



Interview of Robert A. Donnan

Interviewed by:
James Pelkey

Recorded: July 12, 1988
Atlanta, GA

CHM Reference number: X5671.2010

© 2016 Computer History Museum

Robert A. Donnan: I had a number of jobs at IBM prior to 1966 -- '65, 66 -- and was involved in SAGE, if that means anything to you.

James Pelkey: Absolutely. I had dinner last night with Frank Heart and I was with JCR Licklider last week --

Donnan: So you know some of those guys. Got involved in the early part of Mercury, and then early Gemini -- mostly work around Goddard, but still living in -- with IBM, we were living in upstate New York and going down there. 1966, IBM announced the Research Triangle -- 1965 actually. My family all comes from Chapel Hill, so it was: "Robert hold his hand up." I was one of the first ones down there. I had a couple of jobs prior to going down there, but about -- well, we Raleigh announced Bisync in January 1967. Raleigh didn't have any architects, so the up north architects of IBM wanted to tell Raleigh how to do business, so they decided they needed an architect, and they decided they needed an architect, and I got anointed, not even knowing what the term meant. My big job was to get the rest of the company to take in and buy this wonderful thing called 'Bisync,' and 'character oriented line control.' Out of internal arguments -- I formed a group, Communications Product Architecture. That same group occurred again some years later, with the same name, but with a couple of hundred people in it. My group was six or seven people. We started working on Bisync, and all the complaints about: "the thing had been put together in a hurry," for a whole variety of reasons, but we found some things, and finally, in 1969, we made our first disclosure about IBM's SDLC. My concentration was HDLC -- well, SDLC and then the external thing, HDLC, and ADCCP, and all that sort of thing. Up until I left, and everything was pretty well put to bed by the time I went off and moved to France in '77, and went and lived in La Gaude for three plus years. Got very much involved, at that stage of the game, with external standards. You know, you're more than a 100 miles from home, so you're an expert. I did a lot of going to ECMA and CCITT, and a little bit of ISO. Came back to the states in 1980, with the intention of retiring that year, and shortly after I came back to the states, my wife died, suddenly, so -- are you going to retire? Come on. So I looked around, and I had been asked to get involved in local area networks. I said: "Well, no, come on." Well, my wife died, and a few months after that, I said: "God, I got to go and find something to do," and I got involved in local area networks. So I've been involved with the thing -- this would make it -- from about the middle of '81, most of it standards work. IBM has something called, a position called 'Standards Project Authority.' What they do is they go around and anoint people, and say: "For this subject, you're our authority," and supposedly, it's supposed to keep people from running off this way, that way and the other way, and it also has something to do with conserving resources too. Prior to that time -- this is a little bit of an aside -- but prior to that time, of doing that sort of organization, you'd arrive at something like this, and see 15 IBMers, and that doesn't help anybody, from either a (unintelligible) point of view or efficiency. So, as I say, I was Standard Project Authority, as well as starting in with 802, very early on. I wasn't during -- it would have been, maybe, there had been some 1980 meetings, but when I arrived on the scene there were two groups, called 'DELMAC' and the other one was --

Pelkey: Hi-Li.

Donnan: -- Hi-Li, and DELMAC. So we initially started out, there was something called a token and there was something called a CSMACD --

Pelkey: Now, did that happen to you in the Phoenix meeting or the Denver meeting, do you recall?

Donnan: Oh, Lord, I don't know which one it was where they got split. Jerry Clancy was very much -- there was a turf war going on between Jerry and (unintelligible) anyway, all of which I wasn't -- I was just getting on board, so I wasn't really -- (unintelligible) --

Pelkey: But somewhere along the line there was a vote with a two-thirds requirement, and Token got more than 50%, but not two-thirds, and somehow CSMACD --

Donnan: This is prior to my time. I've heard the story about "three more votes and we wouldn't have CSMACD today," and, come on, you guys are kidding. It's a fact of life. You've got to be a little bit more pragmatic than that.

Pelkey: But at one point, when they split between CSMACD and --

Donnan: When I got there, the split had been made.

Pelkey: But there was just one Token committee?

Donnan: There was one Token committee, that's right.

Pelkey: What caused the split into two Token committees?

Donnan: It turned out that, just constantly, the needs were divergent. The Token Ring people saw a very fancy priority scheme, for instance, which is in Token Ring today, with which you can do all sorts of weird and wonderful things with that aren't even possible in Token Bus. This has to do with the nature of everything going by every station. There were (unintelligible) characteristics in Token Ring, which if you say: "You make Ring and Bus look alike." Well, what you do is you reduce it to one common denominator, and you get the worst of both worlds.

Pelkey: Do you recall what meeting -- exactly how the split took place?

Donnan: It would have been either late in '81 or '82, and I don't know exactly -- who the devil would know?

Pelkey: Did you become chairman of --?

Donnan: Yeah, I was the initial chairman, and have been ever since.

Pelkey: And someone else became chairman of the Bus? Did Clancy remain?

Donnan: Clancy was around -- well, when we split up, there were two people that got -- well, three people that got dethroned: Clancy, Nathan Toble, and Nathan's co-chairman -- he and Nathan were co-chairmen; I can't remember his name at the moment -- but he's from up here, from RCA up here. All three of those individuals were made ex officio members. Nathan's working buddy on DELMAC -- no on physical, that's right, DELMAC and physical. The other one was physical. Hi-Li didn't really exist yet. There was DELMAC and there was physical --

Pelkey: It was originally Hi-Li. Then Hi-Li became the two.

Donnan: Well, Hi-Li is really 802.1 today. The high-level interface. Everything that went on outside the thing. As I say, just the evolution of things, Nathan's the only one left. Jerry Clancy eventually didn't attend anymore, and the other guy -- and I can't think of his name -- went away. So (unintelligible), we've got a structure of chairman, and executive secretary, corresponding secretary. We had -- just an aside here -- we had dramatic things, like we had a secretary who just did -- I don't know how he ever did what he did -- he committed suicide, and we had these little events, just living.

Pelkey: Can I ask you a couple of questions about -- I haven't talked to any of the people in Zurich. You're the first person who formally worked for IBM that I've talked to. The book is mostly about the non-IBM world, as opposed to the IBM world. I don't deal with SNA very much -- not that it's not important, it's just that I have chosen to exclude it -- but the issue of where the Token Ring came from, there are some people that think it came because of the Cambridge Ring, and other people think it really came out of MIT, and Jerry Saltzer, doing some consulting. Other people think it came from the Loop --

Donnan: Well, it's very hard -- from an IBM point of view, it definitely came from Zurich. That is their view, and just recently they passed out some kind of fancy money to Zurich people for some initial work done.

Pelkey: Where did the idea for the Token Ring come from at Zurich?

Donnan: Well, it's -- the two guys that got award were Werner Bux and Hans Müller. Now, Müller is a physical kind of guy, who worried about phase locked loops and that sort of thing, and how to take it -- some of the initial work that was done, how the devil to get all these things without wagging the tail. Raleigh had built some rings, 2400 bit per second rings. They had all kinds of trouble, using 201 modems, that sort of thing, all kinds of trouble in trying to track, when you started cascading modems one behind another, but that notion of trying to put things on ring -- there's been the Storer -- the Storer Loop; the supermarket loop; the retail loop; there's the banking system; the Kingston System -- I can't remember any of the numbers. Those things have been around. The idea of a ring -- some of the first work that was done -- the prime work that was done in these Farmer patents, the AT&T patents. I, for some reason or another, was present at that first presentation, down in Callaway Gardens, GA. This would have been '67 or something like that.

Pelkey: Of Farmer's work?

Donnan: Of Farmer's work.

Pelkey: What was the occasion?

Donnan: It was a communication conference of some sort. Callaway Gardens, GA.

Pelkey: Was Dave Farber there, do you recall?

Donnan: I do not. The only one I can remember -- I can't remember who my coworkers were. My boss was with me.

Pelkey: Was that an exciting period of time in terms of these ideas, or was it more conceptual?

Donnan: I don't think it was any more exciting than anything else. The notion of rings, of the idea of block insertion into rings, register insertion --

Pelkey: Register insertion.

Donnan: I've got a couple of patents on that thing. Everybody did. There's more damned patents on the thing, and nobody used it. The register insertion idea; the idea of the Token itself, the guy who claims the rights to it is [Olaf] Söderblom, and he holds the patent. Now his work is the idea of putting some sort of thing on the trailers. Taking (unintelligible) go ahead for the next station down line, or whatever, and that's really the basis of both the retail and supermarket ring. They take the SDLC 'go ahead' is a zero and seven ones, and when you see the seventh one go buy you change it to a zero, and say that's the start of the next frame, and you just pile those frames, one in behind the other. Except that thing's not locked in itself. That's the big change. That's a loop, not a ring. It's an open-ended ring, if you like. So you've got a primary station, a secondary station. He kicks off a poll, or he kicks off a flag, everybody jumps on and everybody in behind it. The big thing was how to get that thing to run it --

Pelkey: Synchronize up?

Donnan: -- I think it was more a physical solve, a physical problem, than it was -- how you build a buffer to accommodate these things and how the devil you can take (unintelligible).

Pelkey: I'm not an engineer, so bear with me. I visualize a lot of the influence, the orientation of IBM towards these issues, may have stemmed back to the way IBM thought about terminals. That is, rather than being this dumb terminal, you had a terminal where you did your error correcting, so very early on, IBM started around this notion that kind of led to the synchronous protocols; this idea of --

Donnan: Really, two notions were of --

Pelkey: -- and that's where this intellectual property that led to the synchronous, and the Tokens and -- they all kind of stem from these fundamental thoughts about how to organize a computer system.

Donnan: Well, they went through a period of time, when logic was still expensive, where you took and you shoved as much logic back in the host as possible. You shared the logic. The IBM banking system is that way. You literally -- you hit a character, it puts a character on your screen, it builds up all the processing that's taking place back there, and that was a period of time when the balance of costs were that way. The opposite of that is, to a degree, the Storer System, where you start putting more intelligence out into the terminal, and of course it's gotten to that extreme nowadays. You put the PC out there, XT, what have you, and the line becomes more important, in terms of how much stuff you're shoving down the line.

Pelkey: It was very natural, because of the synchronous protocols, and then IBM dealt with polled modem networks, in terms of the kinds of communication environments that lent themselves to IBM's way of working. Token Ring was kind of a natural conclusion to that. It was synchronous, it was deterministic, it was consistent with the heritage of IBM.

Donnan: Well, the deterministic -- that's a bunch of conservative IBMers, the whole idea of determinism, but yeah, some of these systems are very old. The Sabre System, which was a polling scheme -- was it American or was it Pan Am? Well, it's unimportant. One of the airlines --

Pelkey: That was American, Sabre. But I heard an old story about the president of TWA and Watson met each other on a plane one time, and that's how TWA, which had the first system, and Watson on the plane in first class, in an old plane, convinced him that he could put a modem network in and make this thing work, and that's what led to the first reservation system.

Donnan: That's possible. My favorite Watson story is about the IBM typewriter. They decided it was losing business, and they announced -- or guess they privately announced -- they were going out of business, and it was either John Hancock or Prudential or somebody who said: "You're not going to make that black model typewriter anymore? Well, take all your machines out." Called up the boss and said: "Take all your machines out." "Pfwheet. I guess we ought to be in the typewriter business after all."

Pelkey: There are people in the world who would say that IBM had its cake and ate it too, with the Token Ring standard. It's a standard, but at some level it's not a standard in that IBM does it a little bit different than the TI chip set does, so while it looks like it's a standard, it's not the same as CSMA/CD, in terms of a wider group of people who all have rights to the same technology. Is there any basis for that view?

Donnan: I can't answer you with any great authority. I know that certainly there are things that got into hardware long before we ever got them standardized, and to that degree -- I know of no differences in the IBM hardware that makes it incompatible, for instance, with other implementations. I can't think of any -- I'm trying to think of a -- well, let me give you an example. We transmit something continuously in the standard, and we provide means of making sure that, before we switch, that these sequences are off the line, because if they're not off the line before we switch, something funny is going to happen. I think the IBM implementation doesn't transmit them continuously. Instead, it transmits one every so often, and with enough interval between them to be sure that this thing does not take place. It's strictly a station on itself, so to the degree that it doesn't affect any other station -- this is a very funny area. It's a very funny area because what your concern is how are stations going to work together, yet when you stop and describe things, you describe the whole process. Some of those processes being -- we'll describe to you how you should tie your tie in the morning, but what the stand is; you wear a tie. There's quite a bit of that in Token Ring, because everything is -- they can get a lot of intelligence out of the stuff coming back, so you have to be very careful about how you get this stuff off the line, and how you insure against some of the things you have done aren't affecting other people. An example: if you transmit n number of frames before you get off the line, there is something called an 'intermediate frame bit' which says this isn't the last one, this isn't the last one. Finally, you get to one that says this is the last one, because you're taking

them all off the line, the stuff you did. Another way of doing it -- perfectly legitimate - - is you could have counted them.

Pelkey: Right.

Donnan: Another way to do it is you could have knocked them out with something called 'abort delimiters.' If you talking about standards, pick one. (unintelligible) IBM implementation doesn't have any of those problems because it only transmits one frame at a time. We've picked one, and yet going on in committee right now, two years after the thing, or three years after it became a standard, they said: "Well, why have we -- it's much easier to count than it is to use the i-bit." Well, the I-bit's in the standard, and one of the arguments going on in the committee right now is: "Should we really go back and allow the counting of frames in order to purge the ring?" So there are those differences, but it doesn't create an incompatibility, and that's the concern. I guess from a commercial point of view, if you go to a product manager, he says: "How does my product fit?" And you say: "Well, your product does fit, but wouldn't you like, cosmetically, to look pretty?" That's awful hard to justify. Now, if it means you don't get a contract, that's a product manager judgment as to whether he spends the money to correct this thing. In today's world, with chips and all the rest of it, the cost of doing that can sometimes be pretty -- I was away from the scene even when I was still working for IBM, because most of my concentration was on this, and sometimes you get too involved in what's going into an implementation, you get --

Pelkey: Buggy.

Donnan: You do: "Don't tell me about your troubles, just implement the standard." They say: "Well, we don't like this." "I don't want to hear about it," and particularly since I've gone." Now it's (unintelligible). I can really claim --

Pelkey: When did ISO begin to impact your thinking?

Donnan: Mine? Well. Probably in the early '80s. We were SNA-ing it. SDLC became the base -- early '80s. Whatever --

Pelkey: I'm getting the sense it is not your decade.

Donnan: I slip all the time.

Pelkey: North Carolina is like that.

Donnan: Well SNA -- there was this architecture group. I was architect for a while. Then a fellow by the name of Sussenguth came in, who was Mr. SNA. He's an IBM fellow today.

Pelkey: Was SNA in '76?

Donnan: No, it was earlier than that. SNA was, maybe '74.

Pelkey: Because the battles of the late sixties were X.25 versus SNA, at some level.

Donnan: No, no, the sixties?

Pelkey: Seventies. Mid to late '70s.

Donnan: We announced in '70? Yeah, '74 I think it was announced. Then they took and they want open systems interchange. Well, what does it look like? What does the model look like? So there was a whole bunch of arguments about what the model did look like. Everybody pretty well knew what data link control was. Everybody knew what the physical layer is. Of course, we find out years later that we drew the line the wrong place, and they're constantly trying to move that line. What they'd really like to do is

have a data link control that is media independent, but they don't. Whereas, if they'd move the line up to where it is in LANs, they'd be ok, and people realize that.

Pelkey: It's a big decision.

Donnan: Well, it's constantly talked about by the various country representatives. It might happen, but I don't know. To go back, the physical level was kind of understood. The data link and network layer was understood. Above that was a great fuzz. The IBM SNA thing just draws the lines in different places, and people didn't want to . . .

[Interruption in Interview]

Donnan: I was not involved in that part of the thing. I'm a physical kind of guy. I'm a data link -- but there wasn't any big argument going on there. It was kind of HDLC was going to be that. That was it from the beginning, and the line was drawn, and that is the precise line between either IBM or ISO. That part of the model is pretty clear. It was these things up above where there wasn't any work going on at all where all the discussion took place, as to what it should look like. Pragmatically, it couldn't look like IBM. There was no way it could look like IBM.

Pelkey: Did you or IBM ever get involved formally, in terms of the NBS workshops?

Donnan: Oh, a lot of IBM was involved, sure. They still are. I was not.

Pelkey: You stayed on the 802 portion?

Donnan: Yeah. Now, since coming back, since 1980, I've been local area networking, period.

Pelkey: That's been a very important process, this getting the standards made, and this 802. It went from not being a standards committee at all in '79, to now being recognized as a pretty good standards making body; very professional and very competent.

Donnan: It's amazing, when you consider our rules and whatnot, how we get anything done, and yet we do. The structure of this doesn't look like any other standards body on earth.

Pelkey: In the sense of how different?

Donnan: Well, you get voting rights by the fact that you're an individual attendant. You're not one vote per company or one vote per country.

Pelkey: It's a vote per attendee, but you have to kind of attend meetings in order to vote, though, right?

Donnan: You have to attend meetings. We keep careful records as to people here --

Pelkey: Someone can't just walk into one meeting and vote?

Donnan: The official rules are 75% in a committee, and two of the last four meetings, but anybody can -- the only rule is that you come to the meetings, contribute technically -- and the thing is somewhat self-disciplining. If somebody is obnoxious --

Pelkey: You find a way to deal with it.

Donnan: Well -- don't repeat the message I gave to Maris -- but that was one of them, and the committee finally, after a year and a half, took care of it.

Pelkey: It strikes me that this process that the group of you have gone through is very important and very instructive, in terms of how organizations in the future may have to function, making decisions.

Donnan: I'm not at all sure that we're not unlike what the Japanese do today.

Pelkey: I'm not sure that we're not either.

Donnan: During the period in the middle '70s, IBM was in the process of opening the Fuji's Island Lab, and it was right around the SNA announcement days. We had Japanese in the Raleigh lab. We had Japanese sprinkles all over the place. They were there for 30 months. We had a lot of fun. An aside: when it came time to go home, none of the wives wanted to go. Just the modus operandi is entirely different; not at all different than this.

Pelkey: It strikes me that way.

Donnan: The only part of the Japanese process that is not like here is that the boss is sheer God. That was one of our prime things to do, incidentally. Nabi Mei, who is today a vice-president of IBM Japan, his talk to us and instruction to us -- talking about the Americans -- will you please teach my people how to tell me I'm full of shit. He says: "I'm just hung out. The only people I can depend on to express on a strong point of view are the Americans that I work with." "Nabi, you don't know what you're doing." I think he's been successful, at least to some degree, really overcoming that cultural problem the Japanese have in that particular respect. I can remember Nabi --

Pelkey: They're real conforming.

Donnan: Well, the boss said: "Well, I think that's what I'd like," and bang, that was it.

Pelkey: I agree. I think the experience that you're going through is not unlike the Japanese, and I think it may point the way to how organizations are going to have to start thinking about working in the future, as anybody can talk to anybody over these networks, and everybody gets connected. Bosses aren't going to be -- information is not going to flow up, it's going to flow across.

Donnan: That has been true at IBM for quite a period of time, and it was more by the technology. There's something called the VM Network -- everybody in the world. I sit in Raleigh, and I've got a compatriot in La Gaude, and we turn documents around every day. I'd, by the end of the day, have something or other. "Pierre, how about reviewing this thing," and sending it off to La Gaude at 6 o'clock that evening, and I came back to work the next morning, and Pierre's reviewed the thing and everything else, and I call it up and there it is on the screen. If you really want to go live, you can even go live on the thing. It's a little -- it's a packet, so those things existed -- they existed in the middle '70s. We were really getting going. Today, that internal IBM network is just out of this world.

Pelkey: You couldn't work without it. But it caused more peer-to-peer, as opposed to this kind of communication.

Donnan: Oh, absolutely. You work in a department, then you're supposed to understand the mission of that department. You're not carrying it over to you boss all the time and say: "Look at this." It's a little too late by then. Particularly in the area I've been working in for years, this architecture business. It's an exchange of ideas.

Pelkey: Do you think we can survive with fewer levels of management as a consequence of this kind of communication structure?

Donnan: I don't know. The best I ever got was the second level. I never understood this -- I never did understand what's going on up there. You would logically think so. 'Levels of management, to me, is some function of 'span of control'. I never felt comfortable, as a manager, never felt comfortable in an

intellectual environment with more than six or seven people reporting to me. Otherwise, you didn't have any feel at all . . .

[Tape Side Ends]

Pelkey: IBM, in the mid '70s, you had some kind of a loop front-end processor that was kind of like a ring --

Donnan: That's what I was talking about: the Storer Loop is a supermarket only retail --

Pelkey: This was a vertical market product. This concept of Token Ring --

Donnan: That was like a Token in a way.

Pelkey: There were enough of these ideas floating around that it was not unnatural for the Zurich group to have --

Donnan: Oh, no.

Pelkey: -- try to embody these in a more formalistic way, plus there was stuff like Farmer's work was out there, and Newhall and Pearce and Farber --

Donnan: It's an idea whose time was right.

Pelkey: I don't think it's a Cambridge Ring, because the Cambridge Ring was a slotted ring, and that's very different from what was developed in Zurich.

Donnan: The Cambridge Ring is very much akin to the IBM banking ring. The Cambridge Ring is built on the notion of real inexpensive terminals -- in my view -- real inexpensive terminals, or the least expensive of terminals, and a central processor. So inside of IBM, there was, as I say, the banking that came out of Kingston. Boy they were for that. That was the right way to go, and (unintelligible) Raleigh, simultaneous. The Raleigh work people were taking a modification of SDLC, change the flag character on the end, and were piling that stuff on, and: "Oh, look how expensive that is," and to me it's really a function of the cost of hardware, that sort of thing, what you can do at any point in time. The Zurich people were just looking around, and that was a good study project.

Pelkey: Thank you very much for your time.

END OF INTERVIEW