

Pittsburgh PA
1:21:01
m4v

Ralph Hollis

An interview conducted by
Selma Šabanović
with
Peter Asaro

November 22 2010

Ralph Hollis: Okay. My name is Ralph Hollis. I'm at Carnegie Mellon University in Pittsburgh. I'm 69-years-old; that means I was born in 1941. I was born in Hutchinson, Kansas, but grew up most of my life in Wichita, Kansas. So let's see. My earliest recollection of being interested in robotics was probably around 1954; and I had read a pair of *Scientific American* articles by W. Grey Walter. I think it was the 1951 and 1953 issues of *Scientific American*. Anyway, as is well known, Walter built a couple of small mobile robots which could move around the house and could sense objects and avoid obstacles and so forth. So I thought that was like the neatest thing ever; like that was it.

So right away – also I had read somewhere about sensing things using capacitance. So let's see, I would've been about 13 I guess at that point. So my friend, Phil Roberts, and I decided to build a robot that could move forward until it could sense a wall and then it would stop, based on measuring the capacitance to the wall. So I had gotten I think as a Christmas present a big old lead acid battery. So that was the first ingredient. And then we went to a local surplus place in Wichita, called The Yard – it's very famous; it's still there – and bought some motors and so forth to make this thing. But it wouldn't really go because the battery was too heavy. And so eventually we gave it up. But we began to understand the problem of carrying around a battery with wheels. You had to – it was a tradeoff between having the battery too big and not being able to move; if the battery is too small, not being able to move. So that was a pretty cool lesson. So of course this was back in vacuum tube days actually as well. So that was my first sort of introduction to robotics.

Later in about I think 1957, my friend Phil and I spent a lot of effort and we built a small robot that was in the shape of a cylinder basically. It was about 10 inches in diameter and maybe 10 inches high. It had lead acid batteries and was mainly made out of old Erector set parts, as I recall, but it used telephone relays for a brain actually. So we had lots of surplus telephone relays on this. And it also had a pair of feelers that could extend from either side and it could roll – and the feelers had wheels – and it could roll along the wall. And as it rolled along the wall, it kept an equal distance from the wall. When it came to an opening like a doorway, the feeler would spring forward and hit an electrical contact, which would trigger a relay to cause the robot to turn and so forth. So it could occasionally find its way out of a room by following the wall around the room; which was pretty cool. And I still actually have parts of that robot; although I don't have all of the robot. Back in those days once we – once something worked we took it apart and did something else with it. So we don't do that anymore. We always keep what we have. If it works, that's it. So well, that was basically high school days I guess.

I went to University of Wichita, which is now I think called Wichita State University, for two years and majored in mathematics there. And one of the big things was my first introduction to the digital computer. Actually I'd started a year earlier in 1959 when I was still in high school programing the IBM 610 computer, which had a drum memory and paper tape and so on, vacuum tubes. And by the time I was done with doing all kinds of calculations with that

computer, a couple of years later, it was – I'd logged more hours in the log book than any of the faculty and other students at the university. So I was quite enamored by computing. I also – at that time the Bendix salesman came around to Wichita University with the Bendix G-15 computer; the size of a refrigerator, and used this cool language called INTERCOM 500, which I tried to learn. And so I programmed that; and it was much, much faster than the 610. And it had von Neumann architecture, as opposed to the 610 which actually did not have the ability to store programs; they were all on the paper tape or on the wiring panel.

So I programmed the G-15 for many months until the Bendix salesman came and took it away; because it was just a demo. So people took pity on me. And then I went – to the local aircraft company, Beech Aircraft, had three Bendix G-15s. So they let me use their computers from midnight to six a.m. So I would spend all the nights over at Beech Aircraft doing calculations. Almost all of the things I actually computed had to do with rocketry; which was my hobby really. I built a lot of rockets and lots of stuff like that. I did a lot of calculations on both the external rocket's flight and the internal flow of gas dynamics and so forth; read a lot of books on fluid dynamics. So well anyway then I went to – after that to Kansas State University where I switched to majoring in physics and got a Master's Degree – a B.S. and Master's Degree. Didn't do much work in robotics I guess at that point. Moved to California; worked for North American Aviation on space flight vehicle simulation. And the reason I got the job there was because of my experience with doing calculations related to amateur rocketry. But this was with the Minuteman III system. So I formed a team of eight people and we wrote this computer simulation of the Minuteman flight from takeoff to impact of the warheads. And it was sort of a robot, if you will. It did a lot of measurements. It had an initial measuring unit and so on. So a lot of the inputs and outputs and computation and so forth; very much like a robot.

So that was up till early 1970s. And – but during that time I still was very interested in robotics. So I started building a computer at that time based on old parts from wrecked IBM 1620 computers, as well as the very newly available Texas Instrument TTL logic. As it turns out the Minuteman missile was the first use of TTL logic from Texas Instruments; and they had a local office there and so I would visit the Texas Instruments people and they would give me circuits and so forth. At that time there were only about four different circuits available from Texas Instruments – hard to believe now with the millions of chips that are available. But using those basic circuits, I built a computer that could be like a couple of hundred thousand additions per second; and also was able to get a surplus magnetic drum memory, which was part of the SAGE early warning system, and a Friden Flexowriter for input/output. And then the idea was to control the robot remotely. So I hung out – as it turns out my friend, Phil Roberts, the same one that we built the robot with me in the 1950s, was a grad student at Caltech. So I was in Anaheim, California. So – in Orange County – so hung out a lot at Caltech. The two of us built another robot, which sort of got half built. And the idea was to interface it to all this stuff that I had; the computer and so on. So that never really happened. It never really quite got together. But the – a lot of the ideas were there. So let's see. That brings us up to about maybe 1970, I

think. So actually that wasn't quite right. It was – I was in North American Aviation from 1965 to 1970.

So I went back to school to get my Ph.D. at the University of Colorado, again in physics, where I succeeded in putting together a superconducting magnet system with lasers – laser optics. At the time it was – became the highest field magnet in the world with optical access; which was kind of neat. And I got a lot of results in solid state physics, light scattering from semiconductors. But I still had a hankering to do robotics. So I started putting together a small robot that had stepping motors; which for some reason I called Newt, 'cause my girlfriend thought that would be a good name. It comes from Shakespeare – “eye of newt, toe of frog”; or something like that. So that became Alpha-Newt. And it was able to roam around and see lights and avoid lights and go to lights. So basically in the 1970s, I was reproducing what Grey Walter did in the '50s or late-'40s. So but from that we – there was another friend I had, Dennis Toms, who was another physics grad student. We convinced the university to let us have a lab room; and with a group of about four of us we started doing all kinds of robotics things, building cellular automata and things like that out of hardware.

And so we wound up building a robot which we called Beta-Newt, the second version, which was – had an onboard computer; which was a significant advancement. And so in – about that time, I guess – so I graduated with my Ph.D. in '75. In '76 I went on to an exchange scientist fellowship to the University of Paris for about 10 months; also got married. So and there I decided to write an article about this robot we had created that had the onboard computer; which was an 8080- Intel 8080 I think it was. And I wrote that for *Byte Magazine*. And at that point I had already submitted several articles to physics journals, like *Physical Review* and so forth. And I was amused by the fact that it cost actually – it cost page charges to publish in some of these journals. I submitted this article on Newt, the robot, to *Byte Magazine* and it became the cover story. And they actually gave me like over a thousand dollars for the article. And as it turned out, it was the most popular article *Byte* had published up to that point. They had a little thing that they rated user responses and apparently it was several standard deviations from all their previous articles. So that was cool.

Q: Do you know what some of these responses were? Did you ever read them or –

Ralph Hollis: Yeah, yeah oh yeah, yeah.

Q: What were people saying?

Ralph Hollis: Well they thought it was really cool; and I don't know, they – it was a long time ago, I don't remember exactly what the responses were. But I started to think: Wow, gee there

must be a lot of interest in robotics, compared to physics, for example; so maybe there's something to this.

Q: Why was robotics – you've obviously had a big interest in it from very early on.

Ralph Hollis: Yeah.

Q: So why did you decide to go into physics rather than something closer to robotics?

Ralph Hollis: There was no robotics.

Q: Right.

Ralph Hollis: There just wasn't. There wasn't even a RadioShack around to go buy parts. I used to travel all over the country to surplus places all over, to get parts. And but there was just almost nothing. There was no Digi-Key to order parts from or Newegg or any of these kind of things of course. So and I did like physics too, for sure, and I was reasonably good at experimental physics. So yeah. So the *Byte* article had a big response. And sometime after that I guess the WGBH NOVA broadcast came out to Colorado, and they spent three days filming Newt – Alpha-Newt and Beta-Newt. And they had a program which I think came out in about 1978 called "The Mind Machines," which featured my robots; which I thought that was pretty cool. And *Wall Street Journal* had a small article on the front page. I got interest from just all over the place; like companies from China that wanted to make them, things like that. So at that point I moved to New York. Yeah, I got a job with IBM Research at Yorktown Heights, which was good. And I worked in physics, in applied physics, doing magnetic field effects in thin wires, doing microwave acoustics and what have you. So that was great and a lot of fun and a terrific place to work; fantastic place in fact.

I wound up spending 16 years at IBM Research and don't regret any of them. So but because of a change in management and a few other things, a change of projects, I wasn't able to continue what I really had started, which was the microwave acoustics. I built the first acoustic microscope in use in industry there, working with Cal Quate at Stanford, who was the real pioneer in that area; and I wanted to continue that. But as it turns out I sort of got a cease and desist order 'cause IBM isn't a university really. So at that time I was kind of looking around for something to do and there was a newly started – fairly newly started robotics effort at IBM. So I decided to – after all this work and education I decided to leave the field of physics and do robotics full-time. And that was a really, really tough decision 'cause I'd invested a lot of energy into physics; and I love physics. But I also love robotics. So all along I guess there's been this duality that has followed me.

Q: What year did you make that switch?

Ralph Hollis: It would've been about 1982 I think, maybe.

Q: Who else was working in the robotics group at IBM at that time?

Ralph Hollis: So Larry Lieberman was one name. Russ Taylor; who you're doubtlessly going to be interviewing, if you haven't already. Now you got me on the spot.

Q: Just the ones you remember. Those are the main ones.

Ralph Hollis: I do remember but I don't; 'cause I'm getting senile.

Q: What were the projects or the applications they were looking at?

Ralph Hollis: Right. Jeanine Meyer is another name. So I'm trying to think of – Mike Wesley. David Grossman; David Grossman was the senior manager. So I wound up reporting to Russ Taylor eventually, after he became manager – a manager there. So at that point IBM had already developed a series of hydraulic robots, Cartesian robots, which became products that were productized in the IBM Boca Raton facility. And they used an interesting sort of hydraulic stepper motor design that a guy name Hugo (Pat) Panissidi had developed. Pat had about 50 patents in IBM, which was pretty nice. So Russ Taylor was responsible for developing the programming language, or at least mostly responsible, called A Manufacturing Language or AML. That was based on earlier work that he had done at Stanford using the Scheinman arms and so on and so forth. So that's another story I guess. So they'd also recently introduced a SCARA robot, the shoulder and elbow and a four-axis SCARA robot mechanism from Japan; but the software from IBM. So it was a pretty exciting time because these robots were starting to be used in many different manufacturing applications within the company.

So my personal effort was in developing – mostly in developing fine motion devices. The robots, any kind of robot you had, they were not very accurate; maybe half a millimeter or maybe a little bit better maybe, 200 micron accuracy. So a lot of IBM products required a lot more precision than that. So the question is: How would you do that? And so what people mostly did is they used very high precision stages that were very expensive; granite bases and laser interferometers and all that to get down in the micron range, which was out of reach of robots basically. So I developed – at the suggestion of Russ, I developed a fine motion device, which we called the Fine Positioner – Fine Positioner, IBM Fine Positioner, which you could mount on the last link of the robot arm; and you could then enhance the precision of the robot

because the robot could move roughly to the position you wanted to go, and then the fine motion device could take over and move from say a millimeter on down to one micron or even better than one micron. So we were able to do a lot of things like precision assembly and precision testing, where we could put an electrical probe on a circuit that was very, very tiny; which would be missed entirely by the robot itself. So we called this paradigm coarse/fine positioning. And I think we were among the very first to really do this kind of thing where you had redundant degrees of freedom on a robot. The robot – the standard robot arm would form the coarse robot; and then the fine motion device on the end of the robot would take over and do the rest.

Q: So is much of this with Russ?

Ralph Hollis: Yeah.

Q: <inaudible>

Ralph Hollis: Yeah right. So and also Mark Lavin – that’s another name – who did some vision work with us as well. So we built a number of these based on flexures, metal springs. I have some with me here actually. It was a big hit. It got on the cover of the IBM Annual Report and like four-million copies in I don’t know how many languages; things like that – which was cool. That early fine positioner was able to move in two degrees of freedom; X and Y let’s call it. The robot still had to do the Z motion. Later I developed another kind of fine positioner that used an air bearing and could in move in X, Y and theta – that is transversely – and also a certain amount of rotation. That was based on a flux steering motor. I got an A in electromagnetic theory; so that was easy. So and that worked quite well and became sort of an internal product to IBM. We made a bunch of them – I don’t know, maybe 12 or 15 – and started to use them in a number of prototypical applications in Japan and France and all over the U.S. and so on.

So for example, one could put a disk platter onto a spindle by having the robot bring that disk platter down to the spindle and then having the fine motion device, which was this air-bearing fine positioner, sense this inductively and be able to guide itself on. So the robot could be off by millimeters but still the spindle would go straight onto micron tolerances, and do it fast. So and all kinds of other applications: stacking thin laminates, like circuit boards, for really high-super-high density circuit boards, things like that. So we built a number of systems around the fine positioner using mostly the SCARA robots and the fine motion devices or fine positioners. Let’s see, about that time I also started wondering – it would be really nice to have six degrees of freedom instead of just three degrees of freedom: X, Y and theta. Why not up and down and rotate around – roll, pitch and yaw. I couldn’t quite figure out how to do the linkages. I guess I could’ve used like a Stewart platform kind of thing. But I thought about it and thought about it for a long time.

And then one evening it occurred to me that: Why not just levitate the fine motion device in magnetic fields? And then you could move it in any direction you want and there would be no suspension or no linkages, no air bearing, no anything. So I sat down and I did the calculations and thought of a few geometric configurations. And at that point there was recently available magnet materials, Neodymium Iron Boron, which had just come out; and then also there were digital signal processors available that could do really fast computation, that I thought we needed in order to keep the levitated part from falling basically. And of course there's no stable equilibrium points in a magnetic field due to a theorem by Earnshaw in 1845. So it has to be a closed loop system with sensing. And I'd already used position sensing photodiodes for doing the sensing on the three-degree of freedom fine positioner. So anyway we started putting together a system that could levitate. We called it – well it was something sitting on the desktop actually; and hooked up to some surplus power amplifiers. And we had like a 9000 dollar signal processor from – with chips from Texas Instruments, and an IBM PC – an XT or something like that. I had hired a fellow from – a new graduate from Berkeley, Tim Salcudean; and he got his degree in Advanced Control Theory from Berkeley. And so he teamed up with me, and also an intern from Duke University named Peter Allen, a student. And we made the thing work. Tim did most of the software and algorithms for control.

And so once it took off and floated above the desktop stably, we were pretty amazed. And we could also move it in X, Y and Z and roll, pitch and yaw over limited angles. So one kind of funny story that I tell people. We thought it was cool because we weren't aware of any other device like that anywhere; and it used Lorentz forces, rather than the usual Maxwell forces that are used in magnetic bearings; so the Lorentz force being the force on a coil in a fixed magnetic field. So the part that floated, which we'd called the flotor – f-l-o-t-o-r; an analogy with the rotor of a motor – that contained six coils which interacted with magnetic fields produced by fixed permanent magnets. The sensing of the position and orientation then was done with light emitting diodes attached to the moving part; and then position sensing photodiodes based on the lateral effect on the stator. So we were able to not only stably levitate it but move it in all those degrees of freedom. So we thought: Well okay, so now it's time to show the big boss, the boss of the whole, I don't know, couple hundred people, the big department, Manufacturing Research Department, to show him that the thing worked and so forth. So we did. We got on his calendar – of course he was a busy man. And he came to the lab eventually on the appointed day and time.

Just before he came, Tim was worried about everything 'cause it was just a spaghetti mass of wire and everything and amplifiers and computers and barely hanging together. And he said: "Well, you know, maybe I should reboot the computer to make sure." And he rebooted the computer; which was the fatal mistake. So the big boss – I don't even remember his name – he came in. We pushed the button and it was supposed to levitate up off the desktop; and instead it just didn't do anything, it just sat there. And he leaned over, kept leaning over, looking at it more and more; and then smoke came up. So he jerked his head back, looked at his watch and said, "That's very interesting"; and turned around and he walked out. And that was – he never

did see it ever work. That was our big chance. It was probably a career-limiting demo. But it turned out it was a software – kind of a hardware/software problem that caused the device to go with maximum force downward instead of upward.

But anyway we got that fixed and things worked after that. So the next thing with that levitation was we built a model that could work on a robot – that we could actually attach to the end of a robot and do all these amazing things we had done before with fine positioners. So again we could do micron level things. But now we could do it with six degrees of freedom. That's really important for things like putting a peg in a hole where you need additional degrees of freedom to let the hole peg slide into the hole, in case the robot isn't perfectly aligned and so on. So as it turned out, that worked really well; and it was like magic actually. It could do all kinds of things. It would act as a force-torque sensor. We could vary the stiffness of it. We could emulate different mechanisms with it that would formerly be required to go to the shop and actually build a specific mechanism for a specific compliance; we could do it all by software. So we tried to think of a name; and we finally came up with the name Magic Wrist. So that became the IBM Magic Wrist. And we proved with a lot of experiments that we could do lots of assembly and test kind of operations, even with a robot that wasn't accurate, and do them faster and with less force than you would do with using just say a robot, even with a force-torque sensor on it. So it became pretty well known I guess.

Another – a little story is that I think it was 1987, I was invited to go to the International Symposium for Robotics Research in Santa Cruz, California. And it's a very prestigious conference; especially back in those days. There weren't very many people who were invited. And I gave – chose to give a paper on the Magic Wrist. That was before the Magic Wrist actually operated. So I was – the deadline came up, you have to go, you have to give the talk, which was mainly on the design. But we hadn't shown that it would work. And there was one famous roboticist there who asked a question: Wouldn't it be easier if you did three degrees of freedom or four degrees of freedom? And I tried to argue: No, it's actually simpler to do all six degrees of freedom because then you eliminate all mechanics and so forth. He was not real convinced of that.

Q: Who was it?

Ralph Hollis: Bernie Roth. And then as I was – and my recollection here is a bit hazy – but as I left the podium and – there was, by the way at this conference the tradition at the time was to give an award or something, sort of a little award to the technology that was deemed to be the most promising technology. And well there were a lot of really good talks; and it turns out Vic Scheinman got the award that year for another project called RobotWorld. So as I was leaving the session I guess, I heard one person say, after my talk, that maybe it should get an award for the project least likely to succeed, the technology least likely to succeed. So that was a funny

thing. But anyway. We did get it to work though; and it worked really, really well. So it became pretty well known. But we only built four of these; and we never quite applied them to actual applications in IBM.

Q: Why did you not apply them?

Ralph Hollis: Well I don't know if you know anything about technology transfer. But either – I mean, we did transfer a number of things. But if you're a little bit too early it's not good; if you're too late it's not good at all. If it's the wrong color or it doesn't have a line cord, it doesn't have a manual, doesn't have an 800 number that somebody can call when it breaks in the middle of the night, then it doesn't work. Okay? So we never developed it into a hardened – what we would call a turnkey system that somebody that didn't know how to program could push a button and make it all work. That would've taken additional resources and so on, and additional time and money and what have you; as is the case with many things. Remember, I was at IBM Research; but we worked a lot with the development labs and the manufacturing labs and different places in IBM. So anyway.

Well anyway so eventually I became manager of Advanced Robotics. Russ Taylor moved on to a different position. And so I managed a group there; and we did all kinds of applications and sensing. And we became interested in manufacturing based on planar motors; which there was a thing called a Sawyer motor, which was invented in the 1960s by Bruce Sawyer. I actually met him once. And we thought they would have a lot of benefit to manufacturing. So we started working with that sort of thing.

Q: What made those motors different than previous motors?

Ralph Hollis: They could move in X and Y directions; and not just in a linear direction, and they would direct drive with a single moving part. And Sawyer had licensed things to a company called Xynetics or Zynetics – I don't know how you say it – and they built plotting machines that would do drawing using these motors, which could travel in X and Y. So you could hook a pen to it or something and you could draw. That company is long gone. That same exact motor design became the prime mover in Scheinman's RobotWorld work, for which he won that best paper, I guess; yeah. As it turns out, we were already working on our own version of RobotWorld at the time; as well as a spinoff from Bell Labs was working on that same thing – two spinoffs, two different spinoffs from Bell Labs. One was a company called MegaMation. They were all working on that; and we were too. So what we decided to do – these were all these planar motors let's call them – they were stepper motors that they could move in X and Y – am I too detailed here?

Q: No, no, no.

Ralph Hollis: Okay. I just have no idea. Okay, okay. They were open loop stepper motors; which means the computer would give pulses to the motor. It would move; but the computer had no way of knowing actually if it actually did move and how far it moved and so on. Plus the resolution was really poor and there was no disturbance rejection; like if you pushed on it, it moved. And the computer never – didn't know that. So I started focusing on making it a closed loop system; so sensing its position on the platen surface that it was flying on and it used an air bearing to fly above that. And people in my group – Jehuda Ish-Shalom, Dennis Manzer, and I; and I had engineers and technicians – we started working on that pretty hard. And we had lots of ideas and we got some interesting results, and so on.

About that time, IBM started going downhill, I think – starting about 1990. And in '92 I guess it was, IBM decided to reduce its workforce in the research division by about 400 people. So I think there were about 3200 people in research in Yorktown Heights and San Jose, California, and Zurich, Switzerland at the time, and maybe Tokyo research lab. So they offered a buyout to everyone. No matter who, you could – it was voluntary, but it was a buyout. And I had no intention at all of leaving, actually. <chuckling> I was perfectly happy, and I had a nice group, and I had a postdoc coming in from France that would spend time with me, and we had customers in our various divisions that were counting on us developing solutions for manufacturing, and so forth. Lived eight miles from work, I drive through the woods, see nature – I mean, it was fantastic. IBM Research Yorktown Heights building is located out in the country and it's this fabulous, this beautiful design by Eero Saarinen, the famous architect. It was a fabulous place. Fabulous.

But it was going through hard times, especially for researchers. And so I told my wife, "Well, I don't know, maybe I should leave." <chuckling> She said, "What?" <laughs> And send my résumé out, and so forth. And so I did. I sent it out to a few places, and I was invited here to Carnegie Mellon on a Thursday. I was met for breakfast. I gave interviews with different people, gave them my talk, and it ended up like ten thirty at night. I liked what I had seen. I had been to Carnegie Mellon a couple of times before giving talks. Went back home on Friday morning, and I think it was Saturday morning a FedEx truck pulled up in front of our house with a written offer. So I went, "Hmm, okay." <laughs>

Interviewer: Who did you meet while you were here?

Ralph Hollis: Oh, Raj Reddy, Takeo Kanade – a lot of the faculty here. I don't know exactly. Matt Mason. So, well, I had that sort of in my pocket, but it wasn't until May that I gave them the final decision that I was going to come here. And then I had – I spent until October – actually November – at IBM, trying to tie up all the loose ends and place people at different

places. Some people retired, some people quit, and some people I moved to different places, and all of that kind of stuff. Also – it's a long story – but I also managed to pack up about 80 percent of the lab and bring it here to Carnegie Mellon. That's what you see here – lots of equipment and instruments and so forth.

Interviewer: So IBM was divesting themselves of the robotics direction at that point?

Ralph Hollis: That was – as it turned out, that was the case, because I was it at that point, basically. Russ Taylor was still at IBM, but he was 11 miles away in the Hawthorne division. He was focusing on medical robotics, which they also decided they didn't need either. You have to understand, IBM got out of so many different things. Like it does almost nothing now but hold customers' hands. Probably have to edit that out, but anyway.

Interviewer: If you want, we can edit that out.

Ralph Hollis: Partly because they got rid of – they did, it turns out, get rid of about 400 scientists, almost all of who went to universities. They did save a little money, and they did emphasize technology much less, and that's probably why they made it through the dot-com bust and have been extremely successful, and I'm glad they are successful. <laughs> For various reasons.

Interviewer: Before that all happened, what kind of applications were you – I mean, you were looking at these manufacturing applications, but were they really for manufacturing various products that IBM was doing, or were they thinking about also selling this technology to other companies for other kinds of manufacturing?

Ralph Hollis: The fine positioner, we considered selling it to other companies to use, and that didn't seem to work out for some reason. I don't know why. But all the other technologies was for internal use, and it –

Interviewer: Like hard drive manufacturing?

Ralph Hollis: Hard drive – for the technologies that we had, circuit board manufacturing, ceramic substrate manufacturing, hard drive manufacturing, printer manufacturing. All the things they don't even do now at all. And you have to understand, dial back to those days, IBM developed most of its manufacturing equipment in-house. It was a very vertically integrated company. Now it's a lot more Sears-Roebuck-ish, because there are hundreds of companies that

produce manufacturing equipment that IBM can just buy. But that didn't exist back in the day, so. So kind of what we were doing became obsolete in that respect as well. So, yeah.

Interviewer: Have you worked with IBM since then?

Ralph Hollis: No, I really haven't. I still have friends back at IBM, and – but I haven't had the opportunity to actually – I've been pretty busy doing other stuff, actually. <chuckles>

Interviewer: Did you bring any of the people from there?

Ralph Hollis: No, no, just some of the equipment in the lab, yeah. Which we've used a lot.

Interviewer: What was the first project you started when you got here?

Ralph Hollis: Well, so I got here at Carnegie Mellon in 1993. By the way, when I was back at IBM, I continued to work on the robot Newt, and doing a lot of things with that. Got it to the point where it could navigate around the basement workshop, and had an onboard gripper, and used ultrasonic sensing to sense walls. I could do a routine like in the corner of a room where it could sense the walls and precisely locate its position and orientation. It could pick up wooden blocks and build simple structures at that point. I was very enamored with a number of PhD theses that were written out of MIT, I guess, and Stanford. Scott Fahlman's work and his system called BUILD was very influential for me. Other work – lots of work at Stanford. Some work at Carnegie Mellon. I came here a couple of times to give talks, and Jim Crowley was a name that – who built mobile robots – Hans Moravec here; Patrick Muir, who I later hired. So I immediately connected with the mobile robot work that was starting here at Carnegie Mellon, but it was in a very crude state compared to what it is now.

So, well, I came here 17 years ago – <chuckles> – and time has gone by really fast. And I don't know what I've done exactly. <chuckles> But I've loved it here. It's a tremendous place, and absolutely wonderful. I think they like me, so that's good. So we raised our two boys here, who are now off in the world, so we're empty-nesters. My wife Beth has been extremely supportive to all my endeavors all these years. So – well, anyway, so getting back to your question about what here – well, so I came here. I had the beginning ideas of how to make closed loop planar motors and possibly make it into some kind of manufacturing system. I also had ideas about how to use the Magic Wrist technology to enable a person to interact with a computer through the sense of touch – haptics. And that whole notion was started with Tim Salcudean as well at IBM, so. We kind of co-invented, I guess, the notion of using magnetic levitation for haptics. About a year or two before I came to Carnegie Mellon, Tim went to

University of British Columbia, where he's a professor now doing many things, a lot of it in medical robotics. We're good friends and keep in touch.

So I had those two basic kind of notions what to do with magnetic levitation, what to do with this planar idea – can we really make it into a <inaudible>, because if we could, we could dramatically increase the performance on it. So, as it turned out, I fortuitously got two brand-new PhD students at the moment I came here. I don't know whether their arms were twisted to work with me as a startup thing or what, but. I also, as I mentioned, brought all this equipment from IBM, from my lab at IBM. Not all the equipment at IBM, but a lot of it. So at the time, right when I moved, the Robotics Institute was taking over this building, Smith Hall, and remodeling it. So I thought, "Wow, that's a good opportunity." So my brother happens to be an architect. He designed the lab we're sitting in. Sent a bunch of plans to me, and I handed them to Carnegie Mellon, and Carnegie Mellon built it like I wanted it. So that was nice. Finally, was able to ship the equipment from IBM. It took three 18-wheel tractor trailers coming from New York to Pittsburgh – 65 crates, 40 thousand pounds of stuff, which my two new graduate students had to learn how to unpack – <laughs> – and plug in, and make work. <chuckles> So they did.

So Peter Berkelman was one of the students, who decided to work in the haptics area, and Arthur Quaid worked in the planar motor area. And I got another student a year later, Zack Butler, who joined us, and so on. Some master students, lots of undergraduates, and so on. So Peter and I took the magnetic levitation technology, we remodeled it into a haptic device, and it worked quite well, and so that was his thesis. Arthur did create the closed-loop planar motor that we were looking for, and it was terrific, and he showed that it performed much better than anything out there. And we started to develop this planar manufacturing system called Minifactory, and I was able to get a large amount of funding – several million dollars from the National Science Foundation – to work on the Minifactory. And I've been working on it – I'm still working on it – after 14 years, and I just have a new grant starting in January, which if it's – we've already created – I don't know, we've graduated a number of students and, I don't know, written like 48 publications, and I think 40 people have worked on the system now. But we still believe that it really could revolutionize manufacturing, especially precision assembly, in the U.S.

Interviewer: What are some of the ideas behind the Minifactory?

Ralph Hollis: Oh, wow.

Interviewer: So just summarize those 48 papers.

Ralph Hollis: <laughs> Well, so, first of all, we had some goals for the Minifactory. The first goal was to reduce the time it takes people to design and deploy, or program and deploy, an automated assembly system for assembling a product of some sort – like imagine a cell phone or a small sensor, a medical device or something like that – a disk drive. There are automated assembly systems out there that build these, but all of those are designed and the assembly systems themselves are pretty much designed and built by hand, and that takes a long time. And a long time translates into a lot of money. So a typical rule of thumb is that if you put all the robots in, and the cost of all the robots – robots make it flexible – multiply by three and you got the total cost. So you have to add all kinds of stuff. You have to add the vision; you have to add the communication; you have to add sensors; you have to add a conveyance for the product; blah, blah, blah, blah, blah. So it takes a long time to deploy something. So our first goal is to reduce maybe the time it takes to deploy an automated assembly system from four months to two weeks. So that's the number one goal. And the question is: What set of hardware and software technologies need to be invented to reduce that time down to a week or two? That was the big question. We didn't exactly know, but we had some ideas of what needed to be invented software-wise and hardware-wise.

The second thing is we wanted to increase precision over traditional robots. About 90-something percent of the precision assembly systems in the world are based on the SCARA robot, the four-degree freedom robot, which has improved enormously over the years but still is not micron-level kind of accuracy. So we want to increase precision anywhere in this Minifactory to micron level or below, anywhere, and do it sort of as a natural consequence of the architecture, which we were able to do, because of the success of the closed-loop planar motor and many other things. We wanted to also reduce the floor space that the thing takes up, and so we're about one-sixth the floor space of the conventional system. So that's kind of a lot of hardware, and truly there is a lot of hardware, and then a lot of software had to be developed as well. So we decided to make the system extremely modular as the key to rapid deployment. So each of the robots in the system we call an agent. It has software that makes it self-describing and self-representing. And the whole system sort of plugs together like Lego blocks in a way that requires no central controller and no central database. So each of the agents talks to their neighboring agents and exchanges information, and so forth, and the product gets assembled. So, so far we've done partial assembly of small things like microphones, like small optical devices and that sort of thing. We don't have enough agents built yet, and we don't have quite the right software yet to build full products, which may require, say, 25 robots or 30 robots. We have enough infrastructure in the factory to do that, but we lack the funding and wherewithal to make that happen at the moment. So that's kind of the status of that.

Interviewer: Where does most of your funding come from?

Ralph Hollis: Most of my funding comes from National Science Foundation. And we're trying to demonstrate scientific principles rather than do engineering, but all of the – all of my projects

involve building hardware. And then what you don't see is all the software and algorithms that go on top of it. In my lab we don't really do fundamental theory, I would say, and I we don't do software only kind of projects. But we always have a hardware component.

Interviewer: Do you collaborate with anyone here?

Ralph Hollis: I collaborate a lot with many different people. On the Minifactory, I collaborated with Mark Kryder in Electrical and Computer Engineering, and with Satyanarayanan in Computer Science, and so forth. And other project areas I collaborate with other people. But in the Minifactory, we're pretty poised to move forward now, and I've had – for the past year I've had two students working on it. They're visiting students from Germany and from China. We're making good progress there. So in my lab we tend to develop what I would call fundamental robotic technology. We actually don't do applications, until just very recently; we're starting to do a little bit of applications. But we typically don't respond to RFPs that say we need such-and-such and then we deliver it. Almost everything in this lab comes straight out of my head, like just because I want to do it. And then of course I have to go <chuckles> make the sale to somebody. Okay. So we don't really do applications and we don't really do theory, but we do what I call fundamental robotic technology, actuation sensing control – real-time control. So those are the things that we I think excel in as a group.

So they – I call it agent-based micro-manufacturing, this Minifactory project. It's a whole – it's not really a project. I call it a program, and I have many subprojects within it. In fact, in the lab we have three programs. Agent-based micro-manufacturing is one. Haptics and teleoperation is another. And dynamically stable mobile robots is the third area, which here's an example, here's another example. So each one of these is what I call a program, and then it has multiple funding sources. And then within that, there are individual projects to demonstrate this or demonstrate that, or graduate this student or graduate that student. That's how the lab kind of works. So in the area of haptics, we – as I mentioned, Peter and I developed this prototype. We also – we collaborated with David Baraff, who is now – for many years now – at Pixar. And when you see *Toy Story*, and so forth, at the end he's on the credits rolling by. He's done very well at Pixar, drives a good car, has a nice house. So. <chuckles> So, but we're able to interface everything to a three-dimensional virtual environment where you can feel the force and the torque in your hand, and so forth.

So the fundamental advantage of that technology is that it's – there is no mechanics. All the mechanical issues are thrown out the window. So conventional haptic devices are basically back-driven robot arms that the user grabs the end of it and moves it around, and that puts a position and orientation into the computer; then the computer feeds back forces and torque..

<break in recording>

– that the editor wanted me to write as an example of how a small company – a spinoff company worked. Because we have a company to do that now. And then there's a Harvard University Press book coming out with a chapter on the magnetic levitation haptics. So lots of stuff is happening there. <chuckles>

Interviewer: So even though the big boss didn't get to see it –

Ralph Hollis: <laughs> Right.

Interviewer: –it's working out.

Ralph Hollis: It's working out. Yeah. So you're rolling now.

Interviewer: Yeah, we're back up. You were talking about the magnetic-levitation haptics.

Ralph Hollis: Yeah, the advantage of the magnetic levitation is that it doesn't have any mechanics. So kind of like graphics – if you have a dirty screen or low-resolution graphics and so forth, you don't see the real thing, really. You need really lots of pixels and so forth. Same thing is true in haptics. You have some machinery between you and the virtual – the three-dimensional and virtual environment. And if that machinery causes friction or backlash in the joint or inertia effects or cogging in the motors, you tend to feel that sometimes, rather than feel what the computer is intending you to feel. With the magnetic levitation, that goes away, and you essentially get a pure connection through magnetic fields between what the computer is doing and what you're doing. So we actually believe it is the most direct connection you can have between a running computer program and a person's hand, other than go straight into the nerves, which is fairly invasive. Probably not going to be a bit hit. So we're real bullish on it.

And a few years back I got a major research instrumentation grant from NSF to further refine the technology, and we did that, and we built ten of them and we sent them out to Stanford and Purdue and Harvard and lots of places. We have some here as well. And they work great, and people like them, and so on. So a little less than three years ago, we spun off a company which is called Butterfly Haptics, LLC, and my wife runs the company, and we have a network of vendors building parts for us and two local companies doing final assembly, so. My wife and I sort of just run the company. She does all the paperwork and stuff, and I'm sort of a consultant, actually. And that's working out well. We're selling systems in U.S. and Canada. We're expanding to Japan soon, and Europe. So that's kind of a side effort, and that's the subject of the article, I guess, in *Robotics & Automation Magazine* that the editor asked me to write. I guess he's starting a new series on spinoff companies or something, that do robotics.

Interviewer: So are they all for research – sold for research purposes?

Ralph Hollis: Yeah. They're pretty expensive. They're 48 thousand dollars each, which is competitive, actually, with other devices that are out there, other mechanical devices. But we have about 25 times the performance of some of the most popular ones, which is a big jump in performance. So what you feel is exquisite, <chuckles>, let's say. And a lot of people don't understand that until they actually try it. So because of that, my wife and I, we go to a lot of conferences. We've been to eight so far I think, where we set up a booth, we put up our banners, and we have our equipment there and we run demos for days to hundreds or thousands of people. And most of the people like it, and we get orders, and we sell them. So it's been a lot of fun for both of us, so. That's a side effort. We're still producing – we're still doing haptics here, but we're no longer developing the product; we're not developing the hardware, really. We're trying to develop a few little applications, actually. So, try our hand at doing something real.

Interviewer: What are they?

Ralph Hollis: So, one is we have a U.S. Navy contract to do bomb disposal robotics, so working with a local company here – RE2, it's called. So we've interfaced one of our Butterfly Haptics systems to their robot, their bomb disposal robot, and you can dig in the sand and you can feel the sand grains and you can feel if it hits a solid object. You're going to get an immediate response in your hand, and imagine that that's several hundred yards away – or miles away, for that matter – and it's kind of a nonvisual thing, because it's buried under the sand. You don't really see it until you hit it, and then if you're going to hit it too hard, it's going to blow up. Maybe that's not good. So, also we're developing a grasping capability, which is soon to be demonstrated, which will allow the operator of the robot to remotely pick a wire or a blasting cap and pull it out without squeezing it too hard and crushing it, which would blow it up.

The objective is to be able to disassemble some of these improvised explosive devices to find out how they're made, and so forth, so. That's one application. I have a PhD student that's doing work related to that, as well as doing that. So, and then I have a postdoc who's doing biopsy needle insertion – so pushing a simulated needle into your belly and going through the skin and the fat and the muscle, and each one of those interfaces you feel a little bit different sensation. If you feel a fourth thing, that's not good, because that's your intestine. That's not good. <chuckles> So, that's going on along really well, and doctors are saying it feels really good, so that's another sort of application. I had a young lady here, a PhD student from China, until just last month doing haptic dental training, so scraping and picking teeth, finding cavities or caries, dental caries, put the probe in the artificial – the virtual mouth – and if the probe sinks in just a little bit instead of scraping across the tooth and it sticks, then – when you pull it out it kind of sticks a little bit – it means it's a decayed area of the tooth, and then you can go in and drill using a Voxel-based approach, different geometric approach, and so on. So that's another

application. We have a lot of interest in developing simulators for eye surgery and ear surgery, those two things. But there's a big step going from the hardware and software that we have to a turnkey application again – it's like the Magic Wrist. So there's a huge amount of software that has to be developed related to the physical modeling and the graphical display. So we're seeking partners to try to bridge that gap because we can't do it all ourselves here with my students. But we're doing a few things – taking a few steps in that direction, I would say.

Interviewer: Are there any other researchers looking at maglev for this kind of haptic?

Ralph Hollis: Well, so Tim Salcudean, my co-investigator back at IBM, also pursued magnetic levitation haptics at University of British Columbia for several years and did some very interesting things there with it. He's no longer doing that, partly because you can just buy it now. <chuckles> But there is one of our systems at University of British Columbia, but it's not in his lab, I don't think. So yeah, that's the only other known case.

Interviewer: Everything else is mechanical?

Ralph Hollis: Yeah. Yeah. There's about maybe five or six companies in the world that are producing systems.

Interviewer: We interviewed Ken Salisbury as well. He's been doing haptics but also mechanical systems.

Ralph Hollis: Yeah, Ken started, well, at MIT. His student has the spinoff company, SensAble Technologies, and I don't think Ken is really long gone from that, having a relationship with him maybe – but yeah. So he developed all these incredible cable-driven devices, robots and haptic devices and so on. It's interesting; he gave a talk in 1993 – again, it was one of these International Symposium for Robotics Research. Same place I – same time I gave it on the Magic Wrist, actually, and the Magic Wrist use of the haptic device. So that was before we actually developed a real haptic device, but temporarily we used the Magic Wrist as a haptic device. So he's turned that into a company and now my idea also turned out to be a company. SensAble Technologies is the world leader in haptic devices, I think. Immersion always says they are but – probably they are in terms of money, but. And SensAble sells these Phantom devices, but we have 25 times the performance that I've found, so we're hoping –

Interviewer: I've tried the Phantom before, but –

Ralph Hollis: You can try ours if you want, if you have time.

Interviewer: I'm just kidding.

Ralph Hollis: Sure. Absolutely. Yeah. So as a company, Butterfly Haptics is – we're selling it on the performance. We view it as a general-purpose scientific instrument that can be interfaced to almost anything. Back at IBM, we actually interfaced the Magic Wrist to a scanning tunneling microscope and we were the first to feel atomic scale surfaces in real time with your hand by moving the scanning tunneling microscope tip over a surface. You could feel the bumps, which were due to small clusters of atoms. We were the first to do that. That was with Tim Salcudean as well. Also, the same technology that we developed from the Magic Wrist and the haptics can be used for vibration isolation. I tried to sell it to NASA for – I went to about five different NASA places and tried to sell it, and couldn't do it. Tim managed to convince the Canadian Space Agency, and Bjarni Tryggvason, a Canadian astronaut, to pursue this Lorentz levitation, we call it – Lorentz force levitation – as a vibration system in space. And so they did, and it flew for three years in the Russian Mir space station, and it also flew on the STS-85 shuttle mission, and it was very successful. A third version is being prepared for the ISS. It hasn't been launched yet, and who knows if it ever will be. But the idea is to isolate a payload from vibrational disturbances in space. On the space station and on the shuttle there's a lot of vibration going on. Fans are running. Thrusters are – astronauts are pushing off of walls and all this stuff. So it's really a microgravity environment, but it's not very good unless it's isolated, which we can do with this technology, so. Another spinoff. We're out of time, I guess.

Interviewer: Almost out of time. We were just going to ask one last –

Ralph Hollis: But we haven't gotten to dynamically-stable mobile robots, and this is the very first one.

Interviewer: Oh no! Let's go there.

Ralph Hollis: So, a few years back, I had the idea of replacing the three or four wheels that a conventional wheeled robot has with a robot that has only a single wheel, and that single wheel would be a sphere, and that's what this is. So this is the world's first one here, called Ballbot. It is – you can – when the three legs come up, it can balance and it can move about from Point A to Point B very nicely. It senses gravity using fiber optic gyros and MEMS accelerometers, and so forth. We're adding a pair of arms to the robot. People around the world are building Ballbots, actually – ETH in Switzerland and also in Japan. This is a different version of Ballbot over here. It was developed by my colleague in Japan, Maasaki Kumagai, there. We have licensed the technology and there will be commercial versions available in a year or two probably.

Interviewer: Do you have a lot of international collaborations?

Ralph Hollis: So, not so many international collaborations, but I do have a lot of students coming from Switzerland and Germany mainly, and some from China. Some from Holland, from Netherlands, and some from France. So. Yeah.

Interviewer: What's the big advantage of this kind of mobility?

Ralph Hollis: So a Ballbot like this allows, with only one wheel, the robot can be tall and skinny. It also can move in any direction, so it's omnidirectional. So something – another balancing machine, like a Segway, has two wheels, but it must turn first before it can move. So a Ballbot like this can be in your cluttered apartment building or something and it's almost impossible to have it be trapped where it can't move, because it can move in any direction. It also – because we only have one wheel, we can make that wheel be rather large. That's eight inches in diameter down here. A large wheel means we can easily go over carpets, changes of height, go in and out of elevators with gaps – no problem. So that's another advantage. Also, depending on how it's controlled, it can be very, very gentle. You can guide it around by hand, you can push it, and you can do all of this without having any external sensors. Although we do have some external sensors, and we're adding more, and we're adding two arms to it. But.

Interviewer: So it's a good potential kind of everyday life personal robot type platform.

Ralph Hollis: Yeah, exactly. So. It has some other – with arms, it can – we haven't demonstrated it yet. Here are some of the arm mechanisms here. It can potentially lift large loads by leaning, much larger than a typical robot can do. It can operate on ramps and slopes that would tip over a regular robot. Most robots in the lab are about this high, which is fine. But if you really want it to be person-size and you want it to interact with people, HRI – you want to look them in the eye and everything – you're going to have a considerable amount of mass up high, which means high center of gravity, means you either move slowly or you tip over. This robot can move a meter and a half per second and do that very gracefully, in fact. So I think there are a lot of advantages.

Interviewer: Are they going to be available as research platforms?

Ralph Hollis: We don't have a plan for that, although we have shared the design with many people, various tips and so forth. We don't actually have the drawings online but we could, I guess.

Interviewer: What kind of advice do you have for young people who might be interested in a career in robotics?

Ralph Hollis: Stay out of it. <chuckles> Well, robotics is, for my life, has come from a hobby to what you might think of as a recognized field, even though still it's not maybe a completely recognized field. People seem to know what robots are nowadays, and that's good if you want to go into that as a career. In the old days, everybody just laughed when you said robots, because it meant science fiction. People don't seem to laugh so much anymore because you have robot airplanes, you have robot surgeons. You have all kinds of robots that are doing things for you. I mean, half your car nowadays is a robot. So it's a good time. I think if you want to do robotics research, that's great. It combines several different fields into one, if that's your cup of tea. If you don't necessarily want to be an expert in physics or an expert in electrical engineering or software, but you want to be a jack of all trades, robots integrate all those fields into one, and increasingly also the human element. Which, by the way, I forgot to mention one long-time collaborator, Roberta Klatzky, who is a psychologist, who – and that's another – we've done continuously a number of psychophysics experiments with her, magnetic levitation haptics, where we've measured human ability to sense textures, to sense deformable objects, and so on and so forth. So, yeah. Probably a lot more – right? – that I missed. But those are the highlights maybe. <chuckles> The highlight reel. I don't know.

Interviewer: That's great. Thank you.