscure in the original specification were clarified. That was a big plus, because engineers are more than capable of interpreting the same statement five different ways if you let them.

Burniece: Right.

Allan: A big goal of that first project was finding out how to write a standard. We were in the learning mode, Gary had done tape but disk is a little more difficult, and we were learning how to express things in a precise manner.

Burniece: Now, just to kind of backfill a little of history, the CDC 9760 SMD was introduced in 1973, and that interface was already in it.

Allan: Yes.

Burniece: You didn't get it published as a full spec until, like, 1982.

Allan: Oh, yes. Because we started it in 1977 **…**

Burniece: So it was about nine more years.

Allan: Yes.

Burniece: Did you fundamentally change it, or did you just really kind of clean it up and get it to the point where it was easily interpreted and could be easily reproduced?

Allan: There was a lot to explain. We finished the standard much earlier, then the approval cycle came because once you've got it done, you've still got a huge number of months involved in just getting it through the process and getting it approved. We didn't take all that long. It just took a long time to cycle through to become formalized.

Burniece: Well, you didn't actually start on it until '78 or '79, I think.

Allan: We were working on that in June of '77.

Burniece: So it did take five years?

Allan: It took five years to hit the standard stage. We were done in two and had started on the 8" Rigid Disk Interface.

Burniece: So the voting process and the final approval process and all that took a substantial amount of time?

Allan: Yes. I don't want to go through that.

Burniece: The base work was done.

Allan: Yes. The rest was Gary's job.

Burniece: Now, in the world of standards, how unique was having a disk drive interface standard? Was that one of the first to be done?

Allan: It was the first hard disk. The floppy existed....

Burniece: Well, I know it was the first disk drive, but was it one of the first standards to be done in the computer world?

Allan: No.

Burniece: Was there already a number?

Allan: There'd been tape standards before. Remember, tape was an interchange medium, so there had been tape standards, and there had been a floppy standard.

Burniece: But most of these are at the media level, right?

Allan: Yes.

Burniece: Not at the drive interface level.

Allan: Not at the interface level, except floppy. So, this was new territory, and it wasn't embraced at first. It became fully embraced when the 8" drives started to appear. Everyone had a different interface. The RDI (Rigid Disk Interface) introduced the idea that the drive could do some thinking on its own, self-configure at power on, modify parameters and such. When that was finished it was a good specification but the market changed. The 5.25" high performance drive came along and it also had a bunch of

interfaces. You had Seagate's ST412HP reaction to ESDI, which was from Maxtor, and you had the Rigid Disk, which was the only really well defined one. Nobody wanted to do RDI because that was 'old' stuff and clashed with the 5.25" pioneer attitude. There was a lot of attitude.

Burniece: Who originally developed that [RDI]?

Allan: The standards committee.

Burniece: Yeah, but who did it come from? Who was the manufacturer?

Allan: It was a blend of multiple implementations. IMI had a lot to do with it. A German guy working for BASF did much of the reconfiguration at power-on. All of the 8" vendors contributed something and RDI didn't look like any of them.

Burniece: So it was sort of a hybrid of different implementations, rather than a clean top-down architecture.

Allan: No, no. It was a hybridization of what existed and Maxtor …

Burniece: Did it result in a published standard or …

Allan: Yes, it did.

Burniece: … did it go away?

Allan: No. It was published. It was built by a few people but the market just disappeared on it. Maxtor won the minds and hearts with ESDI but they couldn't get full participation. "Hey, this is Maxtor's interface." "Maxtor's a competitor." "How can I trust them not to do something that will screw me because of what I've done?" Skip Kilsdonk and Norm Zimmerman bailed me up after one of the meetings, and said, "Do you want to run this?" <laughter> Skip had a different goal than Norm.

Burniece: Well, they had actually started their own committee.

Allan: That's right. And we were all meeting. Maxtor was running into troubles because of this lack of trust and that's why they came to me. Norm wanted out of the job. He didn't care what happened; he

didn't want to keep running ESDI. Skip saw the advantage on the marketing side of having a separate entity, so we started the ESDI Committee.

Burniece: Well, Skip just filled in a couple of cracks. I've known Skip forever, he was the product manager for the SMD.

Allan: Yes.

Burniece: He saw the results of the SMD interface being published and standardized. Then he went to Maxtor. So, he had already seen what had happened.

Allan: Yep. But he didn't want it in the standards committee because he wanted it done quickly.

Burniece: Right.

Allan: That's why we started the ESDI committee, and when we were finished, we gave it to X3T9.2 to become a standard. We took the approach of a power on configuration a little further - a little bit more drive intelligence there. During that same time frame, X3T9.3 had started the IPI (Intelligent Peripheral Interface).

IPI started when I was at ISS. With the changing landscape in the plug-compatible world, we'd been acquired by Univac, and now we had to suit Univac. Univac had the 36-bit 1100 series and the 8-bit 9000 series that was bought from RCA, and many other variations as well. ISS was expected to build dozens of interfaces for the different computers they had and it was terrible. ISS wanted to come up with a single edge, a single sword. The 5046 controller we built could have byte channels or word channels in it, up to four you could mix and match. Univac never shared data between their 9000 and 1100 series, so it was never actually used that way. We didn't know that when we were designing the system, we were thinking about sharing test and diagnostics work between the 1100 and IBM 370, because we were still PCM. The 5046 did one-word reads and writes. We handled the read/update/write in the controller and we built in a Solid State Disk, which eventually became a full caching system. We were developing controllers to bring Univac into the twentieth century, because they always had to finagle things in software before.

Burniece: Just to back up a little bit here, the IPI (Intelligent Peripheral Interface) was one of two or three candidates that got started about the same time frame for an "intelligent drive interface," correct?

Allan: Yes.

Burniece: You basically had the controller processing inside the drive.

Allan: Yes.

Burniece: I started the ISI project at CDC before I left and SCSI came out later, plus a few others. What was the committee's attitude at that point in time towards moving to an intelligent interface and which one should it be? Was it pretty much chaos or was there a lot of agreement?

Allan: There was a lot of angst. ISS wasn't the only company struggling with companies, which had multiple machines to support. You had Burroughs and everybody else still in the game. The ISS proposal was something we had going in the controller group. It was for a single interface that could be a channel replacement, not just a controller replacement. The idea was to suck the capability of the channel into the controller and have that work with the CPU. Suddenly, the channel doesn't belong to the CPU any more. It belongs to the controller. IPI had provision for a lower level interface for the drive, called IPI-2. IPI-1 was physical, bits and pieces for control of the wires. IPI-3 was the channel solution.

Burniece: So in a certain sense, you're now inventing a new standard for what you said was the original gist of the 3330 world. <laughter> A simple drive interface and a channel interface.

Allan: That's right.

Burniece: Which one? Well, we'll do them both.

Allan: Yes. What you had on the IBM system was a B-box for strings and that was a device interface. With B-Box to the controller, as a controller interface, you had the controller to CPU channel interface. We wanted to suck the channel out of all of these CPUs, which were old-fashioned machines put it in the controller, just to get rid of that cable and all that function out of the CPU side. CPU engineers didn't like that, so right off the bat you can see there was opposition. The thing was, since IPI had the two tiers of simple device level, which used the same hardware, it was cheaper to build the channel interface in the controller. It took a very top-down approach; everything that the mainframe people wanted to do became part of the command language for IPI 3. It even pushed to the boundaries of doing some file system things as macro structures on behalf of the file system.

IPI became a platform. Univac really loved publicity, anything positive, so they fully supported the idea of me going to all these conferences, although the company backed off on committing to the ISS interface internally. "Too much investment in something we haven't done before." Univac was willing to take it public, because if it went public, they could buy OEM and never have to put the investment dollars in. This was pretty much Univac management thinking of "Let's not take any risks - get promoted by not

making mistakes." We had a stronger response from Burroughs and IBM, than we did from our own company, and also picked up the major minicomputer outfits like Prime in New England and Gould-SEL in Florida.

These two jumped on that bandwagon because they were looking for a way to get out of the HBA game. IPI did have strong technical acceptance, amongst a small group of thinkers, but general opposition in the marketplace, because "we don't want things to change too much". For the publicity and promotion, Univac subsidized every conference I spoke at, and encouraged every magazine article that was written. So I campaigned all over the place, selling the concept of IPI. When ISI came in as an alternative, it was not hit by IPI, but by the decision that people wanted a standard. The drum had beaten for a year-and-a-half that the next standard was IPI and it was coming out of ANSI. That's what hurt ISI - not that it wasn't a sound interface but that it wasn't the Promised Land.

Burniece: It came from CDC.

Allan: It came from the people who brought us SMD but it's parochial. That was the first perspective, because it didn't come into committee until later. I think CDC expected the customers to embrace it, and the fact is by then, customers had shifted their objective, not from their supplier, but from some other group.

Burniece: I can actually corroborate that, because one of the last things I did before I left CDC in '81 was recommend we don't give it to the committee.

Allan: There ya go! Okay. <laughter>

Burniece: Turned out when they finally did it, it was too late.

Allan: I met Peter Wentzel of Siemens at a show in Europe and he was in your camp. He was going to adopt ISI and after my presentation Peter came up to me, when I was in the booth, and just ripped me up and down about how off-base IPI was. Peter was a pretty blunt guy and the only way I could end it was to say, "Peter, I'll be back here next year. Let's just have whoever loses buy the other one dinner." I'd first met Peter when Univac was doing a merger investigation with Siemens a couple of years before. He had the storage group. He was a good head and a strong leader of their storage strategy. The following year it was all over and we became very close friends. Siemens implemented IPI. It was an interesting period.

Burniece: So what was the culmination of the IPI story? Did that turn into a really good standard that actually was implemented by a few people?

Allan: It was a good standard implemented by the companies that were losing influence in the world.

Burniece: Okay.

Allan: CDC Canada implemented it for tape. I don't know that it was implemented on disk.

Burniece: I don't believe it ever was.

Allan: IBM had some disk implementations and Burroughs definitely did. IPI was what you would call a strong base for what was to come later. Shugart was low on the totem pole for products being integrated when they introduced their 14" single platter disk. It was an integration disaster because it was an analog interface in the world of SMD. OEMs wanted the new drive, because it was cheap, but they couldn't integrate it. Around about this time you started seeing controller companies appearing to solve these problems. They were moving from the shadows to become a little more public. It really became obvious when the 8" drive was introduced in the same form factor as the floppy. SA-1000 became the interface to adopt and a lot of controller companies popped up to solve that integration problem.

Burniece: Now explain what the SA-1000 interface was.

Allan: Basically, an analog interface. No more data separation, which was the value of SMD, the real reason it was successful by making life easy for integrators.

Burniece: Yes.

Allan: With SMD you had zeroes and ones. Any engineer can handle zero and one. Now you give engineers a wavy line and an electrical engineer goes, "Whoa," he needs an analog engineer to make sense of it.

Burniece: SMD had the data separator in the drive; it was very simple, right?

Allan: Yes, and along comes this wavy line. Suddenly all of the HBA (Host Bus Adapter) designers in the minicomputer companies go, "Wee-yaaaaah!" That's when you saw Jim Toreson of Microcomputer Systems and others appear to fill those niches. Shugart looked at the situation when they came out with SASI. The thinking was "How do we win?" Priam was actively selling the SMART interface on their hard drives. Micropolis was selling an intelligent interface, and Pertec had one. They were lock-ins, and those companies did not want to take them public. Shugart, was like "Okay, let's see here. We don't have a controller capability, so let's get it built for us. We have no market presence, because these other guys

are selling their intelligent interface against our disk drives before we were. Let's go public and knock-off IPI to get into the standards." Those were two very effective strategies.

Burniece: So was that promoted by the Shugart guys or the committee?

Allan: Shugart came into X3T9.3 and said, "Kill this IPI thing; we've got a much better interface here called SASI." It was a disaster when Hank Meyer pitched this at the meeting. I think it was September in Minneapolis. The hotel was being refurbed and nobody was in a good mood. The idea of a simple 8-bit slow interface replacing a sophisticated16-bit fast interface was a lead balloon. Pshht! Hank came back early the following year with John Lohmeyer in tow, representing NCR. John was the tip of the arrow "You got it wrong last time. And we're here to say there's silicon being developed. We're going to replace IPI."

Gary [Robinson] recognized that another refusal and pushing SASI away was not going to work. The banging on the door would continue. He took me off to lunch and said, "We need to mollify these people. Why don't you authorize two working groups to address SASI in X3T9.3B. I'll go off and set up another committee, so we can move it over there. And by the way, I'm catching a plane in 25 minutes." <laughter>

Hank and I had been going hammer and tongs all morning. I invited Hank and John to dinner at Anthony's on the Wharf and, over a glass of wine at dinner, pitched Gary's idea. Hank knew his target was "Kill IPI!" so he was not inclined but John realized that might work. SASI started in the enemy camp so to speak. It met for two months as X3T9.3B, while Gary was working to organize X3T9.2. Bill Burr from NBS was the chairman.

Burniece: And SASI then became SCSI.

Allan: Yes, the whole of the first day was spent trying to figure out what the name would be. When they walked in at 4:30 to join the X3T9.3 meeting and said, "We're going to call it SCSI," I was stunned! "That spells "Suck-sy!"" <laughter> Daniel Loski was French and he protested, "No, no, no! It means "Sexy!"" <laughter> Gene said, "Nah, nah, it's "Scuzzy." so they had a choice of three pronunciations.

Burniece: Oh, Gene came up with the name for the acronym?

Allan: That's my memory but that's not the way everybody remembers it. A few people reckon I came up with the name of Scuzzy, though I remember Gene doing it. Gene denied it, when I told the story <laughter>, but it was in that 15 minutes of discussion, when the new name was announced, and it stuck.

Burniece: So "Scuzzy," SCSI, was that fundamentally the SASI interface dressed up a little bit or did it really change dramatically over time?

Allan: There were a number of things in SASI that needed attention. It really was a low-level interface. The thinking was right but, for example, take parity. You could have parity but it wasn't stock. The time was spent doing, what I would call, improving its industrial strength. That was the primary effort and our biggest problems were the SASI manufacturers. The controller guys were adamantly opposed to SCSI.

Burniece: Why was that? Put them out of business?

Allan: Yes, they expected customers to start demanding SCSI controllers, instead of what they had. I invited them all to lunch over at the Pruneyard and over lunch negotiated a deal: X3T9.2 would add their specifications as appendixes to the SCSI draft. If the salesmen were hit with, "How about the new SCSI?" they could answer, "we're part of the standard!"

Burniece: The Appendix was the Burger King version of this, then?

Allan: The companies said, "Okay, that'll work. I mean, here's a way out, right? We can support it, quote, "be in it," and keep selling our stuff." X3T9.2 nuked the appendices a few months later, when the companies had SCSI products, but it bridged that gap. We got the full cooperation of all the controller people. Not just Adaptec and Larry Boucher. Larry was the architect of SASI at Shugart.

Burniece: Well, he's the one that actually originally architected the SASI. But did he actually participate on the committee creating SCSI?

Allan: Yes, he did, actively.

Burniece: So he actively did drive it.

Allan: I have to give Larry full credit. X3T9.2 uncovered a bug in the physical interface; the details escape me now, and Larry's sitting there. He's got hardware and he says "Put it in; we'll fix it." His company had a vested interest in the way it was and he wanted to see it correct in the standard. It was the wired OR glitch, if I remember rightly. Larry is sitting there; he's the biggest barrier to getting it fixed. Boom! Give Larry full credit, he wanted SCSI to be the best it could be.

Burniece: Well, he built a company around it, called Adaptec.

Allan: Yes, unfortunately, they went down a rat-hole years later. But you know, they were...

Burniece: Yeah, much later. But for a long time they owned it [the market].

Allan: They were a critical integration entity for years at Adaptec. So it was a good thing.

Burniece: I'm going to do an interview with Larry sometime in the next couple of months.

Allan: That'll be good.

Burniece: He's already agreed to do it.

Allan: Great! .

Burniece: So tell me a little bit more about that whole SCSI thing. How long did that take to come together - finally become a full-born spec and eventually become a real standard, with people starting to use it?

Allan: It took far too long, because there were so many ideas on how to industrialize it, and not always in sync. There was always something else to fix before we could move forward. We got up to like 17a, 17b in revisions before we came to closure, near closure. In the meantime, we had all these incompatible implementations going on outside the committee.

When I'd left Priam to start ENDL, one of the first things I'd done was what I called the ENDL-facto, which took the IPI-3 specification, which was huge, and winnowed it down to what I thought was an essential implementation for controller people. We held an industry conference, that was well attended, and we came up with a finalization.

I was thoroughly abused by clients because they felt I should have done it for SCSI first but the SCSI world wasn't as near [ready] to do that. Anything done from outside would not have worked. You had too many free-swinging controller companies building their own variation, plus you had lots of opinions in the committee. Everybody was involved in this interface but as a private activity we all did the CCS (Common Command Set). It was only a few months after the ENDL-facto. CCS was a start to bring a rigorous discipline for the disk drive.

Burniece: All right, let's look back on a couple of those things, because you used a bunch of terms not everybody's going to be fully familiar with. You started an organization called ENDL.

Allan: Yes.

Burniece: What did that stand for?

Allan: That's actually my first and second name.

Burniece: Is it really?

Allan: Yes. <laughs> It's Gaelic. My first name is Iain, spelled I-A-I-N. The second name is pronounced Dee-Ell, spelled D-A-L-Z-I-E-L. Nobody gets that out of it, right?

Burniece: So that's where you did it? So it's literally your initials?

Allan: It's my first names. It started because my little sister went to school, and was asked the name of her big brother, and her sister. The teacher understood Kay's name, but didn't understand my name, and wrote down, "E.N.D.L. Allan, big brother." <laughter>

Burniece: I'm glad I asked!

Allan: Family joke!

Burniece: <laughs> So this ENDL became a soap box for you. Basically, a newsletter. Right?

Allan: Yes.

Burniece: And you used it for 20-plus years.

Allan: Yes, we published the last issue in December 2012.

Burniece: To basically try to corral all these different forces in the standards worlds to come together on things. So you called it ENDL-facto. So that's taking E-N-D-L-facto, as your—

Allan: Yeah, as a defacto.

Burniece: A defacto version of the IPI-3. So you came up with that as like a straw horse and said, "Here's what it should look like."

Allan: That market was very amenable to it, because companies were thrashing on how to implement a common core. These were large companies. They wouldn't say what they wanted--Burroughs, IBM, etc. They don't hire people in controller companies to do this work for them and they don't bring them in to ask their advice. What they needed was something to disagree on. By giving them that, they all paid to come to the conference and get what they wanted out of it. We held it at the Los Gatos Lodge.

Burniece: All right, so this ENDL-facto thing, became kind of a precursor of the Common Command Set? Did I understand that right? So give me that bridge, how did that happen?

Allan: I was in both committees, and the ENDL-facto got a lot of attention, so people were saying, "Now, what do we need to do here?" It was pretty much a case of seeding. The controller companies realized they needed something similar but none of them could do it alone. It was SMS, or maybe OMTI, which said, "Hey, let's get together at Mountain View and just talk about this thing,". That became CCS and it was a case of everybody knowing what was needed but nobody could kick-start it alone. What it took was me chatting with the controller people that we needed something like this and they did it.

Burniece: Now explain to us a little bit more about what a common command set means. What do you mean by that term?

Allan: What we did was take everything that was in the SCSI standard (the first one still allowed you to have no parity) and come up with the specification for a SCSI disk drive Host Bus Adapter. That was the goal. It was decided that some commands weren't worth doing. Some commands were essential and you had 6-byte and 10-byte versions of commands. We all sat together and winnowed out the stuff we didn't want and added a couple of things we felt were needed that weren't in the standard yet. CCS took on legs of its own. It was CCS that made SCSI an industry standard, because it was amenable to a disk drive. CCS was not written to be an all-purpose interface. It was a disk drive interface that could be implemented by controller people to attach drives. It also opened the door for integration of SCSI into the drive itself.

Burniece: All right, so the Common Command Set became then the central piece of the SCSI standard that needed to be implemented by everybody.

Allan: Right, and that was the basis for SCSI-2.

Burniece: And then there were options you could go beyond that, right?

Allan: In the standard.

Burniece: In the standard. But everybody had to do the Common Command Set. That's why it's called the Common Command Set. And that idea really was precipitated by your ENDL-facto version of IPI.

Allan: I don't think it was solely that. The fact that the ENDL-facto existed had the people who needed something like CCS ask me how they could get it done.

Burniece: That was the secret then to SCSI really becoming a solid standard. It was the Common Command Set.

Allan: I believe it was, because we were kind of floundering around in committee. What we needed was a single club and CCS gave us that club so you could say, "That's it!"

Burniece: So what timeframe was this when this happened?

Allan: Oh, well, I could tell you exactly. I should have brought copies. We did the 100th meeting anniversary edition ENDL Letters for T10 and T11, in which all of this history is there. I probably should have brought them.

Burniece: Why don't you supply a copy of that to the Museum?

Allan: It's on the web; we can get it any time.

Burniece: We can attach it or reference it on the transcript.

Allan: The blow-by-blow is in there of all of this kind of stuff, and the dates.

Burniece: So it's all there. We'll reference it then [2], [3]

Allan: The thing was, CCS took on such a strong identity; people stopped talking about SCSI. Suddenly the tape people wanted their [own] CCS and the optical people wanted their [own] CCS. John Lohmeyer reined all that in and said, "There shalt be no [standalone] CCS. It will be in SCSI-2." The first SCSI-2 ad

hoc to do CCS under a SCSI-2 umbrella was held in Chicago, at the airport across from the Hilton. That was originally a CCS scheduled meeting.

Burniece: You got any idea what timeframe that was?

Allan: I'd have to give it to you—the years kind of merge together then.

Burniece: Late '80s?

Allan: It would have been in the '84...

Burniece: Middle '80s, okay.

Allan: '83/'84 timeframe, because I was at Priam. After CDC bought ISS, the free rein I had at ISS was going to be chopped. That was very clear.

Burniece: So you went to Priam.

Allan: I was there 18 months. I started ENDL in late '84. So we're talking probably '86, the Common Command Set.

Burniece: Okay. Anything else on the SCSI/IPI stories, those two? We've got more standards to cover here.

Allan: Oh, there's a lot of standards and stories in SCSI and IPI, mostly because there was conflict and also there was unanimity. What I felt was my primary role in SCSI was to ensure that SCSI stayed consistent in philosophy with IPI. For example, priority. Do you want one to be priority? Or ten? If the two interfaces are different, the software can't be the same. My focus was not to let them diverge too far, so SCSI could slip right underneath IPI and arrive at...

Burniece: So the terminology and the way people wrote software would be consistent.

Allan: Yes, because all the IPI investment needed to port to SCSI. That's exactly what Burroughs did. They transported all that investment in IPI straight into SCSI. A lot of my early time was primarily in making sure that as we enhance SCSI, it was an IPI-compatible enhancement at the logical level. At

30,000 feet, it was the same principle. Implementation could be different, but if that was the way you thought about it, it was the way you still thought about it.

Burniece: So you were largely successful with that?

Allan: Yes, SCSI went right in neatly under IPI and grew right there.

Burniece: All right, so let's flip on down to the bottom end. Somewhere in that same timeframe, actually a little earlier than that, ST506 came out. And started a new standard for 5.25" disk drives. Eventually that turned into something that became IDE and ATA. Talk about that whole story, how that all happened.

Allan: Okay. Imprimis did the first what I'll call IDE. Compaq—

Burniece: Go back to the ST506 first. That started with Shugart on the 5.25"

Allan: ST506 started with Shugart. When the IBM PC XT came out, it was a floppy-only machine. The first thing that happened was people wanted more storage and you wound up with third parties implementing adapters that would attach disk drives. They were all over the map. IBM had WD develop an adapter for the PC AT and that became the AT interface. The register set and the instructions were a WD designed host adapter for the AT. Integrating disk drives in those early days was a real bear, because none of the drives were compatible. You had connectors in different locations. You even had different connectors. If you bought a drive, and it had to fit a Compaq, it wouldn't work, because theirs was a different variation for their machine only. You had a huge demand out there [for storage expansion].

When the 3.5" came out, it was suddenly possible to put a drive on a card. I did an article on how the 3.5" could transform the add-in marketplace by actually putting an HBA and the drive in a slot. I had no idea Quantum/Plus was doing this as a product. It was an incredibly well kept industry secret. I published the article and within weeks comes this announcement! <laughs> [Quantum's] "Plus" blew the market wide open. It created a real integration market, because now anybody could install a large disk drive into their computer!

Burniece: As long as it was on a card.

Allan: As long as it was on a card. Now Compaq came along. Plus was killing Compaq's market, right? Their add-on business was going nowhere, because nobody was buying extra disk drives from Compaq; they were too expensive. Folks were buying disk drives from the corner store when Compaq came to

Imprimis and said "Look, we'd like to combine the drive and WD adapter as one." Imprimis did a 5.25" and Compaq sold some of those. For some reason or another, Imprimis didn't pursue that further and when Finis came along, he sought funding from Compaq.

Burniece: Now is this Finis, when he started—

Allan: Finis Conner.

Burniece: What became Conner Peripherals, right?

Allan: Conner Peripherals, yes. When Finis was out looking for money, he sold contracts ahead – X number of drives for an investment from Compaq to get the drives they wanted. Although Finis' first drive was a SCSI design, he never built that drive. That was the one he showed when he opened his new building and everybody paraded through but Compaq had him build a 3.5" ATA drive, the fully integrated ATA/IDE drive, and it took off. Finis had the parochial build-and-sell strategy going. The ATA drive for Compaq wasn't the same one he built for Zenith and it wasn't the same as the one he built elsewhere. Finis was selling OEMs special variations to prevent them being plugged. That created a lot of enmity and that's how we wound up doing the ATA standard.

These things all wove together in interesting ways. Jim Rubino from Seagate approached me at an industry conference on SCSI and said, "You know, the biggest market in the world isn't using SCSI and that's the PC. It's because Microsoft doesn't have SCSI drivers. I want a common driver. We've screwed around at Cipher and we think we've got something that's a basis for a way to write software that could be a common driver for all SCSI devices on all PCs. Why don't you do something about it?" So I started the Common Access Method Committee to come up with an architecture for writing SCSI drivers, in general, but specifically for the PC world. That way the software would be common; people could have hardware that did things a little differently, maybe, but it would all work the same.

We'd been doing that for a while and Jim McGrath pops up and says, "I'd like to do something a little different." He hosted a meeting at which he said he wanted to see the capability in SCSI drives visible over the ATA interface. Quantum was using the same software for their SCSI devices and their ATA devices. He wanted a way to get that SCSI function into the ATA connector. DPT put their hand up and said, "We've been doing that for years!" - months, really. They were selling a PC adapter that spoke SCSI on the peripheral side and ATA register set to the CPU. "We've got the answer." Then Jim said something very interesting "By the way, do we know what the ATA interface is? We're building at least five." <laughter>

Burniece: Where was Jim at the time?

Allan: Jim McGrath was at Quantum.

Burniece: At Quantum, okay.

Allan: Gene [Milligan] said, "I'm in!" At the next meeting, which was only a few weeks later, maybe even a couple of weeks later, DPT turns up with their specification and Gene turns up with a sanitized CDC ATA spec. I merged them both and produced the first ATA draft, which was the CDC spec from Gene and the DPT from Tom Treadway. That became the second thread of CAM.

Burniece: And what timeframe was that?

Allan: Oh, '88, maybe?

Burniece: If you can't remember, you and I did do an article on this whole journey and it's on the Computer History Museum website, so we'll reference that here [4]

Allan: Yes, we can pull the dates from it.

Burniece: We'll reference it, because we had all the dates in there to show the evolution. It's a fascinating story. So the bottom line is now you've got a new interface coming that's semi-intelligent. Not quite [as intelligent as] SCSI and IPI, but ATA became the dominant [HDD interface] in the world.

Allan: That's what happened after we regularized ATA. Now, guess who didn't like it? Conner Peripherals. <laughter> We had a terrible time finding out how to make Master/Slave work with all the variations Conner had out there. Anyway, let's just say that being the first editor of ATA was challenging. Now the other wrinkle that pops in here …

Burniece: So you're saying the irony of this is, even though it got started and was partially implemented by Conner, one of the first guys out there with it, they didn't like it?

Allan: No! For the same reasons as Priam and Micropolis didn't like SASI, right? <laughs> It killed their homegrown control. We thrashed around and basically [did] competitive analysis of the Conner drives to figure out how to get Master and Slave to work successfully on any platform. I don't know how many times I rewrote that section.

WD came in with a pitchfork. They wanted to eliminate anything to do with the enhanced ATA, which was DPT's ability to run SCSI commands through that connector. WD led this campaign to kill EATA and they won, as it became a separate specification of its own. Then it just died and disappeared. It didn't go anywhere; it was not part of the finished ATA specification that we gave to X3T9.2.

Burniece: Which part went away?

Allan: EATA. The enhanced part which was the ability to run SCSI commands for—

Burniece: Right, Jim McGrath's original idea. So that originally does—

Allan: McGrath was Quantum. He wanted an ATA spec he could build one drive for. It was Tom Treadway from DPT who wanted the Enhanced. Jim only wanted the ATA defined and the ability to have the SCSI commands executed by ATA for special things he wanted to do. Jim had two requests and Tom Treadway had the answer to Jim's second request, which nobody else wanted. Or should I say, WD didn't want Quantum to do what Jim wanted to do and set out to kill Tom Treadway's spec. WD managed to get EATA out by saying, "Look, if this thing gets in there, we're all going to have to build this stuff. We don't want to add this controller cost factor."

Now, I don't know if there's any link here, but when the CD-ROM interface called ATAPI was done, guess who was the leader? WD. Guess what they used? The SCSI control block structure. ATAPI threw out the SCSI control structure what they didn't want to do for a single drive. DPT was too early. It was Compaq who told WD "You and the CD-ROM guys get together and give us a CD attached to the second ATA connector." WD did it with a subset of what DPT had done years before - the cycle came full circle.

Burniece: So through this whole evolution, we ended up with a half-a-dozen different efforts going on in various sectors, and we ended up with two basic fundamental disk drive interfaces.

Allan: Yes.

Burniece: ATA, which now, of course, has a serial version [SATA].

Allan: And now we're up to Version 8.

Burniece: And SCSI, which also has serial version, SAS. But we got another one in here, too. We've got Fibre Channel. Tell me about the Fibre Channel story.

Allan: Fibre Channel goes back to when I was at ISS. There were stories coming to me that IBM was implementing a fiber interface based on a Cambridge ring. [IBM] was going to interconnect their controllers to multiple processors, so they could share data with more processors than they could [with Block Multiplexer Channels].

Burniece: IBM was going to compete with the VAXcluster.

Allan: It could have been.

Burniece: Yeah, because that's exactly what [the VAXcluster] was. It was a ring and it shared data [between multiple VAXs] but it was [serial] copper.

Allan: This was coming out of IBM San Jose and it was my job to know, so off I went to Roseville and pitched this idea, told them it was coming. The response was "Ah, don't be silly."

Burniece: So what timeframe?

Allan: This was before I went to Priam, so it would have been late '70s, early '80s. We were concerned, because we didn't know anything about fiber and we didn't know anything about serial. ISS is building controllers for all these old channels and we've got our own [IPI] that we're pushing which would be old-fashioned if IBM was going to do this fiber thing.

Roseville chewed on it and they actually came up with the concept of FDDI. FDDI started as a competitive response to a perceived IBM channel and once again Univac decided, "Oh, gosh! We don't want to commit to this! We don't know for sure IBM's doing that! We can buy it OEM afterwards." So FDDI goes into committee at X3T9.5. Now this is 100-megabit per second interface, and you've got Ethernet running at around six. There'd been a deal cut between X3 ANSI, or should I say X3 and IEEE, that anything above 50 [megabits per second] belonged to X3. We were doing the CI ([VAXcluster] Computer Interface) in X3T9.5 as well because DEC needed it standardized for Federal sales.

Burniece: That was 70 megabits a second.

Allan: FDDI is moving along quite well and DEC looks at it and says, "We don't have to do a network at the IEEE. We can go over there and make FDDI a network!" DEC dived into X3T9.5, took control, and created the network that is FDDI. DEC wanted something faster than IEEE was allowed to do and FDDI went from being capable of being a storage interface to a local area network interface. I stayed active in

X3T9.5 and even worked on the isochronous version (FDDI-2) but it was more of an interest than a commitment.

HIPPI (High Performance Parallel Interface) came along at 100 megabytes per second. It waltzed in from Los Alamos Labs as a very wide parallel copper interface to support graphics. "It was for a super computer and nothing like IPI was suitable, even though IPI could get there in the next generation. That's for ordinary people; super computer people need their own interfaces." HIPPI became an X3T9.3 project.

I re-invented FDDI. I sat down and put together what a fiber interface should be to pull together all of these interfaces, SCSI, IPI, FDDI, HIPPI and possibly an Ethernet backbone. It would be interface agnostic but we couldn't start a project like that, because our X3 charter didn't include it. Fibre Channel started off as an ad hoc committee to do the next generation IPI physical interface and once that was given committee approval we started the discussions going. At that first meeting we went through the list of things that I had for discussion, and as we ticked them off we generated a Statement of Work for Fibre Channel.

Burniece: So that's how Fibre Channel started. It was kind of a consolidation of all the things that were happening with FDDI and HIPPI and even to some extent SCSI and IPI.

Allan: If we did a network, we'd be screwed by the IEEE and the OSI model. To do a network, you had to follow the OSI model. The first line on that list of things we wanted to do said "closed system."

Burniece: You wanted a channel.

Allan: The OSI model need not apply to a "closed," system. Del Hansen of from Hewlett-Packard was a real thorn; he kept repeating, "This is a network!" "No, it's not." "It is a network!" "No, it's not." <laughs> Del was outvoted n:1 all the time. That was the way we started, but we didn't know how to do it. I put a one-year moratorium on proposals. During that year, we contacted every company, which had implemented a serial channel replacement. HP had one. ICL had Macrolan. There were several agency interfaces, which had been done for the government. Rockwell came in and told us about theirs. Not only did everyone tell us how they implemented, but even better, they told us what they'd change if they had a chance to do it again. We were given the insight from what they'd done and what they wished they wouldn't have done, and what they would do if they had another chance.

We asked the optical companies to tell us about parts, and what was very clear is that none of these products had a market. They all wanted to sell to SONET and had a technology, which came down to "We can build this, but there is no use for it." Each company was very parochial in its perspective on how their technology should be used.

We divided the project into multiple tiers - the physical interface and the next level up was the encoding scheme. IBM wound up taking the lead on that. Next we moved up to the command structure. We wound up with me leading the logical interface and Schelto [van Doorn] leading the physical interface. We were struggling hard for a physical technology and one day a guy turns up from Rochester who says he is buying CD lasers cheap as dirt and has them running in the lab for transmission testing. He had been pitching this idea to the optical standards organizations for two or three years and they were ignoring him.

Burniece: I assume this is IBM Rochester.

Allan: IBM Rochester. The cheap laser works and is running with high MTBFs and, coincidentally, IBM Rochester is hosting our next meeting. We go; we visit; we walk through the lab and he shows us his heat chambers and the testing. A tremendous amount of work on a laser that costs virtually nothing. Suddenly, we've got a vehicle to meet the speeds we needed -- 100 megabytes a second to absorb HIPPI, plus the overhead, so it was net carrying capacity. Our spec was 1.0625 Gbaud, not 1 Gbaud but 1.0625. That was figured on a spreadsheet that if you took this size header and that size trailer, and all the gaps you need to carry 100 megabytes a la HIPPI. We were talking bytes at a time when everybody else was thinking bits.

Burniece: Right.

Allan: Suddenly we've got a number, and that number coalesced the entire optical industry around a product goal. Fibre Channel wound up changing the structure of the optical industry by giving it a target. That flowed on from the one guy who walked in and we used CD-based lasers for a few years.

Burniece: Have you recalled his name yet?

Allan: Ron Soderstrom.

Burniece: Okay. So what timeframe we talking about when this is going on?

Allan: We're talking late '80s, early '90s.

Burniece: Yeah.

Allan: The first Fibre Channel meeting was June '88.

Burniece: Okay, that would make sense. So when did you have a spec?

Allan: We were floundering through the early 90's.

Burniece: When did you get to a first level spec?

Allan: Oh, it took two years easy, because <laughs> we'd be up there on the chalkboard and there's this big cloud in the middle. We didn't know what the cloud was but we knew it had to be switched. That was one of the bad things. We knew what we wanted these parts on the outside to do, because they were familiar, and the problem to solve was in the middle. We just drew a cloud and plugged into it. We hacked our way through that cloud and that is what eventually turned into a fabric.

Burniece: Yeah, but the original was arbitrated loop, right?

Allan: No, that came later.

Burniece: Oh, it did? Oh, I didn't know that. So you started off for a switched fabric.

Allan: We started out for a switched fabric. I explained why in all the pitches that I made. "When you pick up the phone at home and call your mother-in-law, she's in Florida, and you're switched through somewhere in Missouri. That switch goes down and you continue talking, but you're now running through a switch in Chicago. You don't know it happened, because there was only a little bit of lag time when that switch went down in Missouri and was rerouted through a different place".

We wanted to have a disk drive hooked up from anywhere and have it say, "Yoohoo! I'm here!" and back would come "Oh, hi! I could use you." The rationale was based upon the home telephone line. First of all, it was very simple to understand and secondly, it didn't need station management. You don't do any station management to log in when you use the home phone. We wanted to plug into the wall, say, "This is my number. And I'm here." We needed the kind of people who understood the telephone-switching network to explain to us what was possible and they came from a company called Canstar. Kumar was in on Fibre Channel from day one, I'd met him on FDDI and—

Burniece: Talking about Kumar Malavalli?

Allan: I'm talking Kumar Malavalli. So he came in—

Burniece: And you just did an [oral history] interview with him.

Allan: Yes I did.

Burniece: Last year [5].

Allan: Kumar had been in the FDDI committee and I'd been up to visit Canstar. They had some interesting technology at the University of Toronto but it did not have wheels. There was no traction anywhere. As we moved along, Kumar introduced his offsider, Bent Stoevhase. Bent had a lot of time in the telephone world and it was Bent and Kumar who brought in a proposal for what turned out to be the fabric. Bent was a deep thinker, so when he brought something in it wasn't just an idea' it was a full-fledged concept that we all triggered on. We went from having a cloud to a fabric with defined functionality. We didn't want to do anything like Ethernet ARP (Address Resolution Protocol) because that would be so disruptive, so we created the Name Server. That's your equivalent to the White and Yellow Pages. Those early concepts were based upon the simplest way things work that we're comfortable with. Next came, how do we implement that in our world?

Burniece: Fascinating. So you used a telephone witching model basically—fundamental.

Allan: Yes, a pretty simple model. Of course implementing it was very different and we couldn't get a cheap switch. All the switch guys kept promising us that we'd get a switch on a chip but they never came through. And it was very clear that if we wanted to hook up disk drives it had to be done differently. You knew Wayne Sanderson [CDC], I assume.

Burniece: Yes.

Allan: Wayne was one of the key guys and Horst Truestedt from IBM Rochester was another. The two of them took it upon themselves to get a cheap disk attach; they broached all kinds of ideas and we went through many proposals. Eventually that's what became AL (Arbitrated Loop), as the disk attachment.

Burniece: So the arbitrated loop came out after you'd conceived the concept of a switch fabric but it actually preceded a real working switch then.

Allan: Yes.

Burniece: Because that didn't come out until the late '90s.

Allan: That's right. Back in that early timeframe there was a company in England building what was called the transputer.

Burniece: Yeah I remember that.

Allan: It was an incredibly simple switch; it had three ports and as you approached the input you decided, "Where do I want to go next?" And for some reason you chose one output port or the other. It was source routing but it could be structured to do destination routing. It went nowhere but that's where I had hoped we would get our switch technology. I thought it was going to be dead easy to get a switch and it didn't come easy at all. That's why we had to do the loop. To sell drives meant that the loop was the first implementation of Fibre Channel.

Burniece: So when did that come to the point where [AL] was now a real complete standard and—

Allan: I think people would tell you it still isn't complete; it's still growing and expanding.

Burniece: Still is? Okay.

Allan: It went through phases. And we went through phases that were very interesting. Paul Rupert from Livermore Labs had studied the eye and he knew that 100 megabytes per second was the fastest speed that the human eye could absorb. The Ancor Class 1 switch was a direct connection - very channel like in that you'd hook up, blast data out to the terminal, and there'd be recurring traffic on that path. This was the very simple computer interface model of a channel put into a fiber network. I send it; you acknowledge you got it. Very, very tight, backwards and forwards. That was the first thrust.

Class 1 merged into the idea that "Well gee we'd like to be able to do the network as well". So we [Fibre Channel] added a Class 3, which is unacknowledged service. And in the middle we said: "Well we'd like to have a real connection but we don't want it to be there all the time. So let's packet switch Class 2, which is connectionless driven, but you acknowledge that you got it back to me, so you and I both know I sent it; you got it." Classes 1, 2 and 3: Class 1 was the initial focus; Class 2 was the next focus; and Class 3 was not really intended for anything but local area networking.

But then along comes the people at Hewlett Packard who developed the Tachyon chip - a networking chip, which implemented Class 3 only. Seagate's supposed to attach to a Tachyon-based system and the Seagate drive only had one 2K buffer, so it couldn't send and receive at the same time. Somewhere in the thinking it was decided to drop the Acknowledgment and go connectionless Class 3, instead of Class 2. It's one of the things that I truly regret about the system because, in fact, the disk drive didn't have to give the ACK. The system was totally responsible for error recovery. When the drive gave data

to the CPU, it was not responsible to recover if the CPU didn't get it. It was the desire to be fully symmetrical that took everything into the Class 3 vein and all of the effort in Class 1 and Class 2 just went away. Class 3 became a hit because [HP] became the primary integrator instead of Ancor. Instead of committing to get something done. Ancor failed by being selected to supply IBM Austin and when you are selected by a major company to be their source, nothing else happens.

Ancor didn't get the chips done they should've done in the timeframe they needed to. They didn't do this; they didn't do that. They went away from supplying Paul Rupert and Ancor lost, because the integration vehicle became [HP's] Tachyon. Emulex designed their own chip as well and the focus then became: "By golly, we've got to make Class 3 work for disk drives." Of course that was another generation, so [Fibre Channel] was implemented three times. We implemented a Class 1; implemented a Class 2 on paper; and [then] had to implement a Class 3, because that's where the hardware took us.

Burniece: So when did that finally come together?

Allan: Late '90s is when it started to really tick. I didn't feel comfortable until '98 that this was really going to work, because in the meantime we had the competition from IBM when—remember when IBM was selling off their disk division - it became ADSTAR?

Burniece: Yeah right.

Allan: That was Ed Zschau. Then Jim Vanderslice came in to run it. He was a printer guy. And the message was: "How do we turn this place around?" The troops were saying: "Adopt industry standards." Whether he transformed that pitch or whether his middle management transformed the pitch, it became "make IBM interfaces industry standards." That became SSA and SSA was pushed on the committee.

Burniece: Now tell us a little bit about SSA. Were you involved in that also? Or was that something you're just watching on sight?

Allan: Oh that was a pimple - turned into a boil.

Burniece: Yeah and never really did turn out to be a success, although it was a great interface.

Allan: It was a great interface for inside a cabinet where you controlled the hardware.

Burniece: Right.

Allan: It was a structured interface. If something went wrong with the cabling you were dead - it stopped working. It was very unforgiving in its configuration but inside the disk cabinet it was dynamite. You could do spatial reuse, so it had a lot of real advantages. It also gained support in the industry. For example, Conner committed to build SSA and so did Micropolis, under the thought that they'd get IBM contracts; that was the thing, to be a second source. As Seagate was number one in Fibre Channel, they needed to not be the same.

The real reason SSA failed had nothing to do with the standards interface and nothing to do with the industry. IBM shot itself in the foot. Oh it's great to get an industry interface but what do you do, when you get an order? If you're building all the drives you can for inside, and you get an order from outside, you have to fulfill it or you're not an OEM supplier. When IBM finally got orders, nobody got drives.

Burniece: Right. Well they had a problem.

Allan: The only company that got SSA drives, was Siemens. Nobody else could get drives. IBM was a real pain; they caused a lot of grief and confusion in the marketplace. The ultimate death was they weren't an OEM supplier. If you ordered stuff from HP, they would screw the internal divisions to supply the contract. At IBM, they screwed the [OEM].

Burniece: Yeah. Well CDC [favored the OEM, as well]. Let's switch over to the small form factor thing. That's another interesting story. Talk about that one.

Allan: Okay. Well once again this was a private group that got together …

Burniece: And what timeframe are we talking about now?

Allan: Oh the first two-and-a-half inch drive. When are we talking? Probably—

Burniece: Early '90s.

Allan: Yeah late '80s, early '90s. Sun Microsystems was going to build a laptop - basically a laptop server. They had two two-and-a-half-inch drives offered to them but they were completely different. Sun couldn't put them in the same box because they were so much apart – they couldn't package the laptop correctly. Sun held a meeting in Colorado and invited a bunch of people to it - a few select drive manufacturers and connector manufacturers, because it was a connector problem inside the case. Sun said: "How are we going to get something we can use

Everybody who hadn't been invited was boiling. They didn't know what happened They didn't know why the get together but they did know Sun Microsystems had held a private meeting on disk drives. Antitrust threats were running around. The answer was a meeting that Seagate hosted in town here [San Jose]. The word went out: "We're going to get together in an open meeting, that Seagate's going to pay for, and we're going to talk about two-and-a-half inch drives." We couldn't start the meeting on time. There were so many people at the Le Baron Hotel it was unbelievable. The hotel had to get us a huge room. We'd started off with maybe 20 and there were well over 50 people, if not 50 companies represented. The upshot of that was things were in turmoil and I walked out of that room and started the SFF Committee; "Okay, let's just get this squared away and we'll do it together as an industry." And that was SFF.

Burniece: And that still exists, right?

Allan: Yes, still ticking along. We started as the Small Form Factor Committee but we didn't last long with that name because, even while we were doing the two-and-a-half inch disk form factor, we were working on connectors. The SCA (the Single Connector Attachment), for disk arrays to plug drives directly in, came out of SFF. We picked up ATAPI because, when it became clear that there was a private CD group going on, that became another furor. The private group had to choose between going to a standards committee or doing ATAPI at SFF. The vote was a 50/50 split to go to ANSI and it was a 99 percent. I think there was only one vote against, or maybe even one abstainer, to go to SFF. We dropped the full name and it just became plain SFF. The acronym was known so well by then we couldn't change the name of the committee to be more encompassing; we simply dropped the meaning of the initials.

Burniece: So what's the total scope of the SFF committee?

Allan: Anything, which does not fit well in a storage committee and is a problem.

Burniece: So it's kind of a—

Allan: If it doesn't fit a standards committee we'll do it for you. When ATA was in X3T9.2, the engineers had great difficulty getting projects started, because it was a SCSI committee. The ATA guys would bring proposals to SFF and we'd start on them; but one of our caveats was that if a standards committee wanted it, we'd immediately hand the project over. Jim McGrath was one who did it several times. He'd come in and make a proposal at SFF, we'd ballot it to say: "What should happen with this?" Jim would stand up in the next X3T9.2 meeting and say: "We're balloting at SFF to do this function unless it's taken by X3T9.2; your choice. If you don't take it, we'll do it over there." Much of the time it was taken up by X3T9.2, but things like SMART started in SFF, finished in SFF, and went to X3T13 ATA, where they've been enhanced and evolved since then. The ATA 80-conductor cable to run 100 megabytes a second was SFF.

Burniece: So you have three standards focus committees with the Fibre Channel group, the IDE group and the SCSI group. Then you have SFF, which does everything else.

Allan: That's pretty much it.

Burniece: Give it to the other guys if they want it; but in the meantime you're going to make it happen.

Allan: Yes.

Burniece: So you're doing everything from connectors to SMART—all that. Are you part of the official committees?

Allan: No. We are not democratic at all. SFF is a totally non-democratic committee.

Burniece: So you're not associated with any of the standards bodies?

Allan: No, the rules were set up by the members. We have a rule that's 'majority of one', one valid technical negative will prevent publication.

Burniece: One?

Allan: One.

Burniece: So it is a veto. One guy could veto.

Allan: Any valid technical negative will veto the publication of the specification.

Burniece: And how well did that work?

Allan: Incredibly well. If you're in a standards committee, all you need is the votes to get through.

Burniece: So it's majority rules?

Allan: Majority—very, very democratic. So you can stack the deck: "Hey, we got a vote coming up and you're selling me cables; I'd kind of like you to help me out here." There are political influences in a democracy.

Burniece: There's a lot of lobbying going on.

Allan: There are none in SFF. If you don't like it, you say so; you win.

Burniece: Wow.

Allan: And the only way we can override your veto is to vote that we do not accept your technical negative as being acceptable. You can't get an engineer to tell another engineer he's right but he's not acceptable. So the issue gets fixed. We have this one goal, to be accurate, and SFF runs like an engineering project.

Burniece: And you're still active in this?

Allan: Yes.

Burniece: You're still leading it?

Allan: Yes.

Burniece: How many people are participating in this?

Allan: We've got about 50 members and they're worldwide. SFF has been the only worldwide standards activity in this arena, because we don't require attendance. You can be sitting in Taiwan, China, Japan, and join the SFF committee.

Burniece: Oh, so you're not required to be there?

Allan: No and you can vote on all these projects. You can send in proposals but you never have to turn up. It's a huge inhibiting factor to come here, speak in a second language that you're unfamiliar with, and express your thoughts in a coherent manner.

Burniece: And have everybody else beat you up.

Allan: Yes, people stop listening to you. They don't understand what you mean; you're stumbling. It just doesn't work. But this way they can sit down in front of their machine, figure it all out and write. Many people read and write English much better than they speak it and objectively take their time to propose an idea. Before we had the Internet, the whole Far East ran on fax machines and we had a fax server for the SFF specs. There would be hours-long connections while people downloaded dozens of specifications.

Burniece: So what are some of the real gems that have come out of SFF? What are some things that turned into real complete standards that the world has accepted?

Allan: When you pick up a disk drive today, no matter who makes it, they all fit the same slot - the same screws, the same position.

Burniece: So the form factors and all that stuff, it was SFF that did that?

Allan: Yes, when we started the committee it was 2.5 inch that was the thrust. Dave White from Compaq became the editor and he got fed up with the disk guys always having a reason not to finish this spec. At the meeting where he showed his final draft, he said: "And by the way I'm really not interested in how well this ballot turns out because this spec will be appended to the next RFP that Compaq puts out to buy disk drives. You don't have to build to it."

Burniece: ... you just can't sell anything to me...

Allan: Yep, this is what we're buying to. And bingo, guess what?

Burniece: Wow.

Allan: Yeah.

Burniece: When was that?

Allan: Oh '91/'92 timeframe.

Burniece: I was going to say—I didn't realize SFF went back that far.

Allan: Yes.

Burniece: The bottom line is that standardization of form factors has been around for 20 years.

Allan: Yes.

Burniece: This was not a new idea.

Allan: No.

Burniece: Everybody's had the same screw holes, the same side rails, the same exact shape, and the connector in the same spot.

Allan: Nobody wanted it at the time.

Burniece: SFF drove that?

Allan: Yes. Dave had had opposition right up until he said, "That's it." What amazed him was three months later when Tom Hanan from Western Digital stands up in a meeting and says: "We'd like to do the 3.5 inch form factor; could you do that for us too Dave?" I mean, 'learning fast', in 90 days.

Burniece: Wow.

Allan: People had known about this specification for months but never committed to it. As soon as that size was committed to, the 2.5, the drive guys realized: "Hey, I can second source everybody; or I can be primary source and they can second source me." They just jumped on board and immediately the 3.5 came next, then five and a quarter, and then the CD. We just nailed it.

Burniece: To me it's one of the more amazing parts about the disk drive industry is how it is so standardized on all those things that really count at the end of the day to integration.

Allan: Yes. All of this is integration oriented.

Burniece: Inside they can have all kinds of invention; but everything that has anything to do with integrating to the world is fundamentally standard.

Allan: That's right.

Burniece: And I don't think any other industry has done that - certainly not to the level disk drives are at.

Allan: I don't think so. This is one of those things; the first standard anything was the IBM Selector channel.

Burniece: Right.

Allan: The commands were the same for all of the CPUs and that was a common thing —

Burniece: That was a standard within IBM itself, for its own products, and the world if they wanted to plug them had to adapt.

Allan: That's correct. And so—

Burniece: They didn't do it for the world; they did it for IBM.

Allan: They did it for themselves, because IBM had all these different empires, which had their own way of doing things. IBM killed that. Years later at Univac we were trying to build products for these different empires and we were having to invent ways to try and get them away from it. We were building SMD drives for Univac, only to stop them building something unique. ISS never sold SMD drives but we were building them. That was the way we could encourage, to use the right term, these organizations to accept it. Don Neddenriep had been instrumental in buying ISS; he was the VP of Development.

Burniece: He was the Univac guy?

Allan: He was the Univac guy and he had a real interest in what was happening out West. There were standing orders that whenever I visited Blue Bell [headquarters in PA] I should stop by Walnut Row and let the secretary know I was in town. At any time in that week I could be pulled out of a meeting to go visit him. This had a very unsettling effect on a structure where everybody knew their place by their title and where they were. I'd be called out of a meeting and have to go up to spend a few minutes with Don, bring him up to date on ISS, and then I'd come back to the meeting. The Univac guys didn't know where ISS stood in the structure or where I stood in the structure. So when we were trying to...

Burniece: They just suspected you had some special privileges.

Allan: Yes. When ISS was encouraging them to use SMD internally instead of their own interface, at the back of their mind was: "Is this Don Neddenriep talking?" That was a nice niggle to get them moving off their implacable determination to do their own thing. You had to use every wedge you had to win and that was a very good wedge as it turned out.

Burniece: So, of all these things you've done in the standard worlds over the last 30-plus years, is SFF the one that you are the most proud of, because of the way it's become a process that has transcended all…

Allan: Yes.

Burniece: … of all the issues with standards committees?

Allan: SFF is a niche player. We're there to help; we're lubrication for the system. Standards committees have difficulty with connectors and they all have stories of connector wars. We do all the connectors now and when we have it figured out the standards adopt them. SFF is an expedient place to have these things done because it helps.

Burniece: So you look at is as lubrication in the standards world.

Allan: Yes.

Burniece: Rather than necessarily the solution. It's the lubrication that provides solutions that they can adopt.

Allan: We're trying to help everything move a little more smoothly. When we figure out a transceiver or a connector or whatever for a purpose, the spec's there. The [standards committees] don't have to argue about it; they just have to decide whether they want to use this one or that one Every electrical engineer knows what a connector should do and they will tell the mechanical engineers how to do their job. That's the psyche, the electrical engineers are in the standards committee and you've got the connector guys who are mechanical engineering, trying to sell these electrical engineers. The electrical engineers say: "Well why don't you do it that way, why don't do it that way?" "I'd rather have it done this way." You've got this kibitzing being done all the time. Everybody misses the days when we used to have these great dinners served by the connector companies, jumbo shrimp to get your vote. That's what it was like, very political. It just hasn't happened that way for years now, because we're the lubrication at SFF.

I think also - taking Fibre Channel from that piece of paper, I steered it through the technical stuff with a terrific team of people. I've been involved but the work's been done by other people.

Burniece: Who were some of the key guys in that?

Allan: Take IBM Austin, Texas; they really decided to get into Fibre Channel. Joe Mathis, Jerry Rouse became editors. Carl Zeitler was the leader. They threw their energies into it. Roger Cummings was an early disciple and the two of us used to go out to all these events to speak. We'd share the rotten tomatoes that were tossed at us, and so on. Kumar, of course. Just about everybody - as you'd go through. We've always had one or two that becomes five that becomes 10 that becomes 15 - who were infected by the people who were working on Fibre Channel.

Burniece: Right.

Allan: And they pooled together their united capability. Even better was that when we did the FCLC, to sell Fibre Channel, I became the guy running the shows. It went from the keyboard to paper to implementation to selling. I wasn't part of the implementation but I got to do the selling side and I ran all the shows for FCLC. We were a huge, huge function; we had million-dollar budget events.

Burniece: Really?

Allan: We'd have 50, 60 Fibre Channel members at all of the industry events - CeBIT, Comdex and such. The events would come courting us because we had so much money to spend in FCLC to promote. That was kick-started by Seagate; it was forced because SSA did a great event at Dataquest. The old Fibre Channel Association had nothing except paper to compare with demonstrations and Mike Fitzpatrick from Seagate stood up at the next standards meeting and said: "We can't sell with this environment" and led the formation of an association [FCLC] to just sell. This was in September and Comdex is November. We had to be on the floor at Comdex and Seagate had the muscle to get us a room. Seagate staff helped us get the packaging together and that room became our first event. It was packed every day.

Burniece: What year would that be?

Allan: Mid '90s - '96, '97, something like that. Once again we can look it up. FCLC took over the shows from Seagate. We had our own contractor, Exhibitree, down in Orange County. I designed the events, the things we used, organized—everything. That was really into the promotion side of it and selling it was probably the role with the most influence.

Burniece: How much influence and even involvement did the people that you worked with in those committees have with iSCSI and Fibre Channel over Ethernet - at all? Or has that all been driven by the…

Allan: We always had crossover. Yeah there was—

Burniece: …IEEE guys and the ITF.

Allan: For example, iSCSI. In the beginning I couldn't sell Cisco on the concept of Fibre Channel at all; they just politely talked to me but they would never make it.

Burniece: It's got to be on IP.

Allan: It's got to be LAN—right? There was no support out there, so we decided to put LANs on Fibre Channel; we've got a standard for that. During this whole cycle, the market size began to boom; it's an $80-billion business today. Kumar says, when it started to get big, Cisco and the others started to say: "Wait a minute here" and they wanted alternatives. That's when the Internet crowd came in to work with us on Fibre Channel. Everybody who's come in has always been happy to work with what's in X3T11 and work with the people. There's never been an antagonistic relationship at the working level. There's competitive spirit at the company level but we've always had what I would call a degree of working together. It's the psyche and the culture that your job—especially a disk drive guy—there's five guys like you in the world who know what you're doing. Those five guys all work for competitors.

Burniece: Right.

Allan: You can't meet with them and talk about what you're doing and what you thought of last week that was brilliant and nobody else had done yet. He who understands is a competitor. You can come to a standards meeting, in neutral territory; you can each contribute to a common goal that you both can live with. You may come in with opposing ideas but the process is to have them vent, so they can get their idea out, because until they get it out they can't listen. You get their ideas out, then you propose some alternative that both of them disagree with. Now they have to unite to prevent that dumb idea getting into the system. Right?

Burniece: The engineers. Right.

Allan: They just became allies. Then they socialize. They're having dinner and so on. The next day or the next day after that or maybe even a month after that the ideas the other guy had trickle into your head and vice-versa. And suddenly you're working together.

Burniece: I think we got back to this first life lessons you got.

Allan: Probably. It's a nurturing environment where you encourage people to express themselves, without stepping outside their knowledge or the company knowledge that shouldn't be provided. To get their technical skill, we're directed in the same way. And everybody buys in. That's not an independent thing; that's a culture of the environment that we've got: T10, T11, SFF. It's all there.

Burniece: For the most part have each of these committees had kind of a consistent set of people that stay with it for all the years or has there been a lot of turnover?

Allan: There's a small nucleus. For example, on T11, Roger Cummings was there in the early days; he's still around. Horst Truestedt is still around. For the most part everybody's dissipated. On T10, Bill Burr was the first chairman; John Lohmeyer was the vice-chairman. When Bill left, John went up to chair; I became vice-chair. John's still running T10. So you've got continuity of a few core individuals; and then you've got people who had been with them for years that take over and then others come and go. When HIPPI was the big thing, Don Tolmie was the chair. Now you've got Steve Wilson from Brocade who's in the T11 chair. But his vice-chair is Claudio DeSanti from Cisco.

Burniece: Right, I know Claudio.

Allan: You've got all of these competitors working together and, what you might call the ugly side of competing, is not in the room. That's neutered; that's out of it and it's a pretty invigorating atmosphere. It's hard for people to believe that, but you socialize, you work together. You've got a tough problem to solve; there's three solutions - which one's the right one? You work it through.

Burniece: What about the SFF group? Has that been a pretty consistent group of people over time? You've got a whole potpourri of things to work on there.

Allan: The core is the connector companies. They want a place where they can develop things for other companies, their OEMs, and they want to have support. They do their individual fighting in InfiniBand and at T10 and at T11; that's where they do their individual competition. In SFF they want their spec finished under industry scrutiny so that it's technically accurate. The fighting is done outside the environment that we're in; so yes, it's pretty cooperative.

Burniece: It sounds like a great journey. We're getting close to four hours on this and I want to wrap it pretty quick. So just give me a couple of closing thoughts here. Are there any other favorite stories that come to your mind that you just haven't had a chance to talk about yet [but] you'd like to?

Allan: I should've—

Burniece: Get it on record.

Allan: Okay. When we took the break, I should've brought it back to Ken Hudson. He hadn't been fired. Ken had been told to disappear for six months. He then came back in a corporate position at headquarters and I wound up working for him. Art's action to inspire his troops wasn't a death.

Burniece: So he didn't kill him; he just put him over in the corner.

Allan: He just made him invisible.

Burniece: And scared everybody else, into thinking he'd been shot.

Allan: Yes.

Burniece: Wow. Any other favorite people that you worked with that you really would like to talk about a little bit?

Allan: I'd have a tough time, because I've been around the periphery of so many great people in all of these things. I've never worked arm in arm on a development program in a company since back in the days of doing Linkage Editor. At ISS, the engineers were the ones implementing. I've never been able to implement the thoughts that I've had; they've always been done by other people. At ISS, Ron Roberts and I would be screaming at each other in his office about stupid ideas and retarded implementers and then we'd go out and have lunch, come back and start doing it again. Nobody in his group understood that he was the reality check on what could be implemented and how it could be implemented, which was often different to where I wanted to get. I knew what would be a good idea, what would sell in the marketplace, and what would be useful. And he's over there knowing what he could do with his resources. We'd figure it out. It's always been that throughout, that there's always been people that I've been able to bounce ideas off of; they've taken the ridiculous and turned it into the sublime.

Burniece: How about some of the key people that you admire in the industry, just from a leadership standpoint? Anybody that stands out there or a couple of people that really stand out?

Allan: I think anybody who's taken a company through the process of going through the vulture capitalist cycle. I went through that once and I don't want to repeat that exercise. Getting that company to the point where it could stand on its own two feet, that's worthy of admiration. Running a business is not something you learn in engineering schools.

Burniece: It's true.

Allan: And so I look at the guys who came out and started their companies and did it. Take Al.

Burniece: Al Shugart.

Allan: When Al was let's say 'encouraged to depart' [from Shugart Associates] he went over the hill, bought a restaurant or a bar. I'm not quite sure which was the most important but very little of the profit hit the bank account. Al was just very morose and depressed and along comes Finis, who gives him another thing to do and suddenly he's changed completely. That's a transformation. So many other people would've taken that first hit and kept going down that slide. It didn't happen; Al came back. He even ran his dog for Congress.

Burniece: Right; yeah he did.

Allan: So just about everybody I've met—this has been a cowboy industry, honestly.

Burniece: Absolutely; I've called it that from day one.

Allan: The problem is that now it isn't. It's stopped being the cowboy industry. The installed base is so large. You couldn't start a Fibre Channel today because there's no hole to fill. There's nowhere to walk in and say, "There's a gap, we need something; let's start it." Those don't exist anymore. We've gotten so much technology. I mean, how old is your youngest CPU, at home on your desk? Is it a brand new one or is it three, four years old?

Burniece: I got one that's five years old.

Allan: Yes. We have stopped having to get the next generation because this old one does so much, and now we've got a surplus of technology. The days that we went through when we started - literally punch cards - punching a few holes and sticking it in a machine and sorting it. Now you pick up a phone and you can do all that and so much more. We don't have a technology gap; what we've got is a society gap.

Why aren't we using these tools to change medical systems completely? Something comes up; it's a $6,000 bill and the insurance company pays $630.

Burniece: Tell us a little bit about that. Where do you see the future for storage and computing? And what do you think are the key problems we ought to focus on?

Allan: The problems I see is that we have a few companies, with a vested interest in the status quo, and they're so strong they can prevent the level of change that their industry needs. The medical industry is an example of that, in that there's no good model. They have socialized medicine in other countries and in some places it works very well; others, you've got complaints. Unless you've got coverage with a company here, it's very difficult to live a healthy life. And we're a very wealthy society. I don't think that should be. I don't know the answers but I do know that it doesn't come from doing things the way you did them in 1932. You don't see the advanced use of anything in a hospital, except the equipment. The scans are better. They're high resolution. What about the process of handling somebody coming in the door as an emergency patient? It's all paper. You travel overseas; the same problem happens again and you're in Bulgaria. If that doctor knew that this happened previously and what the results were previously, he might be able to save your life but he has no idea. There's lots of privacy issues but technology could make it possible that your health should be visible, with your say-so, to any doctor in the world to keep you alive.

Burniece: A truly portable and available personal health record.

Allan: Yes. And that travels on from there. Kumar has a startup going in India where you can walk in for a smartphone EKG, which can be examined by a doctor in New Delhi, and you're somewhere else. That's already working in India and Kumar's part of that. If you'd have told me about social networking, when you grew up and I grew up, if it was behind closed doors, it stayed behind closed doors. Today it's videotaped; it's stuck on the web. That's a little loopy for me but that's transformed the society we live in. It could transform the medical industry too, if they wanted, and that would be a really positive end result. Social computing is fun; but is it an end in itself? I'd rather see that level of interest put into something, which could help, but those are things outside the control of the technologists. Those are political barriers and those are—

Burniece: Unfortunately they are, yeah.

Allan: Societal barriers. We don't use our technology to help all of us, which is unfortunate.

Burniece: So Dal just wrapping up. When you look back on this career you've had, what are you most proud of? You got to go and say what are the things you really are the most proud of you've done?

Allan: It's kind of conceited. But every year that I was in, something happened in that year.

Burniece: Okay. So you feel like you made some things happen each year.

Allan: I've been part of things that happened in every year where I was a minor contributor, sometimes a bit more. The only time I've ever bailed on a project, like AM-1, was that I saw no future for myself in this make-work effort, so I got out to do something different. If you give 110 percent to everything [you do], it pretty much usually works out in your favor.

Burniece: Yeah.

Allan: It's a case of working hard. Everybody works hard and if they do the best they can on a crap job, the next one might be a good job.

Burniece: You're still active, right? And what are you active in? Obviously the SFF and you're helping here with the storage SIG. But what else you doing?

Allan: Oh I've been a beekeeper for many, many years.

Burniece: Yeah? Well tell us a little bit about your bees; that's a fun story to end with.

Allan: It used to be dead easy. When we built the house I planted an orchard, saw lots of flowers and nothing happened. And when that occurred two years in a row, I got bees. I probably had them no more than six weeks when the staff meeting was interrupted because bees had built a nest inside the cable spool - nobody could build cables. It was a Friday and Art Silver looked down the table and said, "We know who's going to fix that. And we'll be back at work on Monday. Right?"

Burniece: So you were now the nominee.

Allan: Saturday I turn up to remove this bee's nest and everybody worked the weekend back then. The windows were full of watchers of me taking the bees out. Everybody knew somebody who knew somebody who knew me, so I ended up being called on to do removals around the valley. One of the callers was Ken Hallam who had a problem at his house; bees were in one of the artificial columns out front. That was a trapping job, so I spent about six weeks on the front porch with Ken after work once or twice a week, with a beer in our hands, just chatting. That led to an association; years later he left WD and joined me at ENDL. Beekeeping is a very social thing. It got me a partner—right? Bees have had all sorts of disease problems in the States for years now with problems coming in from overseas, and now

it's CCD, and so on. It used to be easy: throw some bees in a box and you had honey. These days they have a tough time living so I've been playing around since 2000 with different ways to keep bees, without chemicals. I'd like to get some of that documented.

Burniece: So that's one of your big life passions right now isn't it? Bees of course are critical to the environment.

Allan: Bees have no personality as individuals but a hive is a personality. You can't work your bees if you're in a bad mood. They feel it and they'll get mad and they'll go after you. I'd come home from ISS and I'd be just frustrated but if I went down the back in that mood, the girls let me know. I would have to detox and the act of walking down the back helped flush everything out. When I lifted that lid if I was calm, they were calm. You learn a lot from that.

Burniece: Yeah. Any last words?

Allan: Probably not, I've talked too long already.

Burniece: Oh you did great. Well this has been a delight. It's an absolute delight. We'll cut if off at this point and put it in the can and—

Allan: Okay.

Burniece: Thank you very much. It's been great.

Allan: Thanks Tom.

Burniece: Right.

Allan: It's been good.

Burniece: Thank you.

END OF INTERVIEW

REFERENCES:

1. http://s3.computerhistory.org/groups/ncr-cram-20121031.pdf

2. ENDL Letter 100th T10 Plenary Special Edition (November 2010)
   http://www.endl.com/letter/se100t10.pdf

3. ENDL Letter 100th T11 Plenary Special Edition (August 2010)
   http://www.endl.com/letter/se100t11.pdf

4. http://s3.computerhistory.org/groups/compaq-conner-cp341-ide-ata-drive.pdf

5. Computer History Museum Oral History of Kumar Malavalli, April 3, 2012 -
   http://www.computerhistory.org/collections/catalog/102746234

Initial edits by Dal Allan, February 20, 2014

Final edits by Tom Burniece, September 25, 2014

