

Pasadena, California
38:12
.mp4

Richard Murray

An interview conducted by
Peter Asaro and Selma Sabanovic

6 June 2011

Q: Yeah, if you could just start by telling us where you were born, where you grew up and went to school.

Richard Murray: Sure, I'm Richard Murray and I grew up in El Paso, Texas. All through high school went to schools in El Paso, Texas and then came to CalTech as an undergraduate in 1981. Was here until 198 then went to Berkeley as a graduate student and came back to CalTech in 1991.

Q: So what were you studying in CalTech in undergrad?

Richard Murray: I was in electrical engineering as an undergraduate, more or less getting slightly interested in control systems. Not doing anything in robotics at the time and when I went up to Berkeley, started to get involved in robotics and grasping which was, at the time and now again, something that's kind of interesting -- and then still was doing the control systems side of robotics for my graduate work. Then when I came back down here, I actually switched departments to mechanical engineering at CalTech and then since then, had done robotics things and other things that don't have anything to do with robotics.

Q: So, in Berkeley, how did you get involved with robotics?

Richard Murray: I remember my advisor was Shockra [ph?] Zastry [ph?] and I think he was interested in robotics or various students doing robotics things in the lab. It was kinda the time of manufacturing robotics and grasping was something new that people weren't-- you know, I think the, you know, Utah/MIT hand or the Salisbury hand sort of existed at that time and I think we thought it was interesting from a dynamics and control point of view. So a couple of us in the lab said, "Oh, we should build a hand to do things, where other people were doing things that were a little bit more theoretically robotics. But we thought it was a good control theory problem and had some interesting aspects to it. So I think that was sort of how we got started. Went to robotics conferences, things like that.

Q: So you actually wanted to build your own hand or were you using the...

Richard Murray: <coughs> We did-- no, we built a l-- it was a little plainer manipulator that we called Stix [ph?], that there's a picture of in the robotics book that I co-authored, and, you know, just to have it in there 'cause, pretty much, it did nothing else. I don't even know that we ever wrote any papers that used it-- but anyway we built one, I think, because we wanted to kinda control directly the forces and other things and I think we also didn't have a lotta money and so it was sort of interesting. Got to learn how to use a machine shop, all that good stuff.

Q: So who was working with you on this project?

Richard Murray: Chris Beaster [ph?] helped build it. So Chris is now a faculty member up at UC-Berkeley. He works in Mems [ph?] now. Doesn't really do so much in robotics, although, I guess, at a micro scale-- and it was the two of us that built Stix and were doing things. And then Zexiang Li was a more senior student. He's now at HKUST and was sort of working on some of the types of algorithms that we would eventually implement. Ping Hsu was around. He now is at San Jose State, although, he might've moved actually. And that was, kind of I think, the group of us that were working. Brad Payden was around. He's up at UC-Santa Barbara now and was doing some things on kinematics and different types of representation. So it was a small group that Chris and I were sort of building up the robot, trying to get it to work.

Q: What kinds of things did you do with the robot?

Richard Murray: At that time, we were just trying to understand how you might grasp something and move it around. Manipulate the forces, measure the forces, maintain the forces, so that you could keep a grip on something. We were interested in things like rolling fingertips and other things, but at least in the work that we did on this little plainer robot, never really explored that. It was sorta much simpler than that. So I think we were just trying to understand, you know, what are some of the hard parts of the problem, easy parts of the problem. We were very interested in the mechanics of what was going on and the kinematics of what was going on. So not just sort of grabbing something and holding it but manipulating it inside the finger grasp, things like that.

Q: Is this what went on to be your thesis?

Richard Murray: My thesis had two parts. So one part was that-- which was kind of grasping and manipulation-- and the other part had to do with non-holonomic motion planning, which had some relationship to rolling fingertips and other things, but was a little bit more associated with mobile robots and some of the kinematics of mobile robots and some of the special things that happened there.

Q: Mm-hm.. and how was it working with Professor Zastry?

Richard Murray: It was interesting. He's a very high-energy guy. He had a big group at the time and so we were all, you know, sort of working on lots of different things. So it was quite interesting. It was before he had started taking on various leadership positions like he did later and so he was around a lot and, you know, sort of running a lab that was sort of jointly run with

Jatendra [ph?] Malik [ph?] and John [Canny [ph?]] and Ron Fearing [ph?]. We were all in kinda one laboratory together, which was a lotta fun.

Q: What were some of the other kinds of projects that were happening around you?

Richard Murray: There-- it's been a while. I remember there was-- we had this thing that I actually spent a little bit of time working on also that was called-- I wanna call it Robot World. It was something that we got from Hewlett Packard. They didn't build it but they were using it. It was a big, flat steel table that had these robots that stuck on it and it used some sort of a drive-- Sawyer Effect Drive or something like that-- anyway, that moved these robots around in the plain but then they could sort of dip down and pick things up. So it was used for assembly and manufacturing and things like that. John Houser was there when we were working on-- he and I did a little bit of work on something called the Acrobot, which was a kinda very dynamic robot that would-- sort of almost like a simplified version of a gymnast or something on a barre. There were-- you know, there was an adept arm that was there that wasn't really getting used for anything. I'm sure somebody had been using it at some point. There was work going on in an adaptive control at that time but, I think, more theoretical stuff. There had been experimental work. I think around the time I was there was when the robotics lab at Berkeley, at the time, was kinda getting built up and so, you know, there weren't a lot of robots that were getting used prior to that that I remember. But by the time we left there were several that were around: Robot World and the adept and Acrobot and the little hand which was just really for us.

Q: Did you at all overlap with Rugina Baichi [ph?]?

Richard Murray: Baichi? No, she came quite a while after I left, so.

Q: You mentioned that you went to a number of robotics conferences. What were some of the big conferences at the time?

Richard Murray: Well, so ICRA-- International Conference of Robotics Animation-- was the one that was sort of the conference that we went to as students and I remember going to one in-- it was Arizona or New Mexico or someplace that was quite stimulating. It was right when I was just starting. I think Chris Beaster and I were there together actually. So we went to that. Of course, went to the controls conferences which had quite a bit of robotics content going on. I remember it must've been later because it would've been, probably back when I was a <inaudible> but I remember going to one of the early wafers [ph?] which was a interesting workshop on algorithmic foundations of robotics. But those are the ones I remember. I never went to IRAS [ph?]. You know, I think at that time, ICRA was sorta the place that everybody went. It was the annual conference that everybody went to.

Q: So, how did you come back to CalTech after your Ph.D.?

Richard Murray: Was finished getting my Ph.D., sort of took six months as an instructor at Berkeley after my Ph.D.. Was looking for academic positions. Had a couple of opportunities. CalTech was one of them. Decided to come here, so came here in mechanical engineering. Joel Briddick [ph?] is here. He was doing robotics and was also interested in things related to grasping, manipulation, locomotion and things like that. So it was a good fit.

Q: So what did the robotics area look like in CalTech at the time?

Richard Murray: <clears throat> There wasn't a whole lot. So Joel was here. He was-- he had also an adept robot sitting down in his lab but was a little bit more interested in hyper-redundant robots, so robots with lots and lots of articulated links. Snake-like robots, things like that. Don't remember much else that was going on at the time. When I got here, I continued a little bit of interest in grasping but that died away fairly quickly. Started looking more at locomotion robotics and built some things that were non-holonomic robotics that were driving around, just sort of wheeled robots. But we also built some swimming robots at that time. Tried to understand, kinda from a theoretical point of view, what locomotion systems would look like and what's the mathematics to describe that. And so that was very much done jointly with Joel and some of his students. Gerry Marsden was at Berkeley and then came down to CalTech some number of years later. I was also involved in that work. So there was a group of us thinking about how do robots move through the environment. What are the mathematics associated with that? I also started working on things that are robots but they often get handled sort of outside the robotics community, per se, which were UAVs. Uninhabited, unmanned aerial vehicles doing work with the air force and looking at aggressive trajectories and other things that, again, had overlap with robotics but less of the decision making and other part-- and sensory perception aspects and more the just detailed dynamics of those systems.

Q: What are some of the big challenges for locomotion and non-holonomic <inaudible>?

Richard Murray: Yeah, I mean, I think at the time, we were just trying to figure out what was the right mathematical framework for describing it all. So we knew that there were interesting things you could do. So if you looked at parallel parking cars and other things and tried to ask, "Well, what are the algorithms that are gonna do that?," it turned out that many of the sort of normal assumptions broke and so <inaudible> that had fewer degrees of freedom than degrees of motion. So you might have a car that has sort of two actuated degrees of freedom. You can turn the steering wheel and you can sort of drive forward and backward but you're more or less able to control three things: the position and orientation and the plain and, although that's not particularly complicated, at the time it wasn't something that the algorithms that were out there quite dealt with. And so we sorta found it very interesting to think about that more generally.

What happens when you have under-actuated robots and you have to use some sort of notion of a geometric phase or other types of ideas where you produce cyclic motion in some set of variables. So you might drive back and forth and turn your steering wheel left and right, but if you do that in the right way, you don't come back to where you started. You come to a different position. So if you think about how you parallel park, you sort of drive back and forth once while you turn your steering wheel back and forth twice and that turns out to generate a sideways motion. So we were very interested in these sorts of non-holonomic maneuvers and what could be done with them and, at the time, that really wasn't well understood and it turned out that there were some aspects of that that if you wanted to stabilize a non-holonomic robot around an equilibrium point around some position, it turned out that all of the theory that existed for how to do that broke. And so we had to really develop new theory and a number of people around the world were kind of interested in that. So it was kind of a fun time to sort of think about, you know, holonomy and what that meant and what it meant for robotics and other things. I don't know that in the grand scheme of things it mattered very much but it was very interesting at the time.

Q: Were there other people that you collaborated with around the country or around the world?

Richard Murray: Yeah, there were. So in the area of these non-holonomic robots, which was what I worked on sort of in the first five years that I was here at CalTech or so, there were a number of people around the country who were doing that Zexiang Li was still interested in that. He was at **Currant [ph?]** at the time, before he moved to HKUST. There was a group in France that was working on this, using some techniques called differential flatness that were very interesting too and so **Michele [ph?] Fleece [ph?]** was one. The-- it-- there was a collection of people who were thinking about stabilization and control theory. **Rudolf Sepulchre [ph?]**, for example, was thinking about this. There was kind of an interesting group around the world, a lot in France for some reason. I think that Michele and some people around him, you know, kinda got interested in the problem.

Q: What did you start working on after that?

Richard Murray: It moved from thinking about these non-holonomic robots to moving into unmanned aerial vehicles and unmanned autonomous systems of various sorts, focusing a little bit more on trajectory generation for vehicles that have complicated dynamics. I think that was what we were sort of thinking. The non-holonomic work became a lot about, "Well, how do you generate a trajectory to move from one configuration to another configuration?," and we sort of took some of those ideas-- and differential flatness was one of them-- and said, "All right, well how would we do this on things that are more dynamic. So the difference in some sense between an aircraft and something sitting on the ground is that a mobile robot sitting on the ground, you

can always stop and wait and just decide what to do and then do it and stop in the middle or slow down. With an aircraft, the dynamics are such that, of course, you can't just stop. You're, you know, you have to keep flying forward. You have to keep moving. The dynamics can be extremely non-linear. And so we started thinking about how to generate very aggressive trajectories that really pushed up against the envelope of what the actuation could do and what the **non-linearities [ph?]** were in the system, and looking at how to generate trajectories very quickly. And so these ideas that were coming from differential flatness and trajectory generation and combining those into control algorithms for receding horizon control were kinda the next batch of things that we worked on. Parallel with that, completely out of robotics, we started looking at control of jet engines and aircraft engines. Just, I think, because some of the dynamic issues were the same very complex dynamics-- and how do you control those in an interesting way?

Q: Did these interests come from wanting to work on UAVs or starting to work on UAVs or did the work with UAVs come from the more scientific interests?

Richard Murray: Well, I think that what happened was that-- I mean, this was back in a time when UAVs didn't really exist so-- I mean, drones existed and things like that but UAVs like we think of them today-- you know, sort of Global Hawk and Predator and the things we use for border control and things like that-- weren't really something that were common. And so a number of us were convinced that that was part of the future of where things would go, that there was no reason to have a pilot and an aircraft, and what could be done in terms of the things that a pilot would normally do. And so I think we just sorta saw it as-- because we were involved with looking at dynamics and trajectories and control aspects and were looking for, "Well, where are the next set of problems gonna come from?," we were also very interested in sort of mechanical systems that had energy constraints and other things that had interesting kinematics or dynamics that came from mechanical principles. And so aircraft was sort of a natural sort of progression to look at, and they were very interesting problems. I think that, at the time, the air force was intrigued but it was sort of, you know, not something where-- or at least didn't feel like, from the research point of view, that everybody in the air force said, "Oh, of course we're gonna have, you know, unmanned vehicles and they're gonna have to be, you know, generating trajectories and doing all these things on their own." I think, at the time, it was like, "Well, we'll always have a pilot. That's the pilot's job. Focus on the part of the control system that goes around the pilot." And we were sort of interested at looking something beyond that. Now, of course, lots and lots of our craft and other things that don't have people in them, so.

Q: What year was it that you started working on that?

Richard Murray: Yeah, my recollection was that that probably woulda been in kinda the mid '90s that we were working on those sorts of things and, you know, I'm sure some of it started

earlier. By 2000 or so, I think we had done a lot of the work that, you know, was what we were gonna do and kinda moved on to the next set of problems, which involved more collections of vehicles, so multiple vehicles that were interacting in some way. So, I think, as the early 2000s came along, we sort of said, "Well, the single vehicle problem is sort of done. We know how to control a single vehicle, at least from the point of view of the type of research you might do in a university," and so really started to look at what would happen if you had a bunch of vehicles that had to cooperate. You start looking at the effects that occur with the communication network, that they're gonna send information across. You start thinking about how is it that they're going to interact with each other to sort of do a task that requires all of them to take the right role. So we started looking at interconnected systems of various sorts, often motivated by kinda aerial vehicles and collections of aerial vehicles <inaudible> space satellites that might have to fly in formation, things like that.

Q: Is that what you continue to work on now?

Richard Murray: No, now I do a very-- well, it's a little. So we got involved-- so the research kinda bifurcated a little bit. So on the side that we would think of as still associated with robotics we got involved 2004. So, like everybody else, with the DARPA Grand Challenge series of competitions and-- so that led to some new insights about what worked and what didn't and what some of the challenges were. And so we started thinking about a lot of things that had to do with decision making at higher levels of distraction. How do we describe control problems, still in some sense? Feedback dynamics, but ones where the descriptions were at the level of events and temporal logic and other types of descriptions. How do you drive through an intersection? You need to come up to the intersection based on what's happening at the intersection. You have to do one thing or another thing. Wait until different vehicles have gone through. Go when it's your turn. React to problems. So, you know, you start driving, somebody else starts driving, what do you do? So we started thinking about the kind of theoretical foundations of that and how you combine some of the underlying dynamics with the decision making. And so we still do a lot of interesting work there, which we think of as sort of design and control protocols. Given a specification, how do you design a control protocol? And then, separate from that, started looking at synthetic biology and design at the bio-circuits level. So how do you program DNA to implement given functions, instill that kind of feedback and control problem, but now happening much more at the level of, you know, a very different set of dynamics, very stochastic. How do you put DNA into cells and implement circuits. We do a lot of work there. We still give lots of talks in the controls community about it but not in the robotics community. I mean, in principle, some of the-- you can imagine-- nanorobots that go through and look for environmental contaminants and then go and do something to try and remediate the effects of those. But I think we're a long ways away from-- you know, we're really focused on the underlying technology of how do you actually get some of the pieces put together and working, as opposed to how do you program them as a robot, as we might think of in the kinda robotics community.

Q: So tell me about your involvement in the Grand Challenge.

Richard Murray: Yeah, that was very interesting. So I was on-- ISAT, at the time-- so ISAT's the Information Science Technology Study Group that's sponsored by DARPA. It's a bunch of people in the computer science and affiliated communities and Tony Tether, who was the director of DARPA at the time I think, decided that they wanted to try and do some sort of a grand challenge. He hadn't decided it would be in robotics <coughs>. But the first I heard about it was somebody from DARPA came to ISAT and said, "Hey, you know, we've got this ability to award prizes," and "What would be an interesting one that might stimulate research of interest to DARPA?" ISAT gave 'em a couple of ideas. Steve Koonin, who was the provost here at CalTech at the time, was the chair of Jason's, which is another study group like ISAT-- that was, at the time, kind of giving ideas to DARPA or somehow affiliated with DARPA. And I think it was the Jason's and probably many of others that suggested the idea of a road race, a robotic road race. And so because Steve Koonin was involved and he was here at CalTech, when they finally announced the Grand Challenge, Steve and I talked and said, "Geez, do you think we should put together a team." And so at CalTech, because the research activities in robotics were relatively small-- it was basically Joel Burdick and I and we weren't doing anything at that time that was directly related to driving and, you know, other sort of-- parts of the problem we were doing. But a lot of 'em we weren't. We thought about doing it as something that might involve the undergraduates. So I told Steve Koonin that I'd be willing to help coordinate it if there was interest from the undergraduates to do it. And so we, you know, put some flyers around that said, "Come find out about this thing." And it-- my recollection of it is that sort of 60 to 70 students showed up for the information session, which at CalTech, that only has about 900 undergraduates, that's a large fraction, really, of the undergraduate population showed up. And so we said, "Okay, well, it looks like there's interest so we'll do it." And so we, you know, kinda pulled together a really ragged team and I guess this might've even been-- 'cause the first challenge was in 2004, right? So this must've been in 2003. So anyway-- so pulled together a team and somehow managed to get a vehicle to the competition <clears throat> that first year and made it to the actual race. And, you know, there was always a question about how you count. I think if you count by miles traversed we were, you know, sort of fifth or sixth but we made it 1.1 miles on a 150 mile-ish course or something like that. But I think CMU made it the furthest and they only made it maybe seven miles or something like that. So anyway, that was a lotta fun. You know, that first year was just, you know, sort of-- that anything worked at all was amazing because it really was sorta pushing it. And then we stayed involved and the next year participated and that-- the first year, I think, because we hadn't set it up and structured it within the courses, there were really only about a dozen undergraduates who were trying to get this to work. By the second year, we'd figured out how to set it up within our course structure and other things. And so over the course of that competition-- so that would've been the 2005 competition-- over the course of a year and a half, we must've had, I don't know, 70 or 80 undergraduates who actually participated in it. So, in the first competition, we had, I think, 70 or so who were interested but the number that eventually, in the end, participated was probably 30 of 12 who really did the last bit. The next year, we had 70 who really participated in

the whole thing and we had probably 30 by the end who were really putting it all together. That was another desert race. Very cool, had a lotta fun, didn't do as well. I think maybe made it five miles that year or something like that. Had a nicely spectacular end to the race which was that our vehicle, which we called Alice, decided that a concrete barrier was not really a barrier but perhaps something to go over. And it was right in front of the press area and so we have some wonderful video taken, both from the view of the press-- who were filming, watching this robot come charging over a concrete barrier at them-- but also from the point of view of Alice, watching some of the media people, you know-- and you can go on YouTube and find it and sort of, you know-- you can see people sort of see this thing coming and they're filming it and then you can <inaudible> say, "Oh, my god. This thing's gonna come <laughs> over the wall," and then they start to run and they don't know which way to run 'cause what is it gonna do. Anyway it's really a-- anyway. Noth-- nobody got hurt <laughs>. The vehicle did not. The-- it survived but it sort of broke some of the steering gears and other things but we fixed it, actually, at the race and drove it back home. So anyway that was a lotta fun and then the final year of the Urban Challenge, we also participated. Didn't make it to the finals that year but did make it to the qualifying event. And that year, we had a bigger-- not a bigger team. About the same size but involved undergraduates, graduate students, people up at JPL, and so we had a lotta collaborations with them. But that was a lotta fun.

Q: Have you collaborated a lot with people from JPL over the years?

Richard Murray: Yeah, over the years sort of off and on. I think, when I first started, we were very interested in doing things with them, and, you know, had a couple of dream projects and other things. We've done some space control types of projects, formation flying things, with them. A little bit. The DARPA Grand Challenge was probably the time we interacted with them the most intensely. We, you know, talk to each other a lot and we have lots of students who graduate and go up there. The projects that they're involved with that are on the NASA side are often things that have timeframes and other things that sometimes don't fit a Ph.D. and so we, I think, are interested in talking to each other but may not be working on it directly because the DARPA Urban Challenge was something where we really had to pull together people who had the right expertise. It was natural to go up to JPL and talk to our colleagues up there. And since then, we've continued a lotta those collaborations that <inaudible> started through that. And so we participate jointly in various projects and the JPL researchers actually teach our robotics course-- our kind of the intermediate robotics course here at CalTech so that works out quite well.

Q: Are there particular ones that you have worked with more closely?

Richard Murray: Larry Matthews and his group is probably the group that we interact with the most and Andrew Howard, when he was there, was one of our kinda key contributors on the

Urban Challenge. He's now off at SpaceX. So that's probably the group that we've had the most interaction. Fred Hadiez [ph?] runs the control group there and so some of the formation flying and other things have been with people in his group. So those are probably the two primary groups. Steve Chen a little bit but not quite as much.

Q: So have many of your Ph.D. students gone on to robotics? Did any of them do some specific interest in <inaudible>?

Richard Murray: Yeah, some. So <clears throat> a number of Ph.D. students are now up at JPL working on control problems but because JPL does robotics those are often associated with things in robotics. Francesco Bullo who's now at UC-Santa Barbara does some things and has done some things in coverage and other things that intersect with robotics. Trying to think who else is really doing things that, you know, you would sort of label as robotics. Lots of people doing unmanned vehicles, but again, a little bit more from the trajectory and control. Kristi Morgansen up at up at University of Washington, works in underwater swimming and things like that. So I'd say those students, in my group, have tended to go into robotics and control and so look at that kinda side of things. And then many of them gone just into control, having nothing to do with robotics or the community there. You know, biological systems or networks or other types of problems.

Q: How did you get involved with the <inaudible> bio-molecular <inaudible>?

Richard Murray: You know, I think when I-- so I went on sabbatical for a couple of years and I worked at United Technologies in 1998 to 2000, roughly, and when I came back from that, I was the division chair for Engineering Applied Science here at CalTech, which is kinda the-- our-- roughly, our dean of engineering position's no quite the same, but roughly. And so during that five-year period from 2000/2005, I didn't start anything new because I just didn't have time 'cause of my administrative commitments. But I kept going all the things that was going. But I was looking around, "What's the next thing to do?" And I think <clears throat> what was clear-- I served on an ISAT study in synthetic biology and what was clear was that there was some really new stuff happening. So, you know, the beginnings of synthetic biology and molecular programming are often traced back now only about 10 years. We talk about robotics being 50 years-- to things like the repressilator, which was something that Michael Elowitz built-- that he essentially built a small oscillator out of genes pulled from multiple organisms and put together and then it, you know, sort of inserted into E.coli. And so it was clear during that time that some interesting new results had happened. It was also clear that both natural biological systems, and likely engineered biological systems, were gonna use feedback in a lot of ways and control in a lot of ways to regulate their behavior. And so it was something I started looking at saying, "Geez, you know, this might be an interesting next thing to go look at." And so when I started to have time, after stepping down as division chair, I-- every time a project came up where it was a

joint project with somebody to look at something related to bio-circuits in that scale, we did it. In-- when was it-- probably in 2007 we wrote a proposal to the national science foundation-- maybe 2008-- as part of their expeditions project. A number of us got together and wrote a proposal to do molecular programming so we started something called the molecular programming project, got funded. And so that's a number of us here at CalTech, including **Paul Rodeman [ph?]**, who's the inventor of **D.N. [ph?] origami**, that's just an amazing technology. **Eric [ph?] Winfrey**, **Shukkie [ph?] Brookniles-Pierce [ph?]**, myself, **Eric Clavens**, who's a robotics guy up at University of Washington but now increasingly doing things in bio-circuits-- and just started looking at, you know, what we could do. I think I had always thought it would be kind of interesting to reprogram chemotaxis for example. So chemotaxis is the process by which organisms find gradients in nutrient fields usually and, you know, seek out higher nutrient concentrations. And so it was an interesting motion control problem. Looked a lot, it's clearly something where they're sensing their environment, they're moving their environment. And just sort of conceptually it seems like, "Well, why can't we build systems at that scale that do that and how might we do that?" And so molecular programming for me started as we started to think about, "Okay, well, could we do that?," and, you know, I think if you look at the last 10 years, there's just a lot of interesting trends. So the ability to synthesize DNA cheaply, the ability to build things structurally out of D.N. origami, the ability to build small bio-circuits and get them to work, all the **iGem [ph?]** competitions that are going on. Anyway it's-- as a 10-year period, there's just been this nice explosion of interests and so that area, I think's, gonna be very interesting for the next several decades.

Q: As somebody who approaches robotics from very much a control theory perspective, did you ever find cybernetics to be influential on your...

Richard Murray: Yeah, not so much. I mean, I sort of read Wiener's book, probably only about 5/10 years ago. I mean, you sorta knew about it. You heard about it <clears throat>. It's interesting now because, of course, we're coming back and looking at biology at a different scale, I think, than really, in some sense, Wiener was thinking about. Well, it was a pretty broad view. And so, you know, this sort of combination of looking at biology and man/machine and all of that stuff never really influenced, I think-- directly. I mean, I think, indirectly it did because, of course, Wiener started something that a lotta people were influenced by and it just sort of trickled through. But in terms of the direct influences of it, you know, we didn't really think all that much about it, although my first publication-- journal publication, was in the "Systems Man and Cybernetics" <laughs> and actually transaction one-- "Systems Man and Cybernetics." And I remember **Raft [ph?] Andrea** who was a graduate student here at CalTech and-- it got very involved in the Robocop and all of those sorts of things and is now at ETH. We were talking about the fact that control theory's got sort of a bad name as a field because it doesn't mean anything to anybody and people don't know quite what it is and it's hard to describe and other things. But <inaudible> we should rename it. I remember we thought, "Well, maybe we should go back and rename it cybernetics," right, you know, 'cause it was really about sort of putting

feedback in machines and doing other interesting things that biology does. So-- but that was kind of the only connections that came up with that work.

Q: What do you see as the big challenges and problems facing robotics in the next 10 to 20 years?

Richard Murray: You know, I think there're a lot of 'em. I think, you know, sorting out perception and how to take lots of information about our environment, very complicated information about our environment and reason about that is very interesting. So I think that was one of the things that the DARPA Grand Challenge or really the DARPA Urban Challenge taught us-- was that, you know, if you watch what happens as robots move around in human environments and try and reason about what to do, I think you have to I think a lotta the techniques out there are very interesting but they're not solving the entire problem. So one of the problems we're working on now, for example, is how to get robots to navigate in crowded environments. And so I have a student who's working on getting a robot to move though the cafeteria here, during lunch hour. All right, and so when you do that, you have to take into account it's a crowded environment. There are lots of people around you. They are going to react to you. So you can't just sorta say, "Well, I'll model them as moving in the way that they're going to move, without taking into account how they will react to me, how they will perceive me, how they will perceive each other. You know, if I see two people coming together, one of 'em's gonna slow down. The other one's gonna go by. How do I take that into account. All of that, thinking about how do you perceive the environment around you, how do you represent in some way the environment, how do you reason about that-- that I think is an enormous and interesting challenge. In many of the applications that we're still interested-- we're interested in things where you can more or less certify at some level that correct behavior will take place. And so these things were removed from kinda the lower levels of control to control at higher levels of abstraction but we still want to have specifications of what proper behavior is and synthesize controllers that satisfy those specifications. But the specs are no longer, you know-- track with, you know, one micro inaccuracy or no overshoot past sorta these traditional specifications that might be relevant to manufacturing robotics but really are the wrong thing, if you're talking about robots that are going to be interacting with society or interacting with humans or whatever it is, and you wanna say, "Look if I gonna build a driving robot, how do I design the control protocol that tells me, you know, when to change lanes, when to go through intersections, how to react to unexpected events in a way that I know that it will never run into anything." Right? And then, of course, there's-- you've got some abstracted version of that problem that you need to be able to solve. And then getting it to work robustly and reliably and all that stuff is yet another chance. So I think all of those are interesting.

Q: So it seems like you've worked at a number of places in your career, come across this issue of the robots and somehow interacting with a human or being part of a human environment. So how do you deal with people in that context?

Richard Murray: Yeah, so, you know, I think when we were doing the-- so in the things that we originally worked on, we weren't worried about humans interaction with humans <clears throat>. In the Urban Challenge it became clear that we needed to think about how to do that and that there were some interesting problems. And so there, of course, it was very structured because you're not dealing with humans walking and running and doing other things in that context. It was just human-controlled machines. But the machines obeyed the same dynamics that you did or at least-- you know, if a car is moving at this speed, I can predict over the next second, where it's gonna be with high accuracy, over the next minute with much less accuracy. But those were sort of tractable problems. I think the things where we interact with humans as humans get a lot more complicated and so we're seeing this in some of the work that we're doing in the cafeteria where, you know, how do you represent the way that 50 people interact with each other? And we do this sort of naturally as we sorta move through an environment. You see that somebody's coming up, you slow down, you speed up, whatever. And I don't think we know how to describe how we do it. It's just sorta built in. In fact, we also do some work in insect flight control and insect control systems of various sorts-- and, not our work but work of Pietro Perona here at CalTech and Michael Dickenson at the University of Washington, has looked at how flies move around. And you see things that look very similar to what humans do. That is that two flies are coming along and one will slow down so that the other can pass and things like that. And that's some evidence that maybe a lotta these behaviors are really encoded quite low down. All right, that is that you see some optical flow and that causes your muscles to move in a way that, you know, forces you to slow down. Maybe humans-- humans, of course, are much more complicated but also have some similar types of reactive sorts of things. So how do we combine those reactive approaches? How do we design reactive approaches so that robots react appropriately with humans in the environment, with the more deliberative things where you're somehow reasoning about that? I think those are interesting. There're a lotta people thinking about humans that are interacting much more closely. So humanist systems and things like that. Not a problem I've thought about a lot. I think it's a hard one. So-- but we haven't thought so much about that one yet.

Q: Do you observe the cafeteria initially in order to find out what people are doing or how they react to each other? How do you approach that issue?

Richard Murray: Yeah, so, you know, the set-ups or probably what you'd imagine is overhead camera that are looking at some area and we can more or less track people's location and how they move and start to build very data-driven representation. So, you know, we're trying to build kind of, more or less, conditional probability distribution. So if you think about a trajectory as an object you can say, "Well, okay, well, what's the probability that given this trajectory and given this trajectory that they'll have some reaction with respect to each other. And then given that, I could say, "Well, now if I do this, can I, in a sorta **Besian-like [ph?]** way, reason about the probabilities that that will avoid collision or that people will react in certain ways?" So it tends to be very data-driven and kinda stochastic representations and not trying to say, "Here's a model of human decision making," as much as just "Here's a model of how these trajectories that

humans take interact,” and that seems to, at least so far, be-- it’s giving us interesting results. You know, is it gonna be the end-all way to navigate through a cafeteria? No, I doubt it but, you know, it’s interesting to look at.

Q: What kind of robot are you using?

Richard Murray: In this case, just a Pioneer, a very standard robot, nothing complicated at all. We just, in some sense, needed the simplest robot that could maneuver at roughly human-like speeds. And so we didn’t have to build anything fancy or anything else.

Q: It seems also like people would perhaps have slightly different reactions to a robot than they would to a person.

Richard Murray: Yep, we’re playing with that. So that’s right. So, you know, one of the things you can do is you can run the robot through and see how people react to it and, in some sense, that’s what you care about <clears throat>. In this particular context, you know, one of the things is just to put something on top with eyes or, you know, a camera that looks like eyes and other things, so that people’ll realize that the robot could probably see them and will react to them. Because otherwise, if you just have a, you know, sorta the short Pioneer robot moving around then people interact with it very differently. So, yeah, I think that’s something that-- it was the thing about, you know, urban environments with cars. I think people might do a double take when they don’t see a driver behind-- the reality of the matter was that, if you’re driving in a human environment, you almost always have a safety driver. So most of the time, people don’t even know that it’s a computer controlled car that’s driving through. So that’s a much simpler situation ‘cause they interact in the same way and it’s very structured. But you’re right, something like a pioneer robot with a, you know, pole sticking up and a camera mounted on the top of it, you know, I think-- although, it’s CalTech. You know, somehow <laughs>-- you know, it’s like, “Eh, whatever,” you know, “I need some pizza.” You know, I don’t-- it’s amazing they don’t react as, you know, don’t crowd around it. They don’t, you know, sort of ask. There are obvious people around who are keeping track of what’s going on and they just leave us all alone <laughs> somehow.

<off topic conversation>

Q: One wind up question...

Richard Murray: Okay.

Q: ...which is, for young people who are interested in robotics, what do you recommend for them?

Richard Murray: Oh, I think, build stuff. You know, I think the great thing about the robotics community is that it's been a community that's built stuff and tried things out. And so, you know, I think if you're gonna get into robotics, the way to do it is to start building robots that do things. You'll find interesting challenges. I mean, every robot that we've built, we've always learned something new from. We've always, I think, gotten something out of the experience of building in and trying to implement things and seeing what worked and what didn't. And so I think, you know, starting to build robots early in your career is good <laughs>.

Q: Great.

Richard Murray: Good.

Q: Thank you very much. Thank you.

Richard Murray: Great, you're welcome.

End of RichardMurray.mp3