

Pasadena, CA
1:19:26
m4v

Brian Wilcox

Selma Šabanović
with
Peter Asaro

June 2 2011

Q: So if we could just start with where you were born.

Brian Wilcox: Okay. I was born at Ridgecrest, California, which is right outside the China Lake naval weapons development area where my father was an engineer, and so I grew up among all the – essentially all the kids in my neighborhood, all their fathers were weapons developers for the Navy, developing things like the sidewinder missile and other similar kinds of weapon systems. So we were all kind of brought up with high-tech rockets and things like that.

Q: And did you have any interest in technology as a child?

Brian Wilcox: Oh, yes, yes. I was very interested in all things having to do with technology. My father actually had a Ph.D. in physics, and so I was always kind of brought up with physics. He taught my brother and my sister and myself elementary physics from the beginning. I can't really remember before we would have family discussions about the laws of physics and things like that, so I was always fascinated by that.

Q: Did you build radio sets?

Brian Wilcox: Well I did, later I built a crystal radio set, but what we built a lot of were rockets. All the kids on the street, all the boys on the street, anyway, built little rockets, and that was before you could buy these rocket kits at the hobby store. So we actually built them ourselves and it was quite dangerous. I knew people who got badly injured. And my cousins who also lived there, they had friends who were killed, so it was dangerous. But I knew a guy who blew off several fingers from his hand.

Q: But yet the fascination continued. <laughs>

Brian Wilcox: Danger, yes. Well, yes, I mean in those days if you had to build your own from scratch, and kids of course didn't know what they were doing. When I look back on it now it's all, you know, how could I have possibly been so stupid, but at the time it seemed all the kids did things that were terribly dangerous, but most of us survived.

Q: And so did you go to high school there as well or how long did you stay?

Brian Wilcox: No, I lived in China Lake only until second grade, and then I moved to Washington, D.C., where my father had been transferred to the Pentagon. And we lived there only one year and then moved to Detroit, Michigan, where he had taken a job with General

Motors. And General Motors was forming up a new division to respond to defense solicitations, so my father had been asked to set up a division of General Motors to work in the defense area, and that was eventually located in Santa Barbara, California, near Santa Barbara, a place called Goleta, and so they formed that division in Goleta, California, which is where I lived then. I lived in Santa Barbara, which is the adjacent community, the much larger adjacent community. So I lived there from fifth grade through college actually, in fact, through my first jobs.

Q: And what did you study in college?

Brian Wilcox: I was a physics major, and I ultimately got a bachelors also in math, so I ended up with a Bachelor of Science in physics and a Bachelor of Arts in mathematics at the University of California, Santa Barbara. So I started out as a physics major and then I realized about three-quarters of the way through that I had taken so many math classes that with just a couple of more careful selections I could get a math degree also, so I did that.

Q: And were you involved in research at all when you were an undergraduate?

Brian Wilcox: I would not say research. I was involved with teaching. There was a summer course that was taught at Santa Barbara for math teachers from high school and junior colleges, and so I had come to the attention of one of the professors and he asked me to basically be a teaching assistant. So I was a teaching assistant for these series of math colloquiums in the summertime. And so I would conduct a session where we'd solve problems and grade papers and things like that.

Q: And so how did you decide to go to graduate school and what to actually do in grad school?

Brian Wilcox: Well, initially I went to grad school at Berkeley in physics. I had applied to Caltech and to Berkeley and gotten into both, which were two of the best schools in California, two of the best schools in the world, really, and the third being Stanford. I wanted to stay in California. But I got into Berkeley and so I went to Berkeley. But then I realized about halfway through the first year that really I was an engineer. I'm not <laughs> a physicist, and it's humorous because my father had told me that I was an engineer and not a physicist, which I never believed it because I had loved physics so much. But when I got into grad school the physics was not about things that are practical, right? In other words, as an undergraduate physics major everything you learn is tremendously practical in the real world, but when you go to graduate school in physics it's all about quarks and other things that really don't have very much practical application to real life. And so I realized that I was really much more interested in building things that worked. I had always been very handy as a kid and I had a lot of tools and I did a lot of projects. And so I realized that I didn't really want to spend eight years, which was

the typical time to get a Ph.D. from Berkeley in physics, and I didn't really want to spend eight years of my life – and at that time also there were no real jobs available because all the universities had filled up their physics departments in the fifties and sixties with a fresh crop of post-war graduates, and these were all young professors who were not going to die anytime soon, and so most of the departments were full, and so there weren't a lot of jobs in physics. So I basically dropped out at that time and I went back to work at a little company that I had worked at during college, and there I met other guys and we started a company. So within a year or so after dropping out of grad school or out of the Ph.D. program in physics at Berkeley I had started a company, and it was a little personal computer company. So in 1975 the first personal computers were coming on the market. There was a thing called the Altair in January of 1975. The Altair was on the cover of Popular Electronics magazine, and this was a kit you could buy. And so a friend and I who worked at this little company bought one and built it. And so we got this thing going, and one of the other guys who worked at this same little company, he was an electronics whiz and he said, "Oh, you- you bought this thing? You actually own one?" And we said, "Yeah, we- we actually bought it and own one," because we were more actually into software. And I mean I had done quite a bit of electronics, but I had never done design of computers as such. But he was very interested in the design of computers, and so the three of us formed a company to make products that would go with this computer. So that computer called the Altair was known basically for the Altair bus, and the bus had the property that you could plug a card into it from any vendor. It was a so-called open architecture meaning that it was an open published standard and there were no proprietary limitations, so you could design and build a card that plugged in. And so we started making products, and our first product was kind of a flop, but our second product was very successful, and so we were off and running. Within a very short period of time we had moved out of our garage. When we formed the company we rented a condo and started the company literally in a garage, and so we built our first prototypes in a garage. But within a few months –

Q: What was the product?

Brian Wilcox: Well, the company was called Polymorphic Systems. The first product was an analog interface board because we thought, well, hobbyists are going to want to have analog sensors and analog outputs from their hobby computer. Well, it turned out that product was a total flop. I think we probably sold 100 of them or something over the life of the product. But the second product was a so-called video terminal interface, and what it was, was a card that you could plug into the backplane, the Altair bus, and you could connect it to a TV set and a keyboard and now you didn't need to buy a terminal. Right at that time – people who are used to a monitor on computers have no idea what it was like in those days. A computer was a completely separate thing from a terminal. The terminal was a way to input and output data from a computer, and in those days what we had were teletypes. And so my first computer experience was with a teletype, the ASR-33, which goes clankety-clankety-clank, and it communicates over a phone line, and that was a typical way you'd use it, at 110 bits per second, which is 10 characters a second, which is so unbelievably slow that no one can really imagine what it was

like in those days for this thing to just go clack, clack, you know, pound out characters on a roll of yellow paper. It was very primitive, uppercase only, no mixed upper and lowercase. It was very primitive, just like the old, basically like the old telegraphs. I mean it looked like a tele – you know, when you got your listing it literally looked like a telegram. So, anyway, this idea of a so-called glass teletype had been out for a while where they had taken essentially a TV screen and a keyboard and said, “Okay, well, we’re gonna do what the teletype does except instead of pounding it out on paper with a ribbon we’re gonna- we’re gonna display it on a TV screen.” And so people had started making back in the early seventies so-called glass teletypes, which were just terminals. They were serial terminals that would plug into an RS-232 port, but they still functionally they did the same thing as the teletype. What we did instead is we had a block of memory that was memory mapped to the computer, so it was in the main address space of the computer memory, and now when the computer wrote an ASCII character into that memory it would show up on your screen because our board periodically swept through that 1,000 bites of memory and displayed 16 lines of 64 characters. So that kind of revolutionized computers. Now I’m not saying we alone did this. We didn’t have a patent on it or anything, and many people had that same basic idea, but we actually had the first personal computer that had a memory mapped video display, and every computer now has a memory mapped video display. In fact, it’s almost inconceivable to go back to the old serial terminal model of computing. And we had a bunch of other firsts. We were the first one amazingly enough to have a boot ROM. It actually had a read-only memory instead of toggling in on the switches on the front panel, which is what all of us with those computers, when you first turn the power on every day if you turn it off at night in the morning you turn on the power and now you have to toggle in with binary switches enough bites into memory to get it to read a device reader like a paper tape. In those days we used paper tape readers, and they had a little spool of paper that you’d run through a tape reader, and they had punched holes in them. And as the punched holes went past an optical sensor it would feed in. But we had to load manually something like 16 or 20 bites of program. Amazing today that you would have a program that’s only 16 bites long, but when you had to do it the hard way you wrote very efficient code. And so with a 16-bite program it would just go out to this reader and load the paper tape into memory, and now you could actually do something useful. Now at that same time Bill Gates was writing MITS Basic, which the Altair was made by a company called MITS, and Bill Gates and his partners there at Harvard had written MITS Basic, and so we of course immediately got MITS Basic and said we can do this better. We got another company to write a – they had already written, in fact – we bought from them what we thought was a much better Basic that we sold as part of our product. So we gradually built up to 50 people, 50 full-time people and several million dollars a year in volume. But about that time the whole industry was moving from kits that are sold mail order, which is how we got started, to fully assembled units sold through computer stores, so we started selling through computer stores. The first computer store in the world was here in Los Angeles. It was the computer store down in Santa Monica. It’s no longer there. So within a year or so there were 1,000 stores worldwide selling computers. And our product was sold in many hundreds, you know, 800, 900. The problem is that there was a big shakeout in the industry in the ’77, ’78 timeframe and half the stores went out of business. Unfortunately they went out of business owing us money, and so we lost something like half a million dollars when all these stores went out of business, so we

had to go look for investors. So we were negative cash, negative net worth at that point. So we went out looking for investors and eventually found a group who basically bought the company so that we went from one-third owners each to five percent owners each. And sad for us or maybe it's reassuring, I don't know, but they ran it into the ground even worse than we had been doing, so eventually it was worth nothing. The company went out of business eventually 10 years later, although it took them a long time to completely run it into the ground. But at that time Microsoft emerged and Apple emerged, and so we had kind of met those guys. In fact, Jobs and Wozniak had invited us to dinner one night at one of the big conferences, but we had a plane to catch because we were flying out that night, so the three founders declined to go to dinner with ultimately what ended up as being the big winners. But, anyway, can't play it over again. I wish I'd gone to that dinner.

Q: What was the name of the company that you had all worked in before?

Brian Wilcox: Oh, it was called Minicars, Incorporated [later MCR, Inc.]. It was kind of an interesting little company that was making hybrid – in fact, my first job was with this company in high school. And in 1968 I got a job with this company, and my first job was to work on building a little hybrid car. And the guy who founded it claims to have invented the hybrid car, and I can't deny that maybe he did because obviously he had had this idea a number of years earlier. But the idea was to make very little hybrid cars that were shared used vehicles, meaning that you bring them into a lot and you park them anywhere and you basically turn in your key effectively. And then anybody who's a member of the system, they take the first car available because all the cars are identical. So you don't need to look for your car because you don't have a car. You just take the first car, which means they're all dense packed. There's no space between them in the parking lot. And so basically in the morning down in the central city you just fill the lot dense packed. You all pull up. It's just like returning in a rental car. You pull up as far as you can, then you get out and you go away. And then at night you take the first car available and you take it home and it sits in your driveway all night, which was a great idea. I thought it was a wonderful idea. I think he had a patent on it, but it never went anywhere. It was just one of those millions of ideas that were before its time, I think. I mean just like hybrids. We were building hybrid cars in 1968 and hybrids, when did they become popular? Five years ago, eight years ago.

Q: How did you first get involved with JPL?

Brian Wilcox: Okay, well, so I worked with Polymorphic Systems. We brought in this investor group in late 1979 and then I worked there for a while and I joined JPL in the summer of '82. So in 1982 I was hired by Carl Ruoff, who I think you're going to talk to, and he had been running a robotics group here at JPL, which had recently shrunk actually quite a bit because during the mid-seventies they had been working on what was called the Mars 84 Lander. Now just to

review the history there NASA had spent a tremendous amount of money on the Viking space probes to Mars, which were launched in late 1975, no 1976, and arrived the summer of 1977. And so July, August of 1977 they landed on the surface of Mars. And they had spent billions of dollars. These were by far the most expensive non-human missions ever built at that time and probably to this day. They were tremendously expensive because they wanted to answer the question, “Is there life on Mars?” And there had been a wave of enthusiasm on this topic, which when those spacecraft, ultimately it was determined that they had not found life on Mars many people felt that that had killed exploration of Mars for 20 years. And in fact it was 20 years before the U.S. – or any country, I think, but the U.S. in particular – had any successful missions to Mars. Now we had an unsuccessful mission called Mars Observer. The Russians had some unsuccessful missions I think as well, but it was from the time, from 1976 to 1996 when we launched Pathfinder, that was a 20-year gap in the exploration of Mars. But at the time in the mid-seventies there was this enthusiasm over Mars, and they hadn’t come crashing down – the realization that Viking had not discovered life on Mars had not hit yet, and so there was still this enthusiasm about Mars. And that led to a thing called the Mars 84 Project. Nineteen eighty-four was a particularly good launch opportunity for earth to Mars transit. Because Mars’ orbit is eccentric, every year – well, there’s a launch opportunity to Mars every 26 months, but some of them are better than others. Some of them are actually much better than others, and they repeat on about a 17-year cycle. So 1984 was an especially good year where you could land a lot of payload with a relatively small rocket. So that was the focus of a lot of activity here at JPL during the seventies. So they had built a little robot. Actually they built two robots, the so-called software prototype and the so-called hardware prototype. The software prototype, they connected it through a serial phone line, literally a leased phone line to the biggest computer down at Caltech at the time, and they tried to make this thing autonomous. Well, this was really the first realization of how very difficult autonomy is. If you remember that era people were saying, “Oh, well, the chess – you know, the computers are gonna beat the best human players at chess by the end of the sixties,” right? There was a big wager over that question. Well, of course it took 30 more years before the best computers beat the best humans at chess, and that’s only because they built a piece of special purpose hardware. Really, the first time that a general-purpose computer could have beaten the best chess players is about now because of course no one has actually attempted to do that. But it’s only now that general-purpose computers have the performance that the special purpose computer, Deep Blue, had back in the nineties that was built only for one purpose, and that was to play chess. So people were learning how very difficult intelligence was. And of course now in hindsight it’s obvious that if you look at a neuron a single neuron has about the same computing performance as the early microprocessors. So the first computer that I worked with when we were making personal computers back in the mid-seventies was a so-called 8080 processor. It was the first eight-bit processor, and the 8080 had about the same performance as one neuron, so in terms of the ability to take inputs and process them and do an output it was roughly in the same computing capacity as a single neuron. Well, we have 10 to the 10 neurons, so 10 billion times the performance of a single microprocessor. Now even with Moore’s Law with it going tenfold every five years, roughly speaking, when you have 10 orders of magnitude that’s a lot, so that’s a 50-year – you know, it takes more than 50 years. See, when I got going in the computer business the clock speeds were

one to two megahertz was a typical clock rate. Now clock rates are gigahertz, a few gigahertz, so we picked up a factor of a thousand. Also the on-chip complexity is quite a bit greater, so we picked up factors of 10 or maybe even 100. But we kind of hit the wall on both clock speed and individual complexity, so now all we can do is replicate chips. Well, that shouldn't be a problem because the brain is highly parallel, so if we can figure out how to architect it properly that should be no problem. We just go forward and in another 30 or 40 years we'll have computers as capable as the human brain. But why people were thinking that the computers at that time would approach human performance was just crazy talk. I mean it was just crazy talk because the best computers of that era were not that much faster than the first microprocessors. The first microprocessors were reasonably capable little devices. And, yes, you could buy an IBM 360 that was faster but it wasn't that many orders of magnitude faster. So if people had just done the back of the envelope calculation they would have realized that a person is – that reaching human level performance is another significant fraction of 100 – oh, you've got your own power strip there.

Brian Wilcox: Okay, so anyway, people should have realized but didn't that human performance was a long way away. It's still probably 30 years away from where we stand today, but at that time it was a long, long way in the future. So at the time basically a computer represented the performance of essentially one neuron. Today your laptop represents maybe what a worm has in terms of computing capacity. You know, you don't think of it that way but if you ask yourself, well, what kind of performance should I expect out of something with a single laptop-style processor you shouldn't expect more than what a worm can do. And then we'll gradually go through cockroach and a housefly and so on. And in the last 10 years we'll go through all the furry little animals that one thinks about when one thinks of biology, rats and cats and foxes and so on. But that will go very fast and then of course it'll just pass through human performance like we're standing still, which we are. And within five and ten years after that we'll see a situation where the computers are orders of magnitude smarter than humans, and that's a very scary, scary thought. I won't live to see it. You guys might.

Q: To go back to the Mars 84, so there was that window of opportunity, but this was late seventies?

Brian Wilcox: Late seventies, right, uh-huh.

Q: Or '82 I guess you joined the project.

Brian Wilcox: Yes, I joined, and that project had folded up its tent, and they in fact had laid off most of the people. So I came at a time when they were trying to rebuild the robotics community here at JPL. They had shrunk down to Carl [Ruoff] and just a few other people, maybe two other people who were really seriously engaged in robotics here, maybe a few more than that. But

basically they had lost their funding. Of course, this was before I got here so I'm speculating a little bit, but they had lost their funding because of the collapse in interest in Mars with the realization that Viking had spent all this money and had not found any evidence of life. Now that is kind of an interesting story in itself because the main life detection experiment was to put Mars soil into a nutrient bath and see if it evolved carbon dioxide, which would indicate that it was metabolizing the nutrients. Well, what happened is they put it in and it foamed over. It just bubbled over almost instantaneously, which they interpreted to mean that there was something inorganic in the soil that was causing a chemical reaction, which they later concluded rightly or wrongly that it was a hyperoxidant introduced, that basically the ultraviolet light from the sun was attaching extra oxygen atoms to things making them kind of like hydrogen peroxide. So when you pour hydrogen peroxide on a wound it kind of fizzes up, and this is the same idea. When you put the soil in it's like adding hydrogen peroxide to the nutrients and it kind of fizzed up. However, the principle investigator of that experiment said, "Well, wait a minute, guys. We agreed on beforehand what the success criteria was, and this met the success criteria, so you guys should not be asserting that we did not discover life. It's still an open question. You know, it met the pre-established success criteria of the experiment and that you shouldn't jump to the conclusion that there's no life on Mars," but everybody did, anyway, so his protestations were not heard by any significant fraction of the people. And what ended up is now we believe that, yes, there is this super oxidizing layer on the surface of Mars, but underneath we have no idea whether there could be life just below the surface or not. And certainly if you go down a ways there's almost certainly liquid water, and in that liquid water there certainly could be life because if life evolved at the surface and then as the planet lost its atmosphere, you know, it got colder and colder and that life would then go underground. Just as on earth there is a tremendous amount of life that has more or less recently been discovered down in the rocks underground. It's in the soil. In the broken rock crevices down underground there's a tremendous amount of biomass living underground on earth, and this was not known, certainly not known at that time, but now that fact that life exists on earth suggests that it's very likely that there could be life on Mars. So long-term exploration of Mars would focus on trying to establish if there is life, and many people believe that the conclusion of Viking that there is no life on Mars is not a valid conclusion, that there could well be life underground and Viking would not have tested that question either way.

Q: So what was the first project you worked on?

Brian Wilcox: Well, the first project I worked on was a Mars rover. It was called Mars Rover Sample Return. It was a project, which was supposed to launch in 1996. So here I was a brand new, fresh out person, just arrived in 1982 and I was told to go work on a thing that would launch in 1996. Now that was a shock, and that was a typical shock that people have at JPL, which is they don't realize that the life cycle of a mission can be very long. And if you're in at the beginning of a mission from the time until that mission is over – I knew people that lived through the entire mission cycle of Galileo, which was a big spacecraft that went to Jupiter, and they worked on it from the early to mid-seventies through the end of the nineties, and they

worked on every mission phase from the pre-planning phase all the way until they finally crashed it into Jupiter after 10 years in orbit or something like that. So these missions can go on a length of time comparable to the length of a career. And so here I was, a fresh-faced young fellow, and I was told, “We’re gonna launch this –” you know, if everything goes perfectly we were going to launch in 1996. Now of course we did launch a rover in 1996. It wasn’t exactly that rover, so I got to see the rover grow and shrink and change function and so on. But basically we worked on what was essentially the same rover for 14 years and eventually did launch it. It became what was known as Sojourner. And Sojourner, by that time I had become the group supervisor and so the group that we had, more than half of the people in the group were very intimately involved in Sojourner, and in fact our group was responsible for the software, the electronics, everything except the mechanism itself, which was done in the mechanical section. So I’m in basically the electronics and control part of JPL, and there’s another section that does mechanisms, and so between basically our group and one other group over in the mechanisms area we were kind of responsible for the whole Sojourner flight system.

Q: And so in the time between ’82 and ’96 how did the project develop, and how did the ideas of what the robot would be like or what it was going to do change?

Brian Wilcox: Well, initially it was part of this large mission called Mars Rover Sample Return, and the idea there was that you would send a rocket to the surface – you’d land a rocket on the surface of Mars that could carry samples back to orbit, which would then be picked up and return to earth. That is still a long-term vision for Mars exploration. In fact, the recent Decadal Survey identified that as the next step in Mars exploration, and we either do it sometime in the next decade, the decade starting 2020, or we probably don’t do it at all until we can afford it because there really is no other high priority exploration of Mars left to be done. All the in situ measurements that are valuable have been done. And some people might quarrel with the Decadal Survey, but that’s now the policy of NASA, basically. So at that time we recognized that we were getting close to where we really needed the samples back on earth because you can analyze those samples with so much greater integrity. You know, we have room-sized instruments that can analyze just a relatively few atoms of material, and you can’t do that with an instrument you send to Mars. So this was a large rover. It was bigger than this rover behind me. It was the size of a car, not that different, actually, from the size of the rover that’s going to launch to Mars in just a few months. But at that time it was considered very ambitious, so they wanted to land two of these ascent vehicles, two rovers and have two giant cameras, Hubble class cameras in orbit mapping the surface at the scale of one-foot resolution, which would give you kind of the God’s eye view that was thought to be needed at that time to navigate the rovers around. And so the thinking was that if you had these great cameras, one-meter diameter mirrors in orbit, optical mirrors, that you could map the surface so accurately that you could plan many days of rover actions. And despite the speed of light roundtrip to Mars, which is up to about 40 minutes, you could have that rover continue to function usefully on the surface essentially all through the daylight hours. And so then it would pick up samples. It would put them in the rocket and the samples would be launched back to earth. Well, as we went through the planning

of that mission the whole system grew and grew and grew. It didn't start out with one-meter diameter cameras in orbit. It started out with 30 centimeter – but as the requirements, as always happens it seems within NASA people are very risk averse, and so they say, “Well, you know, but that might not work, so if- you know, to make sure that it will work you have to do this extra thing. You know, you have to take the camera, and instead of it being 30 centimeters it has to be half a meter,” and then somebody else comes along and says, “Oh, but that might not work so it better be a meter.” So pretty soon you've got these humongous things. So the rovers grew from initially at about 200 kilograms each to 1,000 kilograms each. The orbiting cameras grew from 30 centimeters to a meter. Well, it just became unaffordable. So finally they assigned an independent, non-NASA group, SAIC, to do an independent cost review. And their independent cost review came back and said that the total mission cost was about \$13 billion. And \$13 billion was four or five times what the Viking mission had cost, which was two landers and two orbiters. And so it was so much greater than the most expensive mission that had ever been done before that NASA headquarters just pulled the plug and said, "No, we're not going to do that." So that mission concept was cancelled in 1989 roughly, and within a very short period of time came the Mars Environmental Survey which led to Pathfinder. So that concept was, well wait. We're not going to try to do sample return; we're not going to try to put the big cameras in orbit, because we'd had two big orbiters during the Viking era that had mapped the whole surface at the resolution of a football field which is not bad. And so the concept of putting a little robot on that, a little rover was developed, and that concept was very controversial. Now it seems, again, it's one of those things where looking back on it, it seems like a no-brainer, but the idea of putting a little rover on a Mars lander was very controversial. And the reason it was controversial is that scientists see that mass as mass that could have been their instrument developed in their laboratory. So they don't want to give up mass that's going to Mars because in a planetary scientist career you might get one opportunity to send an instrument developed in a university laboratory to Mars. If you're lucky you get one; if you're really lucky you might get two during your entire career. And so scientists don't want anything competing for that real estate, that mass, power, you know, space on board the spacecraft. And so whenever rovers had been proposed before, there was this buzz saw of criticism from the scientific community that said no, no, no, no, no. We don't need a rover. That makes no sense. We have instruments that we've already have been through peer review and they lost last time but now they've come to the top and so we're kind of cycling through the good ideas. And so that was kind of the conventional wisdom within the science community. So Pathfinder was not a science mission. Most people don't realize, but Pathfinder was a technology development mission and the rover on Pathfinder was not even paid for by the science directorate. It was paid for by the technology directorate as a science experiment, as a technology experiment rather. And so it's only by that technique that the fact of – that the idea of a little rover could be floated because it wasn't coming out of the budget of the science directorate and, as a result, the scientists really didn't have any say over whether this was going to fly or not. So eventually, in hindsight of course, we now realize that Sojourner was able to do things that couldn't be done on Viking. Viking landed these two hugely expensive landers that had arms that could reach out, and there were rocks everywhere. In both lander images, Viking 1 and Viking 2, both had just field of rocks as far as the eye could see, but neither arm could reach a single rock. And therefore all the hard rock

science couldn't be done. Now they didn't really have instruments to do hard rock science, but they tried to reach rocks and the ones that they thought they could reach turned out to be just little ball dirt clods. And so it turns out they couldn't reach even a single rock to bring that rock in to analyze it with the instruments that they had. And about the time of Sojourner, there was a guy named Mike Carr, and Mike Carr was a leading scientist. He was – he's long since retired now, but he was leading the Mars science working groups. Many of the Mars science working groups were led by Mike Carr. And Mike started making an impassioned plea to any audience that would listen, saying the history of Mars is in the rocks that are on the surface. The rocks that are on the surface were dredged up by impacts and so they're from deep underground. Now, we don't know exactly where. It would be better if we did, but the fact is that we've got all these rocks spread all over the surface that are from deep underground, and once we get through the weathering rind, you know, once we get through the first few millimeters probably, we'll have the history of Mars captured in those rocks. And if we can get at those rocks, we can determine the history of Mars; we can answer all the major scientific questions about Mars. And so that together with the fact that Viking had landed with long arms, two-meter arms, but they couldn't reach any rocks, you know, led people to realize that maybe having a little mobility. You know, we've gone 150 million miles; maybe we should go the last few meters to actually reach the rocks instead of being so tantalizingly close. And so Sojourner was grudgingly accepted when it became obvious after the mission, technology demonstration mission, paid for by the technology directorate, became obvious that that was in fact the right way to explore the surface. And that led, of course, to the Mars exploration rovers, which are true science missions in the sense that they were built by the Science Mission Directorate and they were Mission Success on Pathfinder. A huge part of Mission Success was just landing safely. Well, in a typical science mission such as Mars Exploration Rover, you get zero credit for landing safely. You have achieved none of – no part of Mission Success just by landing safely, despite the fact that that's the hard part, right? Getting to Mars and landing safely is 95 percent of the engineering challenge and yet you get no credit for it in a science mission because you haven't accomplished any science at that point; whereas Pathfinder, you know, a huge part of Mission Success was just landing safely. So if you remember the videos where everybody's cheering on CNN, you know, you see the whole control room erupt in cheers, that's always true, but it's especially true on Pathfinder because that represented mission success for them and so their job was largely done. Now for us, we are, of course, taking bets on whether the lander would actually work, those of us on the rover team. And we were pleasantly surprised that it did work. We weren't that sure that it was going to work because it was very complicated – it had something like 140 consecutive pyrotechnic events that all had to go perfectly and if any one of them failed the mission would fail, you know, releasing the backshell [of the aeroshell], releasing the parachute, releasing the heat shield; there was a huge number of events all of which had to go perfectly. So when the Sojourner landed, you had that great picture that was on the cover of Time Magazine of the Rover sitting there with Mars in the background and that was great. And now they said okay your show; let's make it go. So most of the people in my group were doing mission ops, so I was on the downlink team. Jack over there was on the uplink team. Jack wrote almost all the software on Sojourner, by the way. So anyway, the fact is that the first time we tried to communicate with it, we could send short packets and they'd get through fine, but all the long

command packets were failing. And so here we were scratching our heads going ew, we've got a real problem with this rover and so we had our radio guy and so on. And we decided – I mean, to make a long story short, the antenna had not been turned up vertical yet, so it was still horizontal, but still it was so close that it should have worked. It certainly worked in all our ground testing. But because of the very cold temperature, we were right at the edge; all the crystals and so on are somewhat temperature sensitive, so we were right at the edge of the band in which the receiver was sensitive to the transmitter. So the transmitter and the receiver were drifting off frequency from each other and we were right on the hairy edge where each one had kind of a square reception curve and those two rectangles were almost non-overlapping at that point. And that coupled with the fact that the antenna was at right angles instead of aligned was causing it not to receive. So anyway, we put together a tiger team and realized that we could get around it. Even if this problem remained permanent, we had a fix for correcting the occasional bit errors, 'cause we were only getting one-bit error in every long packet, so the short packets were coming through with zero bit errors. But the long packets were coming in with typically one-bit error beyond what the error correcting codes will allow. And so we crossed our fingers and we sent the command to turn up the thing and that basically solved the problem. But we had something like 36 hours there where we were very worried because we had not yet – the lander wasn't really ready for us to move off yet, but we knew we were having these problems communicating with the rover. And the media had gotten wind of it, so there was all these stories in the media about oh the rover's got problems and maybe it's not going to work after all. But it did and of course we were able to drive off one of the ramps, as you may recall. One of the ramps did not go all the way to the ground, so it deployed. It deployed kind of like a party favor that just comes out, unrolls like a tube, but it just hung straight and in the low gravity it didn't break it's little – where the membrane was supposed to just pop just like a carpenter's tape when you pull it out too long. It supposed to just [clicking sound]. It didn't do that, so we weren't really too anxious to go in that direction, so we were limited to just the one direction, but fortunately it was free – it was clear of obstacles in that direction and we were able to go off. So we moved off and were able to then explore lots of area. The vehicle was – the first autonomous vehicle on the surface of another planet and it's autonomous in the sense that it had hazard avoidance, it has a find rock feature, and those two, which I designed and Jack coded, so I assume he followed roughly what I attempted to describe in my description of what I thought it was supposed to do. Jack could tell you exactly what's in it. But those systems, each performed flawlessly, but only once, on the surface of Mars. So only one time did it actually have to avoid a hazard. It worked perfectly. Only one time did we give it a "find rock" [command] where instead of avoiding a rock we wanted it to go up and position itself in front of the rock and get that rock, you know, and place that rock in exactly the right position so that we could put our instruments on it and so on. In both cases it worked absolutely perfectly and was never used again because the operations team was so conservative that they didn't want to use those features. They wanted to do everything manually. And this is another great lesson of autonomy and robotics, I think, which is that humans really want to be in control and so just like in the Apollo landings, all six of the successful Apollo landings, the pilot took control away from the autopilot before landing. In only one case was that actually required, which was the first one. All the others, they did it because they didn't want to tell their grandchildren that I was – I landed

exactly once on the moon, but I just let the autopilot do it. I mean, can you imagine a trained pilot letting the autopilot do the one and only landing in his lifetime that's going to be remembered and just letting the autopilot do it? So it's not, you know, it's just not human nature. Now, we have that problem in that when President Bush said we're going to go back to the moon, it was decided early on that we would go back to the polar regions of the moon where the sun is just at a very grazing incidence and there are these long, deep shadows, much deeper than anything you see on earth because there's no sky light to fill in the shadows. So it just pitch black – it's black as pitch in these craters and so even the slightest undulation it'll go from this incredibly bright sun to incredibly black pitch in almost no time at all. And yet, if you're trying to land, you know, what you have is this montage of very brightly lit surfaces that are inclined just slightly towards the sun, and then behind them pitch-blackness. Well, what you really want to do is you want to have a radar system that looks down and it doesn't care about the black areas. It sees just as well into the dark as into the light. And it's either laser based or radio based, and it maps the whole area, gets a 3D representation, does hazard detection and so on. And all that is current technology. So we know how to do those things, but we had this tremendous psychological issue. What pilot is going to let this thing land automatically even though that's the only safe way to do it? So that has become a problem for all of NASA because we're going to get to the point where humans are going to need to land in areas where they cannot land themselves safely; the computer has to land them. And yet human nature is such that they're going to want to land, they're going to want to grab the stick, turn off the autopilot, and land it with a joystick despite the fact that that's not the safe way to do it.

Q: What were some of the biggest challenges in designing Mars Lander?

Brian Wilcox: Well, basically we had, again, essentially the venerable 8080 processor. There was an 8085 which had a few extra features, but it was one of the – I mean, as typical with NASA missions, in order to get a radiation hardened piece of electronics, you'd usually go back two or three generations in electronic complexity. And so here we were in the early '90s designing with a 1976 processor and getting everything to fit in that processor was just a – making it all work, making it – testing it so that we could be confident in it. We had this sandbox where we set up all kinds of challenges for the Rover and tried to establish that in fact it would do what we wanted it to do for every challenge. And it seemed to work. And so we eventually tested it enough that we were all satisfied that it would work. But that's, in general, a problem for autonomy in the future because you get more and more complicated systems validating that they will work in every situation. The combinatorics just gets impossible, so you can't possibly know – you can't possibly test every combination. And yet, you know that on any given mission you're going to be presented with combinations that have not been tested. So how do you convince yourself that the software is, in fact, going to work? And that's especially true as we go to artificial intelligence because the only way, and this is my personal view, is that you're not going to have this explosion in software that we have seen in the past. The history of development of systems over the last 50, 60 years, is that the hardware gets somewhat more complex, but the software gets vastly more complex. And so we're getting to where tens of

millions of lines of code are in these advanced systems: Fighter planes, jumbo jets, things like that, and they're so big and complicated that you have this validation problem. Well, you look at the human brain. The human brain is based on the human genome. Well, we've now mapped the human genome; we know how big it is. It fits on one CD. Amazingly, not even one DVD. It fits on one CD. The human genome fits and what fraction of that is the human brain? I don't know. Ten, twenty percent maybe. I mean, we only differ by two percent from the chimpanzees. Well, chimpanzees are pretty smart. So two percent of what fits on a CD is the difference between us and the chimpanzee. And then maybe another ten times that much is the difference between us and a redwood tree. So everything there is to be human is something like ten or twenty percent of what fits on a CD. Well, that's not tens of millions of lines of code. It's tens of thousands of lines of code. And so whatever the human brain does, it does with tens of thousands of lines of code. And we can guess, of course, what that code is. It's learning algorithms, and it's obvious signal processing algorithms, things that have been in the textbooks probably for 50 years. So my own favorite is Bayesian Learning where you just take the past as your prior and then extrapolate the future based on the past using Bayes' theorem. And that, to me, is probably a huge fraction of what's going on in the brain, together with obvious signal processing type algorithms of the type that have been in textbooks for decades. So what that says is everything is learned. Well how do you validate a learned behavior is going to do the right thing? It's not mathematical. You can't prove it. You have to – what's the difference between Albert Einstein and some day laborer you pick up at Home Depot? How do you validate that one is going to be able to do your job and the other one isn't? You don't know. And you don't know when the system's going to go berserk. And yet, it solves the problem of the explosion in software complexity which has really no other solution. We know that the human brain is based on 10s of 1000s, probably not more than 100,000 lines of code, and yet it performs things that are vastly – have tens of millions of lines of code and that's because we bother to write all that code. We shouldn't in some grand sense. We shouldn't have written the code. What we need is a huge brute-force computer like this one [pointing to brain]. So the team, while we were fortunate enough when I arrived in 1982, there were a succession of people hired, and we were fortunate to have a tremendous number of great people joined the team so that within – when I started there was one robotics group at JPL and that group probably had certainly less than ten people in it, and that includes technicians and others. And within ten years, we had an entire robotics section, meaning that we had four or five different groups each with ten or twenty people. So we had an explosion in the number of people and fortunately we had very good people, and so we were able to do things like Sojourner, that couldn't have been done [otherwise]. And that was partly me, but because I became a supervisor in 1985 and so I was able to then start hiring people myself. But it was also we had a bunch of other people who – we gradually branched off so that one group became two and then two became four and four became eight. And so we built up to the point where we were the size of what is considered a section, which is about 100 people here at JPL. And that's a little bit my doing, but mostly other people who were able to also hire very good people.

Q: And how did – even as robotics became more important like that, how did people decide that it was okay to hire?

Brian Wilcox: Well, these things are always driven by funding. Basically funding had gone to a low [just before I got to JPL]. Robotics in general goes through a boom and bust cycle. Every time you see robotics on the cover of Time Magazine or one of the other big national magazines, you know that you've hit the boom. The bust can't be far away. So people get enthusiastic and then disillusioned and then it takes about ten years to get enthusiastic again and then they get disillusioned. The problem is that the pace of progress in computing, rapid as it is, in orders of magnitude every ten years, is so eclipsed by the performance of biological systems that people assume that robots ought to have the performance of biological systems. And of course science fiction makes this worse because you see science fiction movies with R2-D2 and C-3PO and so on. We've all been seeing these movies where these robots have tremendous capabilities that are very much like humans, and of course that's 'cause the props have humans in them. And so you don't see those kind of behaviors [in real robots], so after a few years of enthusiasm you get disillusioned and people stop funding. So we were fortunate that we're just coming out of one of these bust cycles when I got hired here and that was caused by Viking in the NASA area but also the whole – I think there was a bust just generally coming off the same cycle in industry and so on, where robotics, you know – robots were really first developed in the United States in the '60s and yet now the robotics industry is all overseas because we, in the United States, we just can't make it through the boom and bust cycle with our investment, you know, the way we do investment, the way we do management. All the MBAs come out and say well if I can't get a three-year return on investment I'm not interested, and so on; whereas in other countries, they go for the long term and so they have built a sustained – they've had that sustained attention that leads them now to dominate the robotics industry. I'm not saying research now but industrial robots are very much dominated by non-U.S. companies.

Q: What are some other robotics projects you've worked on since Pathfinder?

Brian Wilcox: Well, this one behind me is ATHLETE. This is the all terrain hex-limbed extraterrestrial explorer and the basic concept of it is that little wheels on the end of legs can do things that neither wheels or legs can do alone. And so you can roll efficiently on 97 percent of the terrain, but if you get stuck – because the wheels are kind of small, it's easy to get stuck, and so we designed the wheels for roughly 97 percent of the terrain – then it just locks the wheels, you use them as feet, and walks out of the problem. So the Mars Exploration Rovers, of course, both of them have gotten stuck at various times. One is permanently stuck now, but the other one was stuck for like five weeks and it took a long time to get out. Well, if that had been an ATHLETE style vehicle, it would have walked out the same day because it would have just realized that it was up to the top of the wheel in yogurt and it would have locked the wheels and just walked out. And so that is the basic underlying philosophy of this. You save so much mass on the small wheels. You save something like a factor of six or seven on the mass of the wheel

and wheel-drive assembly by making it so small that you pay for the rest of the leg, so you get the leg for free. And then once you have a leg you can make it into an arm too. So we put a tool adapter, so every wheel has a tool adapter that grabs a tool and has a power takeoff so that when you spin the wheel it turns the power takeoff, which is a half-inch socket drive, and that turns the tool. So you can grab a drill and drill a hole in the ground or you can grab a gripper and grab things and move them around. So you can do lots of things that you just couldn't do any other way because you now have six manipulators that give you a general-purpose capability that you just wouldn't have any other way.

Q: And do you expect it to ever launch?

Brian Wilcox: Well, we hope so. The launch of ATHLETE is probably been delayed somewhat by the decision not to go to the moon but to go to an asteroid, and then it's not clear where the moon fits in. We were designing specifically for the moon. It works well in any low gravity. The lower the gravity, the better in a sense, although when you get to such low gravity as on an asteroid it's not clear you need wheels at all. You just hop and then you go long distances by just flying. So if you can fly, there's a lot of advantages to flying. But this certainly is good. Everybody seems to believe that the moon will appear somewhere in the roadmap to Mars and of course we'd like to believe that this has good utility on Mars as well. So you almost certainly need to learn to live and work on other planetary surfaces before you go to Mars, someplace that's not as difficult as Mars to reach. If you have to do a mission abort, it's 400 days back to earth. On the moon, if you have to abort it's four days back to earth, so that's a much better place to learn. And so we're very hopeful that this will go somewhere, but right now we're retooling for asteroids, so we're basically using ATHLETE as a kit of parts to build things for exploring asteroids, and some of those versions will have wheels but many will not have wheels. Wheels are not completely stupid on asteroids because they do allow you to move very precisely to, if you see an outcrop or something that you want to study or take a sample and you want to move there accurately, wheels are a very good way to do that on an asteroid; whereas hopping, because the gravity field is so uncertain and you've got the centrifugal force field of the spinning asteroid, getting it to go exactly where you want to go by pushing off is very problematic, so wheels is a nice way to get your final approach. So some of our versions will have wheels, but mostly we'll use it as a kit of parts to study all the different ways. So we're building, here in this lab, we're building with this big, blue structure here, we're building a suspension which will have six cables that give us six degrees of freedom so we can move the thing flying as if it's in microgravity. And then we actually bought an asteroid. You might say how can you buy an asteroid? Well, we're in Hollywood and Hollywood, of course, you can buy asteroids. In fact, we had a prop house call one of our engineers and say, "I've got an asteroid left over from a production shoot, do you want it?" And we said absolutely. So for \$6,000 we're getting a 25 X 25 foot asteroid mockup that is going to go right where we're sitting here and that's what we'll fly this thing on.

Q: Just quickly. We're going to have to wrap up, but I'm curious, where do you see robotics going within NASA and within space exploration?

Brian Wilcox: Well, I think robotics has a tremendous future, especially as the computers get to the point where you can start to approach the performance of biological systems, which will start being really exciting in ten or so years. And then with that, you'll be able to have these human robot partnerships where the robots are not so slow that the astronauts are frustrated beyond belief working with them. See, when – the robots that we've had to date in space have been great, but they're very slow, and they're very slow for two reasons: One is they have almost no power and power is very expensive to get in space. But also they're very slow because they're not very smart. And if you had the smartness, then you would arrange for the power so that they could move quickly. And once you get past the safety issue, which they have to not only be smart enough to move fast, but they have to be smart enough to move fast safely in the presence of humans. So that's a level of sophistication that is way beyond industrial robots today, for example. Industrial robots now are kept in cages literally where if the human enters the cage the robot power shuts off so that the robot doesn't kill the human because the robot is not sensing everywhere and it's just moving at blinding speeds, moving parts around and so on, but if a human were in there, the human would be smacked real hard. So what is needed is ubiquitous sensing and so we see that now with machines like Robonaut, which have gone up to the Space Station. Robonaut doesn't have a sensor skin as such, but it has such delicate force sensors everywhere and redundant force sensing everywhere that if it detects an anomalous force it stops so fast that it can't hurt anybody. And yet it is strong enough to do useful human-like tasks, and so that is kind of the breakthrough that was needed for robots to be useful working in the presence of humans. And so eventually NASA has been pioneering that because NASA needs it more than anybody else, but as it is perfected you'll start to see that in the production environment on earth where robots can work side-by-side with people and not put those people at great risk.

Q: And can you tell us some of the people that you've worked with most closely, both here and outside of NASA?

Brian Wilcox: Well, sure. Carl Ruoff we mentioned originally hired me. And then Carl went off, he went back to school and so on, so he kind of disappeared from the scene for a while. We've had a number of people: Tony Bejczy, he was a major player in robotics for many years. He's now 80 years old. I saw him not too long ago in good health. He retired maybe ten years ago, but he led in teleoperation as distinct from robotics, so the ability to use human gestures and make the robot move identically with what the so-called master-slave approach to teleoperation, where you have a master controller and then the robot is the slave. So he did a lot of work in that area. And then there have been just a huge number of people: Larry Matthies in the vision area, originally came into my group and then formed his own group, has done a tremendous job in machine vision. There's a guy named Guillermo Rodriguez who has, I'm sorry to say, has had a

stroke and passed away, but he developed a thing called spatial operator algebra which has become extremely important in not just robotics but in a broader field. Really, I mean, too many people to mention. It's like the Academy Awards. You never know when to stop, but anyway. That's a good collection of people.

Q: And you took your master's at USC?

Brian Wilcox: That's true, yes.

Q: Who did you work with there?

Brian Wilcox: That was a non-thesis program, so I had to take nine courses and so I did those. JPL has a great program, by the way, where people here can take courses at USC by television. There's a television link, direct link, which in the modern era is not impressive, but 20 years ago it was considered very exciting. And then we just had to go down there for classes, I mean, for the tests. And so it's a great way to pick up an advanced degree, but I didn't really work with people down there as such. I just took the course work.

Q: And do you have any recommendations for young people who are interested in careers in robotics?

Brian Wilcox: Well yes. It may be obvious from what I've already said, but we're now entering perhaps the most exciting period of robotics, which is kind of too bad because I'm getting old, but for the young people it's really the most exciting time because we're now – Moore's Law is now to the point where you're going to see this explosion from worm to cockroach to housefly to ultimately cat and rat and human in over the course of one lifetime, over the course of one professional career. So if you're just starting today, by the time your career is over, robots will have human-like performance. I mean, that is an amazing fact. That's my prediction. It could be wrong. But I think if you just do the back of the envelope based on the computing performance of the human and the computing performance, assuming Moore's Law continues, which I think as they – they will have to go to these highly parallel architectures because they've reached kind of the limits of clock speed. But other than that, you will see this performance. I think parallel performance is possible. The human brain is an existence proof that highly parallel systems, even operating at low clock speeds – you know, the human brain is like 100-Hertz clock. That makes it possible to do things that are just of amazing capability. And so I see a synthesis between electronics and biology as the biologists learn more and more and we try to put those systems and, as I say, I think the software is going to be trivial. So some people are daunted by the software explosion. I am not because I believe that the software is going to turn out to be trivial.

Q: Thank you.

Brian Wilcox: Okay.