Oral History of Greg Ennis

Interviewed by:
David Cohen

Also present:
Marc Weber

Recorded: June 26, 2012
Mountain View, California

CHM Reference number: X6535.2012

Cohen: Hi, I'm David Cohen, and I've been involved in the Wi-Fi industry for some time. I'm currently with a company called GainSpan Technology, and I'm sitting here with Marc Weber, who is part of the Computer History Museum, and we're here to talk today to Greg Ennis, who has had a long history in wireless communications, and 802.11 and Wi-Fi technology. We're going to continue with Greg and we're going to talk a little bit more about some of the early years. We talked a lot about the technology, but tell us a little bit more about yourself. What did you like to study in school? What interested you? What did you pursue and why?

Ennis: I was always a good student. I've always enjoyed learning. To this day I still wish I was back in college. I would love to be a college student forever, but I figured out early that to make some money you do need to go out and work. So as a kid I was always interested in science and astronomy. As a young kid I was good at math. I can't say I was really interested in math, but one thing to mention is that I'm a musician. From a very early age I have been playing piano. I took classical music lessons for years. In fact I can say this here. John Markoff, who you may know, I took piano from John Markoff's father in Palo Alto. So I think that there's certainly a connection between music and math. I'm still a musician. I'm a jazz musician. I've played in jazz groups my entire life. I still am gigging. So that's always been part of it, and I think that ties in with my math, but I'd say I first really connected with math probably in the 7th grade. There was a program, an experimental program. It might have been developed by SRI, Stanford Research Institute. It might have just come out of another part of Stanford, but it was called SMSG Math. I can't remember what that-- it was like the Stanford Mathematics Study Group or something, and it was a new curriculum for elementary school math. Remember, this was back in the days of the space race, right? We had to compete with the Russians, and so there were a lot of efforts to improve math education for elementary school students. I remember we had these little paperback books, yellow covers, and that was just typed. It wasn't typeset it was just like Xeroxed. What we were studying, we were doing base-seven-arithmetic as seventh graders. Somehow I doubt that they're teaching that in schools these days, but it was so radically different. This was the new math, right? You've heard of the new math. This was the stuff where the people were complaining that the kids were coming home with their math homework, and the parents had no idea what was going on. That's because we were adding three plus four that's zero, carry the one. Of course that ends up translating into binary, and octal, and hex with computers, but back at the time it was just an attempt to get kids to recognize that there's nothing magic about the number ten other than the fact that we happen to have ten fingers on our hands, and that you could just as well do all this stuff with a different base. I loved that stuff, and I had a tremendous math teacher back then, Mrs. Bunting, and that really got me, for the first time I think, excited about math. But then it was when I was a freshman in high school taking high school geometry and, again, I'm not sure that they do this now. I really wish they would, and I think they should. But when we were doing high school geometry back then it was all proofs. It was all axioms. It was proving theorems. You had to prove that if you have two chords in a circle that intersect on the circumference with a diameter that that's going to be a 90 degree angle. You have to prove that from the axioms. That was totally the focus of that. We did nothing but proofs for a year, and I was really good at it, and I absolutely loved it. Then we went onto to do Algebra II and trig, and it wasn't quite as exciting because I really loved that real proof orientation that we were taking in geometry. So then I got into Berkeley, and I started out as a physics-major because that was the flashy major to major in at Berkeley. So, of course, I took math and physics both during my first couple [of] years. As I took more and more math, and we started to get to where we're focusing on proofs, we're not focusing on calculations I came to love it again. I quickly realized it's the math that I love here, it's not the physics, and I became a math major. That led into the math that I ended up doing in graduate school, well also as an undergraduate at Berkeley, was mathematical logic and set theory. You know the book *Gödel, Escher, Bach*? The type of math, that Gödel Incompleteness, that's what I concentrated in, in my college and graduate year math studies, and it still is my first love. There's no question that math is. I still read math for pleasure when I'm in bed going to sleep.

Cohen: Do you think that has a direct tie to, ultimately, your career in technology and wireless, or did it kind of interact in any way?

Ennis: I think math education is, I mean, especially as you get up to the upper division math, math beyond calculus, that stuff, it's the best brain training there is. So I don't know if it's particularly-- I can't say that I have ever used the higher math that I learned in my graduate studies, never have used that in my technical career as an engineer, but the way to analyze a problem, the way to understand concepts, that's where I think that all came and helped me as I became an engineer. Of course I was always interested in computers, and I really shifted my focus when microprocessors came out.

Cohen: What drew you to computers? Why were computers exciting at the time?

Ennis: When the Adam Osborne book came about the 8080, when the Intel 8080 first came out, microprocessors first came out, I'm sure it wasn't just for me, but for a lot of people that there is going to be some interesting stuff happening with this. It was right around then that I then went to get a second master's degree at Stanford in computer engineering, because that really grabbed me, and plus I thought I could probably make some money at it.

Cohen: It was after you learned about the microprocessor?

Ennis: Yeah, it was when the microprocessor came out. To me, that revolutionized things for the interest that I had in computer technology.

Cohen: You said you were a musician from early on taking piano lessons and all that, and have continued that through your life. Can you talk about that? Why did you pursue that, and why do you still pursue that, and how does that tie to your training? Comment on that a little bit.

Ennis: Well, I pursued it because, well, I mean for one thing music is a beautiful pursuit, but I come from a very musical family. My grandfather was a professional violinist. He played in the St. Louis Symphony. I have a beautiful formal portrait of the St. Louis Symphony circa 1911 in the orchestra hall, and there's my grandfather right there. Then my mother was like a child prodigy pianist who performed on the radio as a kid. My sister is a singer. I started piano, and I had the best teachers including my mother. I played classical all through my childhood. I still do play classical. I still love classical. As a typical teenager back in the '60s I became enamored with rock music, and I very quickly saw that at least as far as the piano and keyboards were concerned there's nothing interesting going on in rock music. All the action is in jazz, and I became a jazz pianist starting around age 19, probably. I played in bands in Berkeley when I was in college, and I've played jazz ever since.

Cohen: Jazz piano?

Ennis: Jazz piano.

Cohen: Still to this day?

Ennis: Still to this day.

Cohen: That's great. Who were some of your jazz influences?

Ennis: Chick Corea, Herbie Hancock, Keith Jarrett, Wynton Kelly, Bill Evans.

Cohen: A lot of greats.

Ennis: Yeah, all the greats.

Cohen: Other things that you like to do to relax and enjoy besides reading math books at night and then playing jazz, other interests?

Ennis: Well, I currently live up in the beautiful Sierra foothills where I've got an orchard and a creek through my property, so basically homesteading a beautiful piece of California land is a big thrill for me and my wife right now. I enjoy basketball.

Cohen: I think that gives us a pretty good background of your life growing up before you dove into all the, or even throughout the time that you were working in wireless, so that's good. Good, thanks.

[This section below came before the above]

Cohen: So, Greg, welcome. Good to have you here.

Ennis: Thank you, David.

Cohen: Nice to see you. It's been a while. We're going to start off just talking about your early years of your life and your career. Why don't you tell us a little bit about where you were born, where you grew up, where you studied; that kind of thing?

Ennis: Well, I was actually born at Stanford University Hospital. I grew up in Menlo Park, which is right next to the Stanford campus. In fact, the Stanford campus was essentially my playground as a kid, which was a great playground. So I went to schools in Menlo Park, went to Menlo Atherton High School in Atherton, and then I subsequently went for my undergraduate degree at UC Berkeley where I was a math major at Berkeley.

Cohen: After studying math at Berkeley where did you go onto study from there?

Ennis: So I went on to pursue graduate work in math at the University of Wisconsin, and that I actually took a minor in computer science at the time, although, even at Berkeley I was taking some computer science classes. In fact I took Engineering as a freshman. We're talking about 1969 here, so we're talking about punch cards and Fortran-programming, but math was my real love. I went on to pursue further graduate work in math at The University of Wisconsin, but with a computer science minor. I was focusing on mathematical logic, actually, and set theory. So mathematical logic, I guess, has some connections to computer science. That might be what ultimately drew me to outright computer science. I ended up with a master's degree in math from The University of Wisconsin, but then I really got interested in computers, and I went to Stanford where I got a second master's degree, this time, in computer engineering. So that was in 1979, and that was basically when I then started my career in Silicon Valley.

Cohen: What were some of your first professional stints in Silicon Valley?

Ennis: The first job I had in Silicon Valley was with a small company that was literally starting in a garage. It was a pretty nice garage, but it was a garage. It had garage doors. The company was called Sytek. It was started by a group of five guys who all came from Ford Aerospace, so they came out of the defense contracting world, and they were starting up a consulting and contracting company to pursue that kind of work, not just defense related, but also commercial projects as well. I started as a Junior Engineer working for that company. I was the ninth employee. I think by the time I left some years later there were probably over 300 employees, and we did a lot of consulting work, initially, for government clients and for commercial clients, but ultimately the company decided to pursue product development, and in particular the products that we were developing were so called broadband local area network products. Broadband back in those days didn't have the connotations it has today with the Internet. What it really meant was using cable television technology as the underlying wired medium to support both local area network and metropolitan network data and other applications. So I ended up at Sytek as Director of Engineering, so I was responsible for a lot of the product development that went along with that product line.

Cohen: What kind of local area network products were being developed at that time?

Ennis: This was back before PCs, really, so we're talking 1979, very early '80s, so I guess Apple had probably come out with their initial PCs. That was just about the timeframe when IBM was coming out with the IBM PC, but the applications for local area networks, at the time, and this isn't just for Sytek, the company I was with, what we were doing, but there were other companies, Ungermann-Bass and Bridge Communications, and these companies were, actually, they were all located right here near The Computer History Museum. This used to be called LAN Lane over here. Sytek was on Charleston, and I think 3Com was over here. I know Bridge was over here, so it was called LAN Lane. But the applications were supporting what we would call dumb terminals, just a CRT screen with a keyboard that you would use to connect as a terminal to a computer like a UNIX computer or a VAX, a larger computer somewhere. The real market for this stuff was that if you think of a university campus, or a military base, or some kind of a large installation like that, you would typically want to be able to connect to multiple computers. The way that people had to do that previously was that you would have to have a terminal that would connect you to this computer, and another terminal that would connect you to this computer, and another terminal, and this would all be on your desk. So instead, people wanted to have a single computer [terminal] and you would be able to connect to multiple computers, and that's what the local area network technology did that Sytek developed, and then other companies were pursuing this as well. That would basically be a box that would connect to the cable network wires, and then you would plug the terminal into that, and through a command interface to that box you would be able to command it to connect you over here, and then you'd do your session to that computer, and then when you're done with that you would connect over here. So those were really the initial applications of local area network technology. Then, of course, the personal computer came out, and now it's a totally different ballgame, because now you want to be doing file sharing. You want to not just be connecting from a dumb terminal to bigger computers. You want to be connecting between PCs and the file servers. This was right when 3Com really got started with their products that were focused on local network Ethernet applications for personal computers.

Cohen: So what timeframe is this, about mid 1980s?

Ennis: We're talking early 1980s, early 1980s. Sytek actually got the contract from IBM to develop a local area network product line for the IBM PC. It was called PC Network. That was actually where NetBIOS came out of that, that particular effort. I was very involved in the development of that. So Sytek ended up being a contractor to IBM developing this network for the IBM personal computer product family. Ended up not being a huge success on the market because another division of IBM was developing token ring technology, and there was an internal IBM battle, but that's another story.

Cohen: So what happened from there? Obviously, you got some interest and a lot of hands-on development experience with local area networking technology, so where did that kind of LAN technology take you next?

Ennis: I started my own consulting business in 1989, and that consulting business basically as a single person. It's not like I had a bunch of staff. I just operated as an independent consultant, and I had, through the '90s, a wide variety of consulting clients in all sorts of different technology areas. I got involved in some very interesting technology. For example, I was one of the system architects for one of the early wireless airplane phone systems. You still see these sometimes, on commercial jets. In fact, I saw one just the other day where the phone was in front of it, and then they had a sign on it saying, "Not operational," because those things haven't been operational for years. But back at the time that was a real service that was being provided to passengers of commercial airlines. So it involved a tremendous integration of all sorts of different technologies, technologies to connect the phones within the planes, plus all the wireless technologies to connect down to the ground station, the development of all of the base stations, etc. I was a system architect on a project like that. That was just an example of one of the projects. I did a lot of projects continuing to involve cable television technology, to providing voice services over cable systems, providing meter-reading services over cable systems, doing products that today we would call cable modems, although at the time we wouldn't call those cable modems. But what's interesting when I think back on the connection between those projects, and my subsequent Wi-Fi and wireless experiences, cable systems really use radio technology. I mean, when you're transmitting data over a cable system, over a CATV system, it's really a radio that you're using. You're using radio frequency transmission. We would call them radios. It's just that the radio waves are constrained to go simply over the cable system rather than broadcast out into the airwaves.

Cohen: It's essentially RF over Coax.

Ennis: It's RF over Coax. So in practice, much of my career, even though I can't say that in my early days I was necessarily focused on wireless, I was certainly focused on RF. In particular, I was focused in providing communication services, data communication services using radios, and all the issues that come to bear with that type of technology even though, at the time, it wasn't strictly wireless. But then, as part of my consulting work, I started to get involved in some specific wireless projects.

Cohen: What were some of the first projects that you got involved in that were specifically wireless as in non-wire?

Ennis: There were some miscellaneous things, but the first one that I think is probably significant here it was a project for The Chicago Board of Trade, and the Chicago Mercantile Exchange. So, again, we're talking this is probably around 1991, yes, something like that, maybe 1990. What had happened was that the trading floors, you know, these are the big commodity exchanges in Chicago, and the trading floors, at the time, were completely manual. The traders would be shouting at each other across the trading pits to decide to make a trade, and they would be writing something down on a card. It was amazing. I took a tour of this. That was part of the project here. These traders would literally throw a card down on the floor, and there would be runners with color-coded jackets who would come up and pick up those cards. Then they would go into a backroom, and they would try to match these cards up and come up with what the trade was that these two traders had actually agreed upon. That was ripe for fraud, and there was actually an FBI sting operation. I think it was at both The Chicago Board of Trade and The Chicago Mercantile Exchange, and the SEC demanded of those exchanges that they automate their operations, and computerize them.

Cohen: Because of significant fraud that was occurring?

Ennis: Because of fraud. So there was a consulting contract. I was working with other people on this. We got a contract to develop a wireless handheld trading terminal. This is 1991. I mean, basically this terminal it was also going to have a graphic display and a touchpad. They wanted it to emulate as close as possible these physical cards that the traders were doing. My responsibility was doing the architecture of that wireless network. I mean, I previously had experience designing protocol architectures for a variety of different underlying mediums including radios over cable systems like I talked about, so it wasn't a huge stretch for me to be involved in the design of these wireless protocols to support these trading terminals. We're talking about probably a thousand training terminals on a floor, and having base stations, etc. [It was] That experience of designing a protocol-architecture for that wireless trading terminal system for The Chicago Board of Trade that ultimately led to the projects that I got involved in that eventually became Wi-Fi.

Cohen: When you say protocol-architecture are you talking about something that is similar in concept to ultimately the 802.11 networking layers like you got involved in? What would the physical layer be like? What would the medium access control layer be like for this Chicago Board of Trade, Chicago Mercantile Exchange systems?

Ennis: Exactly right. It would be what's called the access method for the radio system so that the devices can fairly share the airwaves bandwidth among each other, and also for the reliability of the communications to make sure that packets, as they're generated, get addressed and delivered to the right recipient, and that there's acknowledgements, all those aspects of the protocol. Back in the 1991 timeframe there were no standards for wireless at this time for this type of a system, so it would be a proprietary design that I'd be developing, but ultimately the standards effort started up within IEEE to develop multi-vendor standards so that multiple manufacturers could develop this stuff, and all work together.

Cohen: What were some of the other companies that were developing - pre-standard - developing some of these wireless access, wireless networking access systems? Were there others that were competing for The Chicago Mercantile Exchange business, or other major projects that were going on that you're aware of?

Ennis: Yes. So there was a company called Proxim who they were very early on in this. The company that I ultimately ended up consulting for was a company called Symbol Technologies. They have since been merged in with Motorola. Symbol, back then, they were in the business of providing handheld barcode scanners for inventory applications, and for warehousing applications, for retail applications where people could go around and scan barcodes. I think Symbol had some of the fundamental patents on barcode technology. That's how they got into this, and then they developed these wireless handheld terminals. So, Symbol was definitely already starting to do this wireless local area network type technology. Another company that was very early here was NCR. Again, same type of applications that NCR was looking at, these vertical market applications for warehousing, and for barcode scanning, and this kind of thing. I started consulting for Symbol Technologies and, in fact, we ended up working jointly with NCR in developing some suff that ultimately became the IEEE standard.

Weber: Were there any that you took as models when you did the system for the Chicago trading floor?

Ennis: Well, not so much on the wireless side. I mean, I had done protocol design, and I was working with other people on those projects who came from Sytek, which was the company that I worked with originally out of Stanford where we had done these metropolitan cable systems and projects like that. So we all kind of came out of a background of knowing how to do proprietary architectures for advanced networks. So the wireless application that we had for The Chicago Board of Trade, although there was nothing that we really could look at that was anything like that, I think the group of people working on this we had the kind of experience to pull that together.

Cohen: Just a clarification question. So how many of you were working on that Chicago Mercantile Exchange Board of Trade wireless protocol?

Ennis: There was probably a core group of about three of us who were really doing the architecture of those protocols, but that project, as you can imagine, involved a lot of things. I mean, there's the actual physical hardware design for the trading terminal. There was a company called Synerdyne in Southern California that was actually developing this handheld trading terminal. There were other people, but I'd say as far as the core architecture of the wireless network it was me and then there were a couple of other guys who were instrumental in that.

Cohen: So now we get to the point where you're working at Symbol, and you've got NCR also working on these vertical inventory barcode scanning wireless applications. Talk a little bit, if you can, about how that developed, matured, and how that ultimately led to a standardization effort.

Ennis: So Symbol and NCR were bitter competitors. There was a guy, one of the head guys at Symbol used to describe them as “compestomers”, or something, trying to combine competitor and customer, because we also sold stuff to each other. So right in that early '90s timeframe, an effort did start up within IEEE to be IEEE802.11. There was the IEEE802.11 committee. That committee was starting to come up with general requirements, and kind of some basic ideas. I started participating in that, I believe, in 1991, and the guys from NCR were participating in that. We ended up seeing that there were some real philosophy-similarities in the way that both NCR and Symbol wanted to approach a wireless local area network, as opposed to the philosophies or the architectures that some other proponents were putting forward, in particular, IBM. IBM had put together a concept that was a very centralized system. So there were these competing proposals that kind of took a radically different approach, and it turned out that the NCR approach and the Symbol approach were compatible. So we decided let's put together a joint proposal. We started working, I mean, in particular I started working very closely with one person at NCR named Wim Diepstraten who was one of their main architects. This was a group at NCR based in Holland. They subsequently got sold to Lucent. I mean, they're a major player through this whole story.

Cohen: That's right.

Ennis: I started working with Wim Diepstraten on the development of a joint proposal that we could potentially put in front of the IEEE 802.11 group. I had had previous experience back in my Sytek days. I chaired a subcommittee of the Ethernet group of 802.3 on broadband Ethernet, and so I had some experience on politically what's the best way to achieve what you might want to achieve within the IEEE committees. I knew it really works out well if you can get multiple companies behind a single proposal and put that in front of the group, and then start to get momentum. I think that's pretty much the standard way that this happens these days within IEEE. So we recognized that working together between Symbol and NCR would be a good thing to put together a joint proposal. Then, and this is almost a personal thing, there was a third company that joined us, and that company was Xircom. The reason why I say this is personal is that at Sytek, my boss, a guy named Ken Biba, he was now the-- I don't know if he was the CEO, or what his position was at Xircom, but he was at Xircom, and they were starting to develop wireless LAN technology as well. So I knew Ken from years. He had a guy working for him named Phil Belanger, and Phil I had previously met many years ago when he was at Corvus just doing local area network development. So, ultimately, Phil, and Wim, and I, put together a joint proposal for really the fundamental architecture of what we felt the 802.11 standard should be for the MAC, the Medium Access

Control layer, which supports multiple different physical layers, but uniformly there is this one medium access control layer that is kind of the top-level architecture of everything.

Cohen: So to recap, it was Xircom, Symbol and NCR joining together.

Ennis: That's right, and actually the proposal that we put together it was called DFWMAC, the acronym. We said this stood for Distributed Foundation Wireless Medium Access Control. Distributed Foundation Wireless Medium Access Control, DFWMAC, but we came up with the acronym after, I mean, we came up with what that stood for after we came up with the acronym. Because at one of the IEEE committees one of the chairs of one of the subgroups, when we ended one of the meetings, he said his fondest hope is that the proponents of these different proposals would happen to meet on their way home at DFW Airport, and they would get together over a beer and put together a joint proposal that they could bring in. So we decided we've got to call this protocol DFWMAC. Then after we came up with that we decided, okay, so what can that really mean. Okay, it's a Distributed Foundation Wireless Medium Access Control. The DFWMAC proposal ultimately won out over the IBM proposal. There was a proposal from National Semiconductor. There were other competing proposals that had been put forward. In November of 1993 our proposal ended up getting adopted by the 802.11 committee as the foundation for the standard, and then the real development of the specification took off from there ultimately leading to the publishing of the standard in 1997.

Cohen: How come it took four years to go from the basic agreed framework of the MAC layer, the Medium Access Control layer to the publication of the standard, and then how involved were you at that time?

Ennis: So as soon as the DFWMAC proposal was adopted, and now we all had a clear course in front of us for how to develop this, I became the technical editor of the 802.11 committee. So I was very involved in the actual writing of that text that ultimately became the specification. But there were a lot of details that needed to be worked out. There was also, still, a tremendous amount of work happening at the physical level. At the time there were three different physical levels. There was a radio based on direct sequence technology. There was a different radio approach based on frequency hopping technology, and then there was an infrared approach. Those were the three kind of different modes of operation within the original IEEE802.11 standard. Ultimately it was just the direct sequence one that ended up taking off commercially and becoming what we now know as Wi-Fi. But so there was a lot of work going on at that level, as well, to finalize the standard.

Cohen: So three different physical layers, but one MAC layer, one Medium Access Control layer?

Ennis: One Medium Access Control part.

Cohen: You were editing both?

Ennis: I was editing the MAC. Well, no. Actually, that's not true. I was editing the whole document. There were subeditors who were contributing different sections, but yeah. No, I was working on the entire document. There's another. Bob O'Hara was also. He and I were coeditors on that.

Cohen: You mentioned, of course, rightly so, that direct sequence ultimately became the PHY layer that became a big part of what we know today as Wi-Fi, but weren't there a lot of frequency hopping implementations at the time as well?

Ennis: There were, and it's still a very viable technology. It's the basic technology in Bluetooth, for example, although that's separate from the 802.11 stuff. But I think what ultimately happened was once the initial standard came out, and the data rate at the time in the initial standard was two megabits a second. Perfect for handheld warehouse and inventory control applications, not so good for office applications, so people saw that if we can get the speed up to where we can support data rates that similar to Ethernet, 10 megabits per second, of course Ethernet at the time you had 100 megabit per second capability coming out, but still people viewed the 10 megabit per second threshold as what was required to really support office applications with wireless. So there were a number of design approaches that people were coming up with to try to reach that speed threshold. Ultimately it turned out that it was the direct sequence technology rather than the frequency hopping technology that lended itself better to evolving towards higher and higher data rates, and I think that's ultimately what made the direct sequence technology the one that ultimately evolved into Wi-Fi.

Cohen: Right, but at the time of the 1997 standard you had both frequency hopping and direct sequence commercialized products for barcode scanning and other applications both at two megabits per second.

Ennis: Yes. That's right.

Cohen: Right, and then the standards committee wanted to look at, as you had said, increasing that speed to something that was a similar speed to ten megabit Ethernet; that's when the focus became more on the direct sequence approach.

Ennis: Right. That's right. So that ended up being 802.11b, which was a new specification for a new physical layer within 802.11, same Medium Access protocols, same MAC, but now at an 11-megabit rate. Before the 802.11b approach was finalized within IEEE, similar to the earlier efforts there were competing proposals that people put forward. Actually, I was consulting for 3Com at the time, and working on some projects for 3Com to do exactly that - to achieve the 10 megabit per second rate and to try to address the personal computer market, the office application market and hit that Ethernet speed.

Weber: Before 1997 can you just give a brief thumbnail of who was actually using wireless networking at the time, and in what industries?

Ennis: Before 1997...

Weber: Or around '97.

Ennis: Yeah, it was very much a vertical market. It was these warehouses. For example, one of the early applications of this technology was for the car rental companies. When you go up to a car rental company and they have their remote wireless terminal where they're checking you in. That stuff got started to be introduced back in those timeframes, and it was originally proprietary networks. There was a company called Telxon that was developing, a competitor to Symbol and to NCR, that was developing products like that. Things like UPS. I don't know specifically that UPS, I mean, they certainly not much later than that started to have systems like this. So it's for those kinds of what we'd call vertical market applications in the wireless industry, as opposed to the broad horizontal market that everybody ultimately wanted to go after, which was putting it in the laptops and into PC cards that would plug into laptops, and putting it into peoples' homes and offices for standard office type applications. Remember, this is really before the Internet was big. That's another thing to keep in mind. We're talking about the early '90s. We were just starting to see the beginnings of browser technology and web. It was the Internet. You could do e-mail. You could do file transfers, but we weren't doing web browsing. It wasn't the web. That technology was being developed, really, at the same time that we're talking about this wireless technology being developed. So those types of applications, the horizontal applications for Internet access for office suff, really wasn't going on in the wireless industry back then.

Weber: In '97 that was still the case?

Ennis: In '97 that was still the case, that's right.

Cohen: So you were doing some consulting at 3Com, which I know, because that's when I met you when I was working at 3Com. So 3Com's, I guess, consulting direction to you was to do exactly what regarding the standard?

Ennis: So 3Com, at the time, they were the dominant local area network manufacturer. Very successful selling plug-in cards for PCs, IBM PCs, and so they saw wireless as an obvious next step for their traditional markets. I mean, they were definitely into that horizontal market for their wired Ethernet local area network technology, and they wanted to get into that for wireless. 3Com had an internal project to try to develop a wireless technology that would cross that magic ten megabit per second threshold. I was initially brought in by a person that I had known previously, Jeff Abramowitz, who was at 3Com, to help evaluate this internal technology, and to determine if that's something that we potentially could put forward to IEEE as a proposal for a higher data rate standard that could be adopted by IEEE. So that was my initial task, and as a result of that project I think we all came to the conclusion that there are other approaches that people are talking about out there, and we need to look at those. Because at the same time, IEEE802.11, the wireless LAN committee, they had started up a project called 802.11b to come up with a higher data rate version of the standard. There were proposals, most notably, coming in from a company called Harris, which was a major semiconductor manufacturer producing semiconductors for wireless local area networks, and then another proposal coming in from Lucent, which was really NCR. Lucent had bought the division of NCR in Holland that we were talking about earlier in the 1993 timeframe. So there were these other proposals for 10 megabit per second speeds in wireless LANs.

Cohen: Was that the agreed upon goal in the 11b committee? Was it a formal or informal goal to reach a 10-megabit barrier?

Ennis: I think it was a formal goal. Within IEEE there is a project authorization request that a particular effort needs to have to get authorized to do work. I'd have to go back and check that, but this 802.11b was called the high rate standard, and I suspect the 10 megabits was built right into that project authorization request.

Cohen: I guess formally or informally everyone acknowledged that that's the goal to reach.

Ennis: That was definitely the goal. So my work at 3Com changed to really evaluating some of these other proposals that were being put forward within the IEEE committee to try to see if partnering with one of those companies would make sense for 3Com. So as part of that I started to represent 3Com on the IEEE802.11 committee, and to start to work with these other companies on seeing what we could do to come up with a standard there.

Cohen: You mentioned here are basically two proposals. I guess the 3Com specific proposal was abandoned at some point, and there were two major proposals, one from Harris and one from what was then Lucent after NCR was acquired. Can you describe, roughly, the two different approaches, and what the differences were?

Ennis: So it was differences in the modulation technique that was being looked at, and the encoding technique for the physical layer of the radios. Both approaches did not require any change to the MAC protocol, and so it was really focused on how the bits really are signaled over the airwaves. It almost got into religious wars. There were other proposals as well. It wasn't just those two, but those were the two big ones. Ultimately it was the Harris technology[1] that, with some modifications as a result of all of the discussions, but ultimately it was the Harris technology that got adopted within IEEE for the 802.11b.

[1] [Interviewee's note] In fact it was a combined proposal from Harris and Lucent that got adopted for 802.11b.

Cohen: Can you tell us about some of these passionate religious wars and some of the fun stories that may have ensued during these standards battles of, "No, it should be my proposal, no, my proposal." Talk about both maybe some of the drama that played out as well as the technical arguments that were going on to the extent that you can remember.

Ennis: I'm not sure I can necessarily remember the exact detail here, but there was a very dramatic moment in one of the IEEE meetings working towards the selection of an approach for 802.11b where on one of the critical votes there ended up being accusations of block voting. So IEEE, the way IEEE operates is it's an organization of individual engineers. Obviously everybody who's there they're working for their individual companies, and in practice people bring in their own biases based upon on the companies that they work for, but formally the IEEE is a committee of individual engineers, and organizing all of your employees to vote a particular way within IEEE that's a real no-no. Evidence of that, if it can be brought forward, is a severe problem. There were accusations on one particular vote of block voting, and I remember that created quite a disruption in the process.

Cohen: Was one proposal about to win, or was it just an accusation that was leveled that had to be dealt with?

Ennis: I don't recall the specifics there. We could go back and look at the minutes.

Cohen: The meeting minutes, yeah.

Ennis: This would have been in, probably, in early 1999 I would suspect. No, probably 1998, I would think.

Cohen: You mentioned the Harris proposal eventually won out. Do you recall any sense for why ultimately that one got enough favor, because for everyone to understand it has to get a certain amount of votes from all the individual members, so at some point some votes had to be swayed, I suppose, if it were previously deadlocked.

Ennis: I don't remember the details of that. I do recall that ultimately in my work at 3Com, working with the people at 3Com, that we decided that it was the Harris proposal that seemed to be the best technical fit for what 3Com wanted to do. So I know that there was discussions that we had as part of my effort there, but I don't know what other people-- but ultimately the Harris approach did end up getting selected.

Cohen: Well, what happened? From your recollection what happened next? Everyone voted for the Harris proposal, and then, when was the 11b amendment officially published?

Ennis: The 11b amendment was officially published in 1999, and so there was an effort that had started up that was a competing effort to the IEEE efforts that was called HomeRF. HomeRF, there were a number of companies involved in that. Intel was in there. Proxim was very influential within HomeRF, and their goal was to develop some kind of a standard multi-vendor radio technology that could be used not just to support local area network applications, but also to support voice applications. So their concept was that this was going to be kind of a universal network, wireless network within the home that would support cordless phones, and also file transfers. There was a particular technology that they were pursuing for that, and that was different from IEEE. It was not being put forward within IEEE as a potential standard. It was really potentially a threat to all of the companies who had been focused on having this true international standard come out of the IEEE efforts. So I was also consulting for Symbol Technologies at the time, and with Symbol Technologies I was asked to participate in the HomeRF efforts. There were separate meetings going on to develop things. Of course, I was still very actively involved in IEEE. My task was to try to convince HomeRF that for the higher speed version of what they were doing that they should adopt 802.11b. I gave a proposal. I gave a presentation at a HomeRF meeting on this, and was essentially told, "No. We're too far along. No. That won't work, and besides the IEEE802.11 stuff is way too complicated, and it's too costly, and it'll never fly on the market," and all these types of things. There were, at that meeting, other representatives from Symbol. There was a representative from Harris. There was a representative from 3Com. As I recall, we ended up going out and having Sushi in Vancouver during this HomeRF meeting, all of us.

Cohen: I think it's okay to say who it is.

Ennis: So it was David Cohen from 3Com, and Jim Zyren from Harris, and then I believe that Sarosh Vesuna was probably there from Symbol, and me. It might have been only three of us.

Cohen: I think it was three of us at the sushi if I recall that.

Ennis: So we were all proponents of IEEE802.11b at this point, and we decided we can do what HomeRF is trying to do here. We need to form an organization that will certify interoperability of 802.11b products, and so that we can prove to everybody that this stuff can be implemented in a multivendor way that everybody can build, and that was the genesis, at least my recollection, of what ultimately became the Wi-Fi Alliance.

Cohen: So, Greg, we've talked about some of the standards bodies, and some of the players, Harris, Symbol, Proxim in another camp and all that; NCR, Lucent, all that; talk a little bit about the complexities involved in doing a wireless local area network for the 802.11 effort because this is certainly a challenge beyond wired local area networking.

Ennis: Absolutely, and in fact there were many naysayers as this effort was progressing, saying that it's just not going to be possible to really come up with a standard here that people will be able to implement and that they'll be able to interoperate with each other, because it's just too complicated. The 802.11 spec, just the original 1997 802.11 spec ended up being that thick, whereas the original Ethernet specification was like that. Even though we used to call the 802.11 stuff wireless Ethernet, well, why is it so much thicker and more complicated? Well, that's because it's wireless, so you have issues. I mean, for one thing, just the communication medium it's much more difficult to communicate a bit from Point A to Point B when you've got all sorts of interference going on, when you've got competing users. Remember, this technology is operating in unlicensed bands, and that was one of the real...

Cohen: What does that mean?

Ennis: So that was one of the real instigating factors that kind of leads to the proliferation of this technology originally. So the FCC in the United States, and other regulatory authorizes in other parts of the world, they're responsible for coming up with the basic rules for how equipment can transmit over the airwaves using radios. Some of the bands that the FCC regulates are called licensed bands, which means that you actually have to pay money to get a license for a particular set of frequencies within a particular geographic area, and nobody else can use that frequency in that area.

Cohen: Kind of like a TV station, for example.

Ennis: Like a TV station. Like cell phones, I mean, cell phones today, the different cellular operators they have licenses for specific frequencies in specific geographic regions and it's only their equipment that can be using those frequencies. So in that type of an environment now you're in pretty good control. It's still a challenging environment as far as just transmitting over the airwaves as opposed to having a nice clean wire that you're transmitting over, but at least you're not putting up with interferers that are competing in the same frequencies. So in contrast, the wireless LAN technology was authorized by the FCC, and the same thing in other parts of the world, to operate in unlicensed bands, which means that a user doesn't have to go out and get a license from the FCC. That's why today you can go out and buy a laptop or a Smartphone and it's got Wi-Fi in there, and there's nothing that you needed to do in order to obtain that license, or the provider of that Smartphone, actually. The complication is that that means there can be interferers in there. There can be other users of those same frequencies in the same geographic region who are competing for that same bandwidth. So the protocol, and this really comes down to some of the complications that go into the Medium Access protocol, which was the original proposal we did back in '93, you have to do things in order to operate effectively even in the presence of interfering users. So that means that you have to have the built-in acknowledgements. Ethernet doesn't even bother with built-in acknowledgements. It relies on higher-level protocols to handle that level of reliability because the wire itself is just so reliable.

Cohen: Yeah, so essentially you have two challenges. One is you don't have a wire, so the medium is a little less certain, and also, like you mentioned, you don't have license control over that particular spectrum. You just have to deal with an unknown number of other devices in that spectrum.

Ennis: Yes. That's right.

Cohen: So talk a little bit about how the standard dealt with that challenge.

Ennis: First, the underlying physical radios would be designed to be able to operate in that environment, but at the Medium Access level the protocol would be designed to, well, this is similar to Ethernet, although it operates differently, to sense whether there's another transmission going on, and to not transmit if it's sensing that. But then, also, there's built into the protocol for this stuff, as it ultimately became standardized, there's something called collision avoidance. I'd say Wim Diepstraten from NCR he was one of the main developers of some of these concepts. Before you start transmitting there's actually a mechanism built into it so that if one of your peers might want to be transmitting at the same time you would kind of randomize your start of transmission. So that was built into the protocol, but they also built in acknowledgements, which at the time in the 802 family of standards, having acknowledgements built right into the 802 standard was typically thought of as not necessary.

Cohen: What does an acknowledgment actually do?

Ennis: So it makes sure that a transmitter gets a positive feedback that the receiver actually received that packet. So that's built in there, but then in addition there's other vagaries of the wireless environment that make things complicated. So, for example, you could have a station over there trying to communicate with somebody over here, and another station over here that's also trying to communicate, and these two stations are so far away from each other, that even though they're trying to communicate with the same guy, they can't hear each other. So there's complications in the protocol to be able to deal with making sure that these transmitters don't step on each other even though they can't hear each other. That involves something that actually there was the Xircom contribution into the joint proposal that we ended up putting out dealt with that. That's the so-called hidden terminal problem, but then in addition there's roaming issues. I mean, in a wired environment a device doesn't typically roam from one location to another, whereas in a wireless environment, I mean, that's the whole point. In fact, you're connected to one access point that's giving you access to the rest of the network, and you can roam and connect to another one without losing your connectivity. So all of those mechanisms in order to deal with roaming, I mean, that's another complication of a wireless environment that you don't have within the wired environment, so that's why that document became so thick, because it had to deal with all these issues of the wireless environment. Security's another reason, right? Security is way more of an issue within a wireless environment than it is within a wired network. So there's security mechanisms that get built right into the standard that weren't viewed as necessary for Ethernet, for example. So we ended up with a very thick specification, and because of that people were skeptical that we would be able to have independently developed implementations really interoperate just because it's such a complicated specification. People figured, "Ah, it'll never work. People are going to interpret the spec differently." That's why ultimately we started the Wi-Fi Alliance to make sure that everything interoperated properly.

Cohen: We'll talk more about that soon. Is it possible to talk, without getting really lost in the technical details, to talk about how the standard dealt with some of those issues like, for example, the hidden terminal problem? What was it in the standard that helped resolve that? Is it possible to get kind of a high-level explanation of that resolution?

Ennis: Yeah. In the original standard, I mean, basically the MAC protocol, although it's had some evolution, today we're basically using the exact same MAC protocol that we came out with in the mid '90s.

Cohen: The existing, yeah, the '97 802.11 MAC.

Ennis: There's an optional packet exchange that stations can do. It's called an RTS and a CTS packet.

Cohen: That stands for?

Ennis: Well, it comes, actually, from the old RS232 days that you'd plug-- now we're going way back, RS232 is the wired connector standard for connecting a dumb terminal to a local area network back in the '80s, or connecting a terminal to a modem. There's multiple signals within those wires and one of them is called RTS. It stands for Request to Send, and the other one is called CTS. That's Clear to Send. Part of the RS232 mechanism there's a little handshake that goes on so that the transmitting side doesn't send something that the receiving side isn't ready to receive. So this was really a play on words to call these packets RTS CTS, because it was a little handshake. But basically, the way that it would work, this was for communication, for example, to an access point, and you'd have one station that's over here, and another station that's way over here, and they don't hear each other, but the access point hears both of them. So basically the way this would work would be if a station had a packet to send it would send an RTS packet first, very short little packet, to the access point, and within that packet would be an indication of how big the packet was that he wanted to eventually send. The access point would retransmit that out as what's called a CTS, still including that information in there saying how big the packet is that this guy wants to transmit. Since that's being transmitted from the access point that would be heard by everybody, even if they can't hear themselves, so this station over here, it may want to transmit something, but it first sees a CTS coming in from the access point that indicates to it that there's some other station that wants to transmit a packet this big. So what it's going to do, it's going to wait for that period of time until it knows that that other packet got transmitted. I simplified it a bit, but that's basically the idea.

Cohen: It makes a lot of sense that the second station can hear the CTS packet, but it didn't know originally that the other station existed.

Ennis: That's right.

Cohen: What about roaming? Part of the whole point in wireless is roaming, and if you have to be stationary it kind of defeated the purpose. You may as well use a wired network, so can you similarly give maybe just a high-level overview of how the roaming handoff worked, or was it borrowed from another concept? How did that get developed in the original 802.11 MAC?

Ennis: Kind of the standard configuration for Wi-Fi these days, and this is true way back in the earlier days, was where you have some kind of a wired network already installed within an office environment, or factory floor, or whatever, and what you're trying to do is you're trying to add these wireless devices to that environment, to connect them into the wired environment. The way that you would do that would be to have these special devices called access points that would connect to the wired network directly like through Ethernet. I mean, you still have these in your homes. You have access points that have an Ethernet connection to your DSL modem, but these attach to the wired network, and then they serve as the wireless access point, that's why they're called access points, for all the wireless devices to connect to. But in a large facility like a factory or a big office complex you would have multiple of these access points. The radios can only transmit, say, a hundred meters. It depends on the thickness of the walls and all of that. So you'd have multiple access points, and people would potentially be roaming. So built into the protocol, and this is another big difference between what we had to do for the wireless standards versus what happens in, say, Ethernet, built into the protocol is that notion of a station associating with a particular access point and then using that access point while it's within range of that access point. But if you roam to where you're starting to get a better signal from another access point it would be better for you to switch your association from this access point here to that one, because you'll have better performance. So not only is there this handshake to actually achieve the association, but there's mechanisms built into the protocol to be doing background scanning of the environment so you can sense, "Oh, I'm currently connected to this access point, but guess what? On another frequency channel in my background scan, that I did when I wasn't transmitting any data here, I've detected that there's another access point that actually has a stronger signal, and so I'm going to reassociate now with that access point." So, all that mechanism, the scanning, the determination of what the best access point is, the actual association with the access point, that also impacted the thickness of that ultimate specification.

Weber: Was there much cross-fertilization with the telecom world like GSM and CDMA? I mean some of the issues you're talking about have, certainly, analogs there.

Ennis: Right. So there was, and certainly some of these ideas aren't unique to the wireless environment for the wireless LANs. So there certainly was cross-fertilization. I would say in those days, though, the people participating there was the telecom world, and remember, even in those days cells phones weren't as prevalent as they are now. So there was, obviously, work going on in that world. There was work going on in the local area network world that really came out of the LAN efforts that didn't have necessarily strong ties with the telecom world. So there was some separation, but everybody involved in this, everybody's engineers, you all keep up on what's going on in all sorts of aspects of technology. So there certainly was cross-fertilization, and today there's tremendous cross-fertilization because Wi-Fi is incorporated in all the smartphones, and the big wireless operators are all also the big hotspot operators. There's a tremendous cross-fertilization, but back then there was some separation.

Cohen: So one last, I guess, quasi-technical question. There's many chapters to this, but you mentioned that compared to a wired network in wireless you have to be concerned about security; just if you could briefly explain. It feels right, but explain a little bit about why that's so on a wireless network, and what were some of the early efforts to address security.

Ennis: If you think about a wired environment in an office, let's say Ethernet, where you've got all your computer terminals back then, then later your PCs, all your servers and everything. Of course, today, it would all be Wi-Fi, but let's take a wired environment with Ethernet. So with that type of an installation, the only way somebody could tap into the network to obtain the data, and to snoop, and to compromise the privacy and the security of your data going on, on that network, the only way they'd be able to do that would be to actually enter the building and to put some kind of a snooping device on one of the actual connections to your wired network. Typically, an office installation you would prevent that by not letting just a random person come into your office. You would have locked doors, and you would assume that if the person can't get into the facility they're not going to be able to compromise my data. So consequently, back in the '90s, a wired network would typically not be encrypted at all. It wouldn't have any special security. You'd have passwords on a file server, or whatever, but as far encrypting the actual data, that typically wasn't viewed as a real strong requirement. Wireless goes through walls. Somebody can park a truck out in the parking lot and conceivably could snoop on your traffic without having access to your physical security. So that's why the base standard of 802.11 incorporated encryption. It was actually called Wired Equivalent Privacy, WEP, and the meaning of that is that this should provide the same level of privacy protection as you get from an unencrypted wired network that is protected by physical security mechanisms like locks on the doors and things like that. So, that was the initial encryption mechanism that was put into the 802.11 standard right from the start.

Cohen: Yeah, and all that, as we said earlier, got incorporated in 802.11b also where the MAC layer's still the same.

Ennis: Right, and went through significant evolution.

Cohen: Right, which we'll talk about probably a little bit later.

[SESSION BREAK]

Cohen: Great, we're back again. David Cohen here, talking to Greg Ennis. And we left off, we were talking about, how the standard got developed, and then we realized that maybe we needed something more than just the standards body. The standard was great, a lot of great intelligence went into it. But we realized maybe we needed something more to promote the standard and, you know, get some promotion behind the standard, ensure interoperability and all that. And I think there were a couple of events that led up to it, and we were just talking about one of those. At one point, I think, you know, you, as a consultant for Symbol, me, as an employee for 3Com, Jim Zyren as an employee for Harris Semiconductor, went down to the HomeRF meeting, the other standard that was bubbling around, to see if we could merge efforts and see if we can get them to adopt their Version Two, as IEEE 802.11b, so we have a nice kind of merger and harmony. So that was our mission anyway. What happened at the meeting from what you recall, Greg?

Ennis: We were shot down.

Cohen: Succinctly put.

Ennis: Yeah, and I don't think it was for technical reasons. Yeah, so I gave a presentation, I recall, that went through what the IEEE 802.11b high rate standard was all about and trying to convince them that this made sense for HomeRF, as their next generation standard as well, and like you said, so everything would converge. But I think they were already pretty far along on developing a specification for their next generation stuff, and they were expressing, I think the typical concerns that many people had, which was, you know, the IEEE specification is way too complicated, it's not going to end up allowing products to really interoperate, because it's going to be so difficult to implement. It's going to be really costly and it's just not going to work, and we need something that's more closely focused on our specific requirements, which included this voice support requirement. And they basically said, we're going to continue on our path, and you guys can do what you want.

Cohen: There's the door. Now remember, some of our-- I remember you went over very detailed technical presentation. I thought there might be like a counter technical presentation. There wasn't. It was like, thanks for coming, you know.

Ennis: Right.

Cohen: And you know, and to your point, they had some legitimate concerns, I think, in terms of at least the perceived complexity, the fact that they wanted to make it voice oriented, and IEEE 802.11b at the time, wasn't really voice oriented. Those quality of service enhancements came later. So I think it was a good effort, but you know, they didn't really want to adopt our proposal.

Ennis: And I think they did have legitimate concerns about the ability to-- for multiple independent implementations to be done, and to interoperate. And so that was really the fundamental problem I think we needed to solve at that point, in order for this standard to be successful.

Cohen: Yeah, I think they were right about that, and for HomeRF, they were basically, you know, essentially taking one company's – Proxim's technology, and so they didn't have to worry about interoperability as much. But then again, it wasn't an open standard, necessarily, like the IEEE standard, so yeah, I think that after that meeting, we realized that there was going to be no intersection. I think that was, when we sat down over sushi, you, Jim [Zyren of Harris Semiconductor] and I and said, okay, we just got shot down, now what? Maybe we should do something a little bit different, you know, and I think that was kind of like the last straw in trying to merge, right?

Ennis: Well I think we realized that what HomeRF was doing, was actually good in many ways. It was an organization of companies, rather than a bunch of engineers just getting together, and those companies had commercial interests, and they wanted to see it succeed, and they wanted to put resources behind something in order to get it succeed, and I think we realized, you know, some kind of a sister organization with the IEEE Standards Committee that focuses on those aspects of this technology that IEEE didn't deal with, like interoperability testing, really made a lot of sense.

Cohen: Yeah, I think that was the genesis for forming what would later become the Wi-Fi Alliance, was, yeah, we had two charters, right, one is let's establish multi-vendor interoperability, right, we're just-- that's, if anything, the reason for being, right. It's, like, as people may or may not know, just because you build to the standard, it doesn't necessarily guarantee interoperability. It might seem like it does, but if Company A says I build my very complex, thick, standard based product on IEEE 802.11b, and another vendor says, well I built my product to IEEE 802.11b, and then they ship them both to a customer, and it doesn't work together, now whose fault is it, right? You get into that problem, a lot of finger pointing. And so solving true compatibility and interoperability is a really important goal.

Ennis: And so I think you know, the way that the Wi-Fi Alliance then started here, I guess the first thing we did was come up with a name, it was going to be the Wireless Ethernet Compatibility Alliance, because we didn't have the name Wi-Fi yet. And we were trying to come up with-- we were viewing 802.11b running at 11 megabits as the wireless equivalent of Ethernet, so that's the Wireless Ethernet Compatibility Alliance, but actually you might recall all this, that 3Com, had a project to develop a product line around that, and they contracted much of the development out to Symbol Technologies, and the chip that the Symbol Technologies implementation was going to be actually based on, was the Harris Semiconductor chip, and so-- and I was consulting both for 3Com and for Symbol at the time.

Cohen: Right, that's interesting.

Ennis: I think there was an agreement between you guys, so--

Cohen: There was, there was.

Ennis: And so it really ended up initially being 3Com, Symbol and Harris, getting together. That's, you know, like at the sushi meal, you me and Jim, talking about forming this organization to do interoperability certifications, and that was the start of WECA, and subsequently, we ended up inviting--

Cohen: Aironet

Ennis: Aironet which later got bought by Cisco, and Lucent, which was the successor organization to the NCR guys who had done all that early wok that I talked about before, and then Nokia--

Cohen: Yeah, Nokia came in kind of at the end, right before we launched, I think Nokia joined in as a customer of Harris. You know, I think it was--

Ennis: Yeah, well Nokia had a development group in Cambridge, England. Just a small little development group doing local area network development, and they had been active in the IEEE, standards. So we knew those-- I mean, we didn't know people at the, you know, at the Helsinki headquarters, but we knew the Cambridge England guys, and we thought that that would be really great to get Nokia in, and it was. They've been--

Cohen: Yeah, that was the genesis for forming the Wireless Ethernet Compatibility Alliance, and we didn't have the name Wi-Fi yet. That was actually a project of the Wireless Ethernet Compatibility Alliance later, to come up with a name, and then I think that was part of it, too, like you said, Greg, it wasn't meant to emulate the IEEE as a group of individual engineers working on a technical standard, it was much more about multivendor interoperability, it was absolutely about commercial interests, whereas, like you said earlier, when there was a voting block or something, it was all supposed to be individual engineers, that was no-no IEEE, and WECA, and now these are companies with commercial interests, they want to commercialize their product, and the joint belief, even though I guess some of these companies that we mentioned are competitors, is that the market would be bigger if they were all interoperable, so let's fight over a larger piece of pie rather than a smaller piece of pie. And I think it was definitely a commercial marketing and business interest, as you know, it was driving a lot of it to try to get the technology to go from these vertical applications that you mentioned earlier, to something that is much more mainstream. And then we all had a strong belief that interoperability was going to be key to achieving that.

Ennis: So, but the first thing we needed to do was come up with a name, because the standard at the time was called 802.11b High Rate.

Cohen: Right, IEEE, 802.11b High Rate.

Ennis: Right, and that wasn't going to fly in the market, so--

Cohen: Yeah, I remember these discussions, yeah.

Weber: Could you talk-- sorry-- What year-- so I know '99 was the sushi meal?

Ennis: Yeah, this is '99.

Weber: So this is all within a few weeks or months.

Cohen: Yeah, this is all '99. A few months, I'd say, and then there was another aspect, too, that it be a fair dimension as the history, was the sushi meal where we realized that we were never going to merge with HomeRF, and then either shortly before or after that, that I also had a conversation my boss at the time, at 3Com, Jeff Abramowitz, who was also instrumental in this whole history, just a conversation saying, yeah, we have to form an organization and we need to reach out to some of these other key players and do that, and he kind of put me in charge of being the 3Com guy, you know.

Ennis: So you were an original board member, you know.

Cohen: that's right.

Ennis: And I, you know, as a consultant, and not working for any of the board companies, you know, it was figured that, okay, the way we'll have Greg fit in here is we'll make him technical director. And so I was kind of like the only-- I wasn't an employee-- well but I was not a board member, but in charge of the technical work leading up to the actual interoperability certifications, and all that.

Cohen: Yeah, and that was a huge part of it, is we didn't want to be IEEE, at the same time, you know, we didn't want to be just another PR announcement and marketing announcement that like, hey, we're here, we're doing wireless LAN. I think there was also-- I think there was a nice synergy, even from the very beginning, that we have some technical meat, right, we're really solving a technical problem. We're not writing a standard, but we're adopting the standard and solving interoperability, while, at the same time, we're also, you know, a voice to promote it, you know, wireless LAN, 802.11 has a home, it has a voice, you can talk to somebody about it if you're the press or you're an analyst or something like that, and you want to learn more about it. It now has a home, rather than, go download this, join IEEE and go download a spec, you know. And I think promoting it was really key. So yeah, this all formed in-- all this work was going on in early 1999. I think I remember all that leading up to our public announcement.

Ennis: So do you remember some of the alternate names that were--

Cohen: You know, I get asked that question a lot, and I feel really bad that I do a bad job of answering it. So I'm going to hope that you're going to help me remember.

Ennis: Oh, I'm sure--

Cohen: Because I can remember presenting [the name] Wi-Fi to the board, and thinking that it wouldn't fly, but it did, and I was very happy, but I can't remember the alternatives-- all the alternatives.

Ennis: So, the Wi-Fi Alliance, several years ago, had its ten year anniversary, so it must have been in 2009, their ten year anniversary. And so I actually gave a presentation, I'll have to show you this.

Cohen: Oh, please.

Ennis: That on the ten year anniversary. And what I did was I did the whole presentation at the beginning, pretending that we hadn't selected the name Wi-Fi, but instead we had selected the other names, and so I put up these logos with the standard Wi-Fi logo but with these other names up there, rather than Wi-Fi. And of course most of the people in the audience, you know, they were very confused about what was going on here. But-- so I can name, you know--

Cohen: Please.

Ennis: Dragon Fly?

Cohen: Oh yes, Dragon Fly, I remember that.

Ennis: Torch Light?

Cohen: I do remember that.

Ennis: Kinect, K-I-N-E-C-T. Elevate.

Cohen: Elevate.

Ennis: So remember, it was 802.11, at 11 megabits per second, so hence, Elevate.

Cohen: That's a little too cute.

Ennis: Right.

Cohen: Well I remember one board member had a suggestion, maybe I should or shouldn't mention his name, but I think I will. But somebody said we should call it Flank Speed.

Ennis: Flank Speed, that's right.

Cohen: That one I remember. And then someone said, that sounds like flank steak.

Ennis: Yeah, that one--

Cohen: That was Jim.

Ennis: That one didn't fly.

Cohen: That one didn't fly.

Ennis: Well none of them flew except for Wi-Fi.

Cohen: That's right.

Ennis: And so Wi-Fi, of course, it was a play on words of Hi-Fi. Of course, nobody talks about Hi-Fi now, but back in '99, Hi-Fi was still a term that, you know, high fidelity, and so you know, we figured we needed to say wireless fidelity, although it really meant nothing. And we very quickly abandoned the wireless fidelity and just went with the Wi-Fi.

Cohen: Right, my recollection is, it was kind of a, it went either way. You know, first of all, Wi-Fi, we thought, as a play off of Hi-Fi, would be catchy, easy to remember, and be much easier to remember, of course, than I-- again, you know, even knowing that it's pronounced IEEE 802.11, you have to be on the inside to-- to the outsider, it would be the symbols of IEEE, 802.11, just an absolute mess of a name. And that Wi-Fi, even if you didn't know what it meant, it certainly would be catchier. And then you might be able to guess that it had something to do with wireless, because it had Wi, and it would just be catchy, and then if they really knew that it stood for Wireless Fidelity, hey, all the better, you know, but it wasn't really-- to your point, dependent on that, and later that was kind of dropped, because originally, remember, the tag line was, Wi-Fi, The Standard for Wireless Fidelity. I still have shirts that have that. Then we dropped that later on. It wasn't necessary.

Ennis: So when we got the proposed names put in front of us, by this branding company that we had hired to help us through this process, and I remember that they were all presented, and then we were all to go home and think about it, before we came back to make a decision. And so I remember, I went home and my daughters were young then. I think they were probably I don't know, 5 and 11, I'd have to figure out how old they were. But I remember sitting down with my daughters, and just saying, so what do you think of these names? And you know, Dragon Fly, and I remember they liked Wi-Fi, they didn't know why, but they liked Wi-Fi. I think we all reacted to it that way.

Cohen: That's right.

Ennis: But, so I like to say that my two daughters had a hand in that name.

Cohen: In picking the name, that's great. And I remember well, I remember, yeah, the name company gave us a big list, and I think I was, if I recall correctly, you know, responsible for presenting it to the board and I really-- for some reason, I thought that I liked Wi-Fi, but I figured it just wouldn't fly. You know, I thought it might be too esoteric, but everybody gravitated to it right away, you know, after we went home and thought about it a little bit. I think that was an easy-- much easier conversation to settle on the name, than you might have thought, you know.

Ennis: So the board at that time, I mean, it was-- I mean, you were on the board, and then Jim Zyren, who we've talked about.

Weber: Oh, sorry, give a little background on him. He was sort of the third--

Cohen: Sure. Jim, yeah, Jim was, Greg can add, but he was-- worked for Harris Semiconductor, which later spun off as Intersil, in wireless LAN semiconductors. So he is an IEEE member, a very technical guy, a lot of radio design background, and you know, was their standards body guy also, that went to a lot of the IEEE meetings, and was trying to work out the 802.11 standards, and you can give a little more to Jim, to cover it.

Ennis: Well, I still see him occasionally. He's-- his company that he ended up working for, got bought by Qualcomm, and--

Cohen: That's right. Yeah, and I talk to Jim occasionally, too, and I have to give credit to Jim, because when I was brand new at 3Com, without any wireless background, it was actually at one of the earlier HomeRF meetings, not the one that we finally decided they weren't going to do anything with us, and we went to sushi, but an earlier one, Jim and I were talking, and I was still brand new to 3Com, and he did an amazing job in just explaining some basics of wireless. And I remember this conversation, because it was a really weird setting, because the HomeRF meeting itself, at that time, it was like adjacent to this big horse barn. And so we were having this conversation about the basics of radio, literally sitting in this barn, you know, it was like a break in the meeting, when he was explaining some very basics of wireless. I still owe him to this day for kind of patiently explaining some of the very basics of radio modulation to a newbie in the industry. So-- he gets a lot of credit.

Ennis: And then also on the board, was Phil Belanger, who you know, we talked about earlier, because he, along with Wim Diepstraten and myself, we were the trio who put together the proposal back in 1993, the DFWMAC proposal, right. So you know, I had been working with Phil all those years, and he had since gone off to work for Aironet, which was a spin off from Telxon, actually. But they-- right when we were forming the Wi-Fi Alliance, Aironet got bought by Cisco, so some of our early materials, right at that time, talk about Aironet, and then suddenly it starts talking about Cisco. That was Cisco's big entry, of course I mean, they've been a major player ever since.

Cohen: Sure. No, that was their entry. They had no, you know, as a big networking company, they had no wireless play, and so they bought Aironet in '99, around that acquisition.

Ennis: Right, and then there was Angela Champness, from Lucent.

Cohen: That's right.

Ennis: Who came from the-- that group in Holland, that had been originally NCR, and then got incorporated into Lucent. So she came from that group. And then there was John Ferrari, from Nokia.

Cohen: Right, yeah.

Ennis: Who had very little wireless LAN background when he came in, but he ended up being totally instrumental in a lot of the early marketing stuff.

Cohen: Sure. I know, he was a great addition to the board.

Ennis: Yeah.

Cohen: Yeah, we had a good mix of marketing, business and technical talent on the board, I think, and that kind of-- I think everyone got that it wasn't-- it was both a technical and a marketing effort, and we were able to kind of get a little traction that way.

Ennis: So once we picked the name, then the other big task that we had, was to really come up with a real interoperability certification program.

Cohen: That's right, the meat.

Ennis: And that kind of fell on my shoulders to organize with all these different companies, and the very first interoperability demonstration of this equipment, of what was to become Wi-Fi equipment, was at our public announcement of Wi-Fi at NetWorld Interop in Atlanta, in September of 1999.

Cohen: That's right, and June of '99, just to establish the timeline, we announced that WECA was an organization. I think that was the timeline, and that was our press release, like, we're here, and this is who we are, and this is what we're going to do. And then in September 1999, was our first, yeah, Atlanta NetWorld Interop was our-- we lifted the veil and said, you know--

Ennis: It's going to be called Wi-Fi.

Cohen: It's going to be called-- that was two things, yeah, lifted the veil and we told the world that the name would be Wi-Fi, that was officially announced, and we-- that's when we had our first interoperable equipment demo, I think.

Ennis: Right, right, because I had worked with the guys from the six founding companies, to develop a demo where we were showing interoperability. And remember, this-- there were the naysayers out there. The HomeRF people were saying this will never happen, this stuff will never interoperate, and so it really was a major thing that we had to demonstrate.

Cohen: Oh yeah, we proved a lot of people wrong, so that this very complicated technology could actually work in an interoperable manner, and--

Weber: So what was the demo?

Ennis: It was just laptops with PC cards plugged in, and you know, from-- one from Lucent, and then another laptop with a card from--

Cohen: 3Com.

Ennis: 3Com, and you know, an access point from Symbol, and you know, this different equipment and just showing that we could transfer data back and forth. Very simple.

Cohen: The client device, the station device could connect to all the different access points, and the access point could connect to all the different brands of station, you know, client devices, and you know, just simple data packet forwarding, nothing too whiz-bang. We weren't running video or anything exotic, it's just, you could make a connection, you can actually operate on this local area network, as part of this local area network, and among all these vendors.

Ennis: Yeah, and ensuring interoperability has been the focus of the Wi-Fi Alliance since the very beginning now, and I think that's a big part of Wi-Fi's success, is that this stuff works. You can take your smart phone to Cairo, and connect to that internet café.

Cohen: Yeah, I agree. I think that's a-- it's-- to the lay person, it's a small and subtle point, perhaps, but without it, you just would have none of that, right, if you think about it, if you had to know what brand of access point was sitting behind the network, and make sure that that matched, or was from a-- some compatibility list and you know what kind of Wi-Fi device was in your smart phone or your laptop or your iPad or something, and you know, just none of this would have really taken off, right, it's kind of just like you said, it was just kind of a no brainer, that you can just show up and connect.

Ennis: So the-- I guess the third thing that we needed to do, and which we did unveil at that NetWorld Interop show, was our certification lab, our selected certification lab that was going to do the actual testing. And that was a little company based in Silicon Valley, called the-- SVNL

Cohen: SVNL, Silicon Valley--

Ennis: Silicon Valley Networking Lab?

Cohen: Yeah, that's right.

Ennis: SVNL.

Cohen: Steve Bell.

Ennis: Steve Bell. And so I remember, I think you and I were the principle people negotiated that contract.

Cohen: That's right, Steve Bell worked with them, and--

Ennis: Yeah, and so I remember Steve, as we were working up the RFP for this, and he asked, well, how many products are we going to certify? Because he needed to factor that into how much money his revenue going to be and everything. So, we said, oh, we kind of scratched our heads, so 30. Yeah, so we've now done over 14,000 product certifications.

Cohen: Yeah. Right. No, but here's another part of the story. We said 30, and he was like, No way. He's like, you're crazy. Probably never be more than five. You know, he thought we were off our rocker to say 30. But anyhow, so he really didn't believe us. We had to kind of prove him over time that, you know, there was a market for this, and that there actually would be companies coming in with their devices, because he was calculating, okay, you know, two from Symbol, two from 3Com, whatever, you know, you've got five board members, you know. You'll be lucky if you break ten, you know?

Ennis: But other companies started joining up right at the beginning, I mean, it's been a ramp, and it's still-- I mean, it's continuing-- there's 500 member companies now, in the Wi-Fi Alliance.

Cohen: I just really keep pointing out, it really did catch fire, in the sense that once we kind of had this demo and showed interoperability, it wasn't that whatever, we were being interviewed by major news networks the next day, but I think there was a lot of companies that were interested in wireless LAN on the sideline and started to move off the sideline and said, okay, this looks like something interesting is happening, you know, and of course they knew that 802.11 ’97 standard finally got done, and ‘99 extension for 11b, and now you have this organization that's actually going to stand behind interoperability. I think you know, in the period of '99 through '01, '02, whatever, and you know, as you say, it still continues, but certainly in that first year or two, we saw a lot of companies come off the sidelines, and start to join and say, this is for real. Companies like Apple, companies like Microsoft, you know.

Ennis: Yeah, Dell was a very early--

Cohen: Dell. Big major companies, influential.

Ennis: Yeah, and ultimately the, you know, Intel who had been one of the main companies behind HomeRF, they ended up joining the Wi-Fi Alliance.

Cohen: Yeah, that was a big moment.

Ennis: We felt that that was a big moment.

Cohen: It was, it was.

Weber: When was that?

Cohen: What's that?

Weber: When was that, approximately?

Cohen: That was-- I know exactly when that was, because their representative, Barry Davis, talked to me, and this is the meeting that was in-- it was in the Netherlands, I think, this meeting, it was actually in--

Ennis: So that was probably about 2001.

Cohen: Amsterdam. Yeah, I think it was 2001, and I think the meeting was in Amsterdam, and I remember their representative, Barry Davis, saying you know, because he'd been coming to the meetings as an observer, you know, just to see what's going on, and he finally said, I think we're going to join the Wi-Fi Alliance officially, and drop the HomeRF effort. And he came to check with me, and said, you know, I hope you guys aren't going to, like, you know, kind of in a kind of-- dance around in celebration or like, you know, kind of rub our nose in it, or anything, or like, put out some announcement, like, Intel finally does the right thing or something. And he was really concerned about the PR aspect of it. I got a good laugh out of that, in this conversation over the hotel bar, you know. I assured him that no way, no, we're just glad to have you aboard, and we have no interest in making anybody feel bad. We're just glad they finally joined kind of the open standard. And that was really the thing, too, you know, Wi-Fi wasn't just somebody else's proprietary push, it was based on the open standard, so we had a lot of that going for us, too, so-- But you're right, when Intel joined, there was also a big moment. Then you had virtually every major player by that point, you know, Cisco, Intel, Microsoft, Apple, 3Com, you name it, all the semiconductor players--

Ennis: Speaking of Apple, so that was a major milestone in the Wi-Fi story also, was when Apple, you know, announced that they were incorporating Wi-Fi into their products, but not only that, but they were incorporating it into their laptops at a 99 dollar price point. And up until that point, nobody thought you could build a Wi-Fi device and break that 100 dollar barrier.

Cohen: That's right, that's huge.

Ennis: And that was a huge additional step that made everybody out there suddenly realize, oh, the cost of this technology is coming down, and it probably will continue to come down.

Cohen: That's right, that really validated Wi-Fi in a big way, because you know, the other-- you said, 802.11 had a lot of knocks, it was too complicated, it's never going to be interoperable, and the other one was going to be too expensive, right, and then, in a sense, that was part of HomeRF's reason for being, is we're focused on the consumer, it's going to be cheap, it's going to be consumer friendly, and then all of a sudden, you know, Apple comes out, has a Wi-Fi and 82.11b in a 99 dollar device, and it's like whoa. It's now not expensive, it's interoperable, it's, you know, it's everything that the naysayers said it wasn't, you know, and they really just disproved all the negatives. And it was definitely a major milestone. That caused a lot of-- also, computer manufacturers, notebook manufacturers to take notice. I said it was going to be-- was it the iBook at the time, or whatever it was, you can add a 99 dollar card to it. All of a sudden, Dell and you know, HP and all these other notebook vendors really started to take note, as-- and started on the path towards first offering network cards, and then embedded Wi-Fi in-- right into the notebook.

Ennis: Yeah.

Weber: Could you talk a little bit about the Wi-Fi Alliance as an organization? What were you officially, what was your revenue model, how many companies, how many employees--

Ennis: So the-- The Wi-Fi Alliance is a-- it's a nonprofit organization. It's an alliance of the individual companies. There are a set of companies who are sponsor companies, who, you know, by virtue of being sponsor companies, they have the privilege of designating representatives on the Board of Directors. Today, there are 16 sponsor companies, you know, like, I better not start to name them, because I'll probably forget some of them. But it's all the major players that you would anticipate. And there's like, 500 member companies, and the revenue stream supporting the activities of the organization initially, came completely from membership fees. The member companies pay a certain amount per year.

Weber: Roughly how much?

Cohen: Well in the beginning, the sponsor companies and part of it was, yeah, one of the original reasons for sponsor level companies and regular member companies is we knew that not everybody would put up this big fee, and it was a big fee. I think it was like 100K, you know, when we originally founded the-- you know, we needed some starting money.

Ennis: From the sponsors, yeah.

Cohen: From the sponsors, right, so the sponsors each put up like, 100 K per year, to fund the organization and get it going, and then the regular members, I don't remember the original fee, but much less, like, something on the order of you know, 10 K or something, or 15 K. You know, big, big, big difference. And Greg's point, yeah, the revenue model wasn't-- you might think, oh, they probably charged some of that for certification, but it wasn't the model at all. It was just really sponsor member fees and regular member fees, and then the money, at least initially, went to Steve Bell's SVNL Lab, you know, the one that thought he'd get no more than six certifications, right? In the very beginning.

Ennis: And it went to pay for me. I mean, I was--

Cohen: Yeah, right, right.

Ennis: I was the paid worker here.

Cohen: You were the one consultant. Although, I think, technically, and Greg, please correct me if I'm wrong about this memory, just to get the record straight. I think technically, you were still being paid by Symbol.

Ennis: So the--

Cohen: Not WECA. I don't think WECA wrote you a check. Not then.

Ennis: Right, well not right at the beginning, you're absolutely right that actually Nokia financed separately the marketing activities, and Symbol financed separately the technical activities, which was really the work that I was doing. But then I-- it wasn't even a year before we transitioned that, and then I had--

Cohen: Then we had a marketing director, a technical director, that were funded by WECA.

Ennis: That were funded by WECA, that's right. Right, I will say that some of the revenues for the Wi-Fi Alliance now does come from certification activities.

Cohen: Okay.

Ennis: So it's a mixed revenue model and there's a lot of stuff going on. It finances a fair amount. And that's because the technology keeps on advancing. I think when we first started, WECA, the Wi-Fi Alliance, you know, maybe you had a clear vision, but it wasn't clear to me what this organization was going to be doing you know, three or four years from now, but what ended up happening was, that the technology continued to advance, and there continued to come new amendments to the base standard within the IEEE standards activities for higher data rates, for the--

Cohen: For frequency bands.

Ennis: To operate in different frequency bands, then later on there was stuff for quality of service enhancements to support multimedia applications.

Cohen: Security, which is still an enormous topic.

Ennis: Security, right. And so all of this ended up just getting folded right into the Wi-Fi Alliance activities, to develop certification programs, not just for the base standard, but for these additional extensions to the standard in these various areas. And you know, now there's, you know, I'd say there's, today, you know, probably about, 20 different activities going on, within the Wi-Fi alliance, leading towards you know--

Cohen: Certification?

Ennis: The launch of potentially new certification programs that would be additional ones.

Cohen: Twenty new ones?

Ennis: Yeah, and there's just a tremendous amount of activity going on in all sorts of areas. We keep on launching new programs for higher data rates or for voice support or for supporting sharing of displays, using Wi-Fi, getting into all sorts of application areas, actually.

Cohen: Yeah, that's something we never envisioned that would get to that level of complexity, when we started. We just wanted to make 802.11b, you know, successful. But it was a natural step, at least early on, when we, you know, when companies, early companies like Atheros wanted to do 802.11a, which is in a different frequency band of five gigahertz, and they came to the Wi-Fi Alliance and said, you know, can we do certification around that as well? And the-- after some limited debate, I think the answer was pretty quickly yes. The alternative being to have some other organization start up and do certifications around a different 802.11 technology, I think we decided, correct me if I'm wrong, but I think we decided pretty quickly that it would be better to have that under the, you know, under the WECA, as we still called at that time, that under WECA, and later the Wi-Fi Alliance umbrella. So then we took on 11a, and then when G came out, at that point, that was kind of obvious that the Wi-Fi Alliance would back that as well.

Ennis: Yeah, so I think another one of the real major factors here in the success story for Wi-Fi, has been the complete adoption of Wi-Fi by the traditional cellular industry, and that was not necessarily a real easy hurdle. Then, you know, as smart phones began to come out, and companies such as Nokia, Nokia was very instrumental early on in this, but others as well, saw that they-- that smart phone functionality could benefit, not just from having a cellular radio in the device, but from having a Wi-Fi radio in the device for Wi-Fi applications, that introduced some real complications. And challenges. And one was that all of our certification techniques were oriented towards basically certifying PC type products.

Cohen: PCs in a Windows environment, or things like that, yeah.

Ennis: And so the Wi-Fi Alliance had to take on a whole new set of test tools and test environment, just to be able to certify these hand held devices. But then also there was a whole issue of having two radios implemented-- you know, a cellular and a Wi-Fi radio implemented right next to each other within a single device, where even if they're operating on different frequency channels, because they're so close to each other, they can be interfering with each other. And so that led to a whole effort where we worked jointly with CTIA, the Cellular Telecommunications Industry Association, to develop a joint certification that would make sure that when Wi-Fi is implemented in smart phones, that it really works well with the cellular radio that's also implemented there. And that really brought in the whole operator universe into the Wi-Fi Alliance, and ultimately was very successful. I mean, now you can't buy a smart phone without Wi-Fi.

Cohen: Sure. Yeah, there was a time where, you're right, initially, early in the smart phones, a lot of the providers were really kind of anti-Wi-Fi, and they saw it as, oh, it's a competitor to my service, I want to sell my whatever, at the time 2.5G, 3G, you know, service, and I don't really want Wi-Fi in my device. There was some kind of pushback early on, if I recall, you know, there was really some, you know, demand for Wi-Fi, but not necessarily carriers agreeing to do it, but that certainly has changed dramatically in the last few years.

Ennis: Right. Well, I think the iPhone was a big part of that, because when the iPhone came out with Wi-Fi, and the, you know, and the 3G cellular service, with kind of this, you know, unified access to the internet, whether you're on Wi-Fi or on 3G, and people fell in love with the data application so much that the cellular networks, the 3G networks couldn't handle the demand that the iPhone users had for all of this bandwidth. And so Wi-Fi suddenly became something that wasn't-- it's not just this, oh, kind of interesting little adjunct, but it suddenly became a critical component of the smart phone, because there's no way that the 3G network could satisfy-- could really handle the applications that people wanted to do with that. So Wi-Fi became the preferred way of doing that.

Cohen: Yeah, it went from being something that the carriers didn't want, to almost being a savior technology, because of their bandwidth got totally overloaded, especially, as you say, as data apps took off, and people stopped using their cell phones just for talking, and went to everything, texting and you know, other data apps, Facebook posting, all this kind of stuff, you know, then the carriers got dramatically overloaded, and yeah, the iPhone had thousands and thousands of applications, and data oriented that want a network connection. I think you're absolutely right, I think Wi-Fi has played a tremendous role in helping carriers to roll that out in a way that doesn't bring down their whole network. I mean, they're still struggling with, you know, not enough capacity for data applications, and thinking about how they can move more stuff to Wi-Fi, rather than less.

Ennis: So what's interesting there is that, yeah, there are advances in the cellular technology to get increasing data rates, you know, over the years. But there's also advances in data rate within Wi-Fi. I mean, we've gone from the two megabits per second, of the original 1997 standard, to the 11 megabits per second of the original Wi-Fi, and then 54 megabits per second for 802.11g, and then up to over, you know, hundreds of megabits per second for 802.11n, and now there's new things coming out of the IEEE standards that will come out over the next couple of years, that will support over gigabit rates. So the thing that's interesting is that yeah, the cellular data rates may increase, but so does the Wi-Fi data rates, and the cellular rates are never going to catch up to the Wi-Fi rates at any snapshot in time, you're always going to have much more performance out of the Wi-Fi part of your device than out of your cellular part of your device, and so consequently, people will use that, and there will always be a demand for, you know, the highest performance. So even as the cellular technology improves, you know--

Cohen: Wi-Fi is still--

Ennis: Wi-Fi is still going to be the higher performing, I would say, preferred alternative for-- certainly for the high performance applications.

Cohen: Yeah, now if we could only get a local broadband to keep up the with all the [Wi-Fi data rates], but you know, but you're right, Wi-Fi is evolving extremely quickly. Think about, you know, the forthcoming 802.11ac standard they're talking about, yeah, so many variants and options, but you're talking about many of those link rates of well over a gigabit per second. That's pretty amazing when you think about it-- just had gigabit Ethernet not too long ago as a breakthrough, and now we're doing you know, gigabit plus Wi-Fi soon to come, so that's pretty wild.

Weber: So from early years, like 2001, when it really was clear that you were growing Wi-Fi, what were the main kind of stages between then and now.

Ennis: Well, I can talk from my perspective of the certification program.

Weber: From yours, and from the Wi-Fi Alliance.

Ennis: So one of the-- this is kind of an operational transition that ended up having a big impact. You know, when we started the Wi-Fi Alliance, we had this-- we were working with this little company, they subsequently were bought by Agilent, and they continued to serve as a certification lab for the Wi-Fi Alliance. But very quickly, we swamped their capacity to be able to certify products, number one, but number two, the industry took off in a huge way in Asia. In particular in Taiwan, but also in Japan, Korea and now in China in a very big way. And so the-- and as the Wi-Fi Alliance grew in its membership and ended up with a very global membership, in fact, now it's-- half of the member companies are Asian companies and then the other half would be North American and European, and we needed to support certifications local to the-- to where those companies were developing their products. So we changed our model from having a single lab, doing certifications, to having multiple labs, and I mean, this was one of the projects that I did, was a-- you know, kind of a grand tour of the world, looking at potential lab partners that we could contract with, to take on Wi-Fi certifications, and you know, we now have a network of 17 labs worldwide, that do Wi-Fi certifications. You know, we have a lab in India, we have multiple labs in China, we've got multiple labs in Taiwan, we've got a lab in Spain, a lab in Germany, you know, multiple labs in the United States, and that just-- it was absolutely necessary, operationally, just to handle the volume of products that was coming on. So--

Cohen: But I think the operational model changed, too, from you know--

Weber: What year, roughly? When did you do the grand tour?

Ennis: That's a good question. That's--

Weber: Roughly.

Ennis: That was probably around 2002.

Weber: So early, wow.

Ennis: Yeah, fairly early, because I'm thinking .11g came out in 2003, and we had the labs in Taiwan in place at that time. So--

Cohen: And at some point, the operational model changed from the Wi-Fi Alliance you know, having a single lab, to actually labs could come and get certified by the Wi-Fi-- or you tell me what the right method was, but labs could kind of come in and apply to be-- right.

Ennis: Yeah, essentially. We still don't let just any lab do Wi-Fi certifications.

Cohen: They have to be a member of the Wi-Fi Alliance.

Ennis: I mean, there are other organizations that are certification programs where they-- it's kind of a free for all, where labs can sign up to do certifications and-- but we've always had the approach of maintaining some level of control over the labs that actually do certification. I think that's--

Cohen: That's a key part of keeping the quality up and then making sure that-- although you could lose interoperability and lose both.

Ennis: Right. But we should probably say, when we talk about these certification programs, those get developed within the Wi-Fi Alliance and it basically is a test plan that identifies a set of devices that are used to verify that other devices can be interoperable with them, and it's a tremendous amount of effort that actually is primarily the volunteer effort from the member companies within the Wi-Fi Alliance who form committees within the Wi-Fi Alliance, to develop this, and then there's a Wi-Fi Alliance technical staff who ends up implementing what gets developed into a real operational test environment, and rolls that out to these 17 labs worldwide, so there's a tremendous effort that gets contributed from all the member companies in developing all of these programs. And that was the model in the beginning and it continues to be the model today.

Cohen: Yes, any member company can join with other member companies to form a committee, right, and if they're interested in a particular kind of certification, right, they can, you know, join together and say, hey, we're interested in this process, and those people were approved and all this stuff, that if there's a need for a certain type of certification, like the example you gave earlier, where you have smart phones and you needed to actually have joint cooperation from the CTIA, you know, there's kind of a method for being able to do that. And other specific certifications, there's voice certifications, and what are some of the other examples of certifications that we've done recently?

Ennis: Well, there's-- there's what we call the ease of use certifications, which has to do with being able to-- it's, you know, essentially user interface conventions that products can get certified against, to make sure that networks can be set up very easily by uneducated users. There are certifications that are specific to support applications like video, and other multimedia applications that-- where the bandwidth needs to be handled in a prioritized way, to make sure that the higher priority traffic is being dealt with properly,

Cohen: It's a quality of service, certification.

Ennis: Quality of service, that's right.

Cohen: In Wi-Fi speak, it's Wi-Fi multimedia, right, so--

Ennis: That's right.

Cohen: So I think it might also be good to explain that, you know, when you think about how many certifications are done a year, thousands of certifications a year, right? And represent 14-- Wi-Fi Alliance up to--

Ennis: I think it's about 3,000 a year.

Cohen: Yeah, that's a lot, and people think about it, they might think, well, you know, when you bring a product in for certification, do you have to make sure that it's interoperable with, you know, 3,000 other devices? How does it actually work? I know the answer, but I think it's good to explain how the certification philosophy works. How do you prove that it's interoperable?

Ennis: So yeah, and actually, the model that we have here has been replicated by other organizations in other technologies, so I mean, a lot of people look upon what we do here as really the right model. And it involves a series of-- although it's wireless, we still call them plug fests. Maybe we should say unplug fests. But it's, you know, relatively formal get-togethers on the part of the manufacturers of equipment, prior to any certification, in fact, often prior to even the finalization of the relevant standard in this particular certification program. Where the companies get together and they determine whether or not they can interoperate, they do tests and they do all sorts of examinations to make sure that they're interpreting the specification correctly, very much like what we did to-- at the very first interoperability demo back in Atlanta in 1999, of getting the companies together and making them play well together. And there's a whole series of these events that are managed by the Wi-Fi Alliance, and most of them are held at our Wi-Fi Alliance lab, which is in Santa Clara. And at the-- while we're allowing companies to test out their products against each other, we're also simultaneously developing a test plan to actually ultimately certify that. And at the end of the process, a set of devices is selected to form the core test bed that will end up getting replicated to all of the labs worldwide, and those form the core devices against which interoperability is checked, and we always-- we have, you know, like at least five independent implementations. And very often, I'll say this, the companies who very heavily participate in this part of the development of the Wi-Fi certification programs are actually the semiconductor companies who are producing the Wi-Fi chipsets, and it ends up very commonly being their reference designs that are full implementations of a real product, but it's a reference design showing people how they can build a fully featured Wi-Fi product, using a particular chipset. Those products very heavily participate in the plug fest series, and commonly end up even within the test beds. And consequently, you know, the interoperability that's assured for that particular chipset reference design, ends up carrying on in all the customers to that chip company, who then go on to implement their own products using that chipset. Of course those companies still need to certify their products, but there's been a head start that they've gotten because they know that the fundamental chipset has already gone through this process.

Cohen: Yeah, I think that's a good point. By having the semiconductor companies so active, you know, and that's, I think, to mention the names, you know, who are some of the most active players, and you know--

Ennis: You know, it's Broadcom, it's Qualcomm Atheros, it's Intel, it's Ralink, it's, you know--

Cohen: Yeah, yeah. At the risk of leaving out one or two, yeah.

Ennis: Marvel.

Cohen: Marvel, yeah. And it won't be the same one for every certification program, but a lot of the same players show up, and yeah, it's a good way of going and explaining how interoperability works, and if you can get those basic reference designs interoperable, you know, Broadcom's reference design access point works with Qualcomm Atheros' reference design station device, and on and on, and then all these other OEMs essentially build off those reference designs, well you've got a really great head start for interoperability. Yeah. I thought it might be-- people would want to know, is that when a product comes in for interoperability testing, you know, they're not literally testing against the existing, whatever, 3,000 products that year, or 14,000 products over the history of Wi-Fi. That would obviously take forever, and be unfeasible. They come in and test against that reference design test bed, or whatever is in the test bed, specifically reference design, and it's kind of proxy interoperability, right, so if you come in as a station device, and you're interoperable with these four or five access point designs, then, you know, by proxy you'll be interoperable as you go out there in the real world.

Ennis: Right, and that's why it's important--

Cohen: It's a small logical leap, but in reality, it seems like it's actually worked very well.

Ennis: It's worked tremendously well and that's why it's important that the devices that do go into the test beds are truly independent designs. And you know, if you can work with five different independently designed implementations, then you're-- and then you're very likely going to interoperate properly with other equipment that similarly interoperates with all those five different independent implementations. It's proved to be a model that is very effective, and has really stood the test.

Weber: When you talk about a test bed, this is every time there's an update to the standard. I mean, how often do you create one of these reference sets?

Ennis: So we launch a-- probably something on the order of five new certification programs each year, and so each one of those new certification programs will be covering some new aspect of the-- either the IEEE standards, or it could be a specification that actually gets developed within the Wi-Fi Alliance, for example, Wi-Fi Direct, which is a peer to peer technology to allow direct data sharing among devices without needing to go through a hotspot, or be connected to the internet at all. That's a mode of operation for Wi-Fi devices that actually got specified and developed within the Wi-Fi Alliance, not within IEEE, but so that's an example of a certification program that we launched a couple of years ago. And then, you know, like I said, there's-- every year, there's approximately five new certification programs that have to do with some new-- either some new core technology standard, or some new application that we end up now allowing people to certify their devices for interoperability on.

Cohen: I think there was also a big milestone, you alluded to that briefly, in the history of the Wi-Fi Alliance, initially was, we're certifying based on IEEE 802.11 standards, and then, well my recollection, it was security that first broke that model. It was such an urgent need that we actually developed what became WPA Wi-Fi Protected Access, you know, in that case it was kind of before the standard, you know, the actual 802.11, in that case, 802.11i standard was ready, just for the urgency. But then I think that set the model later for doing things that were completely independent from 802.11, everything from Wi-Fi Protected Setup, to now Wi-Fi Direct.

Ennis: Right. Another good example is with 802.11n, which I mean, right now, .11n is the highest data rate version of Wi-Fi, and it's implemented in all the retail devices that you're going to purchase these days, but when that version of the IEEE specification was going through the process within the standards committee, and being finalized, it was going through a process of draft standards and balloting and you know, various conversations about one technical solution or another on some minor point of the evolving standard. And what was happening was that companies were starting to come out with products that were based upon an unapproved draft standard, and calling them 802.11n, and there was a real worry within the industry that these products were not going to be interoperable, and that was going to give the whole Wi-Fi industry a black eye, and was going to sour people on the interoperability promise that Wi-Fi has always offered to people. And so the Wi-Fi Alliance decided to actually start its certification program in advance of the finalization of the standard, by identifying a particular draft specification and saying, we're going to certify to this draft specification, and that was a year in advance of the ultimate 802.11n ratification in IEEE, but it turned out that there were only negligible differences between the ultimate standard, and what we started to certify on, and that ended up giving a huge kick start to the whole .11n market.

Cohen: Yeah, I think it's a great example of the Wi-Fi Alliance being flexible, and realizing again, the market need, right. It's not a standards body where you always follow the exact same rules, but if, you know, if it's in the best interest of the industry, the best interest of commercialization, the best interest of the end user experience, and not breaking the promise of interoperability, sure, we can do that. We can do-- and I think it was what, draft 2-0, that was adopted as the interoperability program first, and then later, it was just the final 11n standard, and really, at that point, it was really no big deal to go from the draft standard to the final standard, and you know, you had a smooth transition, you had interoperable devices, good user experience, versus the alternative, to have, you know, the mess that you alluded to that probably would have happened with a bunch of devices out there, not interoperable, and then waiting for the final standard, you know? And I think, you know, this is from my recollection, even from very early on, there was a spirit in WECA and in Wi-Fi Alliance that, you know, we're not held to some official rule book. We want to operate fairly to all the members, but it was-- there's a nimbleness, I think, built inside WECA and the Wi-Fi Alliance that, yeah, we could react to these market forces, and to the interest of the members and the interest of growing the market and solving real problems, where it was-- we wouldn't be tied down by, you know, an IEEE charter, or a PAR, or Roberts Rules of Order, whatever you have, it was kind of, as the Wi-Fi Alliance grows really big, that may be harder to do all the time, but I think even your recent example proves that, that there's still kind of a philosophy from the very beginning, of, we're here to solve problems, we're here to grow the market, we're here to serve our members, we're here to ensure good end user experience, and we'll do what we have to do to get that done. So I think that stayed throughout.

Ennis: Yeah, so I think another thing that, you know, in thinking about the key transition points in the evolution of the Wi-Fi Alliance, was that, you know, originally it was just the Board of Directors and me, basically, and we decided, you know, as we grew in members, that it was time to bring on a real organization. And we hired Frank Hanzlik, to be the first managing director, and--

Cohen: Actual paid staff.

Ennis: Actual paid staff, and with actual Wi-Fi Alliance office, yeah, besides my home office. And yeah, and Frank led the organization for some very critical years, through the .11g, and he brought on a tremendous staff, including Edgar Figueroa, who is the CEO of the organization now. I mean, Edgar started on the technical operations side, so he's-- and he's now the CEO of the organization. So that was a-- that was certainly a major-- of course every organization goes through that from its startup phase to professional management.

Cohen: No, no, you're right, because in the beginning it was all, you know, we had the board, which was all volunteer run-- well you as a consultant, but all volunteer run, and you know, that was-- it was a burden on everybody. It was needed and we're glad we did it, but yet, at a certain point, you know, it was like, hey, it's time to actually hire some people to run the shop.

Weber: What is your relationship like with the relevant parts of the IEEE, the 802 group?

Ennis: So-- we are-- I think we would be referred to as the certification authority for the IEEE 802.11 family of standards. There's nothing formalized there. It's not like there's some higher level body that has designated the Wi-Fi Alliance as--

Cohen: There's no signed agreement, as far as I know.

Ennis: Yeah, it's just simply the way that it is. I mean, nobody would question that. And the-- so it's a, you know, to some extent, it's sister organizations. In practice, many of the same individuals participate in IEEE as they participate within the Wi-Fi Alliance. Within the Wi-Fi Alliance, you know, as we mentioned before, it's an organization of member companies so you-- you know, when people are participating within the Wi-Fi Alliance, they are clearly representing their companies' interests, and voting is by company, for example, whereas when they're participating in the IEEE committees, they're there as an individual engineer, and voting is by individual without company affiliations. But nonetheless, it's, you know, a lot of the same people are participating in both. And so there's a lot of very good cross fertilization. So for example, when a specification is being developed within the IEEE for some new extension to the standard, part of the feedback that goes into that process comes from the plug fests that the Wi-Fi Alliance is hosting for early prototype implementations of this draft standard, and then those people recognize, oh, you know, you interpreted that line of the specification in this way, and I interpreted it this way. We need to get that clarified in the spec, and that will end up going back to the IEEE specification. So it's a very close relationship, but there's actually nothing formal.

Cohen: So, I think we talked about a lot of the different aspects. Maybe just, you know, can you pontificate a little bit on the future of-- from your perspective as technical director all these years, the future of the Wi-Fi Alliance? What's the Wi-Fi Alliance going to be like in 2014, 2018, and 2020?

Ennis: Well I can tell you that we just signed a lease with, I think a seven year commitment on lab space that is four times the size of our existing lab here in Santa Clara. And so clearly we're all anticipating, you know, tremendous continued growth. And it's not just in the, you know, the numbers of devices that get sold, but it's in the diversity of devices, you know, basically Wi-Fi has become the default common language that all high tech devices use to communicate with each other. I mean, you think of it. You've got a cell phone, and you've got, you know, so there's other-- you've got a cellular radio in there, but you also have Wi-Fi. You've got your HDTV flat screen television, you know, it's got HDMI connecting to it, but it's also got Wi-Fi. You've got, you know, you're laptop, it's got all sorts of connectors.

Cohen: But it's got Wi-Fi.

Ennis: But it's got Wi-Fi. So you know, there's all these different communication methods, and they all have their place, and they all solve some kind of a specific problem. But what Wi-Fi has kind of evolved into, is the common one that goes into all of these devices, and you know, Wi-Fi long ago passed that knee in the curve of the positive feedback loop, you know, I think it's called Metcalfe's Law, right, that the-- I'll be misquoting this, but it's something like the utility of a networking standard is proportional to the square of the number of devices that implement it. Well Wi-Fi long ago passed the point where the utility is just tremendous for everybody, because all these devices have it. So I think that's what-- that's one of the things that we're going to continue to see is just Wi-Fi going into essentially every device.

Cohen: Yeah, any device that needs to be connected.

Ennis: Right, and so this will include sensor networks, this will include applications to support smart energy, for example, you know, appliances in the home, all sorts of-- I mean, we can't envision what the devices are that we're all going to be using five years from now, but I'm pretty sure that that device is going to have Wi-Fi in it. It may be different from the Wi-Fi that we see now, but it will be backwards interoperable with the Wi-Fi that we have now.

Cohen: That's right, that's been a great promise, too, is that it really held to that promise of interoperability. You can even take an old Wi-Fi device, and show up into a network, and all likelihood if it's Wi-Fi certified, it's still going to work, you know, that is really quite a promise that's been delivered on. So I agree, in lots and lots of devices, lots of different types of devices, lots of different types of certifications, that may aid, and depending on the kind of device, and the application, like you said, whether smart energy, voice, video, things like that. Yeah, it's been exciting, and it will continue to be an exciting road, I'm sure.

Ennis: It sure will.

Cohen: Well good. Well thank you, Greg, very much for joining today.

Ennis: Okay, well thank you, David.

Cohen: Sure, pleasure speaking with you.

Ennis: My pleasure.

END OF INTERVIEW

