

Oxford, UK
1:29:10
m4v

Ron Daniel

An interview conducted by
Peter Asaro

March 9 2013

Peter Asaro: Okay. So if you could just tell us where you were born and where you grew up.

Ron Daniel: So my name's Ron Daniel. I was born in Maidstone, Kent, England, and that's where I grew up.

Peter Asaro: And where did you go to school?

Ron Daniel: Went to school in Maidstone, what was then called the Maidstone Technical High School for Boys.

Peter Asaro: Technical. Did you learn any engineering while you were there?

Ron Daniel: Well, yes. It was an unusual set of schools that only really existed inside Kent. Which were schools which offered people engineering education as well as the more traditional education available in grammar schools.

Peter Asaro: And where did you go to university?

Ron Daniel: I went to University at Brunel University. Which is in London.

Peter Asaro: Yeah. Did you study engineering there?

Ron Daniel: I studied electrical engineering there.

Peter Asaro: And then you went on for your graduate work?

Ron Daniel: My graduate work was at Cambridge, King's College, with Professor McFarlane, on control theory.

Peter Asaro: And what made you decide to go into control theory?

Ron Daniel: It was the part of the course that interested me the most. I tend to have a mathematical background, and I tend to have a mathematical approach to things. So that's what appealed to me more than the other parts of the course.

Peter Asaro: And what got you interested in robotics?

Ron Daniel: Oh, a long history of interest in robots and mechanisms. I also was quite a science fiction fan as a youngster. And I read a lot of Asimov and so forth, "I, Robot," and Arthur C. Clarke. So I had an interest in robots. There's also my name, which I found amusing, given one of the main characters in the novels was called R. Daniel. So it felt right to go <laughs> in that area.

Peter Asaro: So what do you consider your first robot? Or robot project?

Ron Daniel: Well, that's an interesting question, because I got into robots via teleoperation, not straight robots. So I was actually interested in man/machine interfaces rather than autonomous systems. And so after I finished my Ph.D., my main focus of my Ph.D. was how to enable designers, design control systems. And I was doing my Ph.D. right in the middle of a lot of excitement going on in the subject about a new subject called H-infinity, which was about defining control systems in terms of their robustness properties. It was interesting to look at what people were trying to do, but it seemed to factor out the human being from the design process. And my Ph.D. was on trying to get the human involved with the computer much more, such that you can actually use the two sets of skills. From there I looked around at robotics projects, and the main one that I found, that interested me, was working with Central Electricity Generating Board on the remote maintenance of their nuclear power stations. So the UK at that time had two types of power, power station. We had the Magnox reactors and the advanced gas-cooled reactors. Magnox was getting old, and there was a pressure to extend the life of these reactors. And so one of the issues was how to do remote maintenance and inspection of these devices. So I came to Oxford to join a research team, a research fellowship, working with the Central Electricity Generating Board at Marchwood, which is a lab just outside Southampton. When they did most of the development work on the design of remote manipulators. And that's how I got interested in managing interfaces, how you connect up human beings to machines, so you can get them to cooperate.

Peter Asaro: And did this lead to your interest in kinematics and haptics?

Ron Daniel: Yes, it did. I was quite interested in knowing, or trying to work out how you actually could match a human being to a remote manipulator, performing a complex task. So the early work was mostly associated with how on earth you get a manipulator down an 8-inch pipe through 30 feet of concrete. And then reach around and perform useful tasks. So the early work was mostly about solving the kinematics problems quickly, given the computing resource that we had available at that time. So there's a lot of interest in dealing with things such as singularities, representing manipulability, how you go about computing inverse kinematics quickly and in real-time. There's also the issue I also became interested in, which was the fact of how do you

perform remote tasks at considerable distances from a support point? So if you can imagine, you've got something about that diameter, which has gone down 30 feet, is anchored either to the floor of the reactor or to the ceiling of the reactor. And then it has to reach out 10, 15 feet in an arm which has to unfold itself and then perform useful tasks. And the sort of useful task we did was remote drilling, replacement of things called top hats, thermocouples, replacing bolts holding the thermo shielding into the concrete pressure vessel. So the issue was how do you go about performing these tasks? Or the method we developed was basically on a almost self-contained robot in itself, which had a number of machines, like drills, nut runners, all encapsulated inside something about this, which would be carried at the end of the arm, thrust up against the, say, the ceiling. The operator would press a "go" switch, and the machine would carry out basically like a machine <inaudible>, would drill a hole, insert a nut, run a nut on it and so forth. So I became interested in how would you go about performing such tasks at a remote distance? So there were two issues that were interesting me. One was the effect of the robot machine's compliance on the behavior. Sort of how can you accurately position a robot if it's bouncing around carrying significant loads subject to machining forces and so forth, and how could you transmit the information to the human operator such that the human gets a feel for what's going on? Because you can't really afford to have such machines make mistakes inside a nuclear reactor. It's rather expensive if you break something down there. So there were issues about reliability and transmission of information to the operator and so forth.

Peter Asaro: Now, what year roughly was this work?

Ron Daniel: This was between 1982 and about 1985.

Peter Asaro: And so what kind of a video feedback were you using, or were you mostly focused on the force feedback, haptic feedback?

Ron Daniel: There were cameras also inside the nuclear reactor, which fed information back. But it was not processed in any way. It was just we looked at using stereo systems, for example, to get a sense of depth. There was very little haptic feedback from the device. The CEGB were interested in trying to build haptic devices, but the concentration on how people built these things was basically they tried to build some mini robots. They had problems with stick-slip and all sorts of problems associated with mechanisms while driving something which moves about and transmits force to a human being. So that was the start of when I became interested in, firstly, how would you design a mechanism such that you can transmit forces to a human being? Such that the mechanism designed to match the dynamics of the human hand? So the device that I came up with was based on a parallel rather than a serial mechanism. I used direct-drive rather than geared. And that work was done mostly with the UK Atomic Energy Authority at Harwell. And the device which I came up with, we call it the bilateral Stewart platform, or BSP, which was a six degree of freedom parallel mechanism. Direct-driven, very, very hard mechanical

bandwidth, the order of small motion, 100 hertz. You could get reasonably large motion at 10 hertz. There's a lot of power inside the system. And we then tried to couple that to a Puma 560, which we were modeling as our remote manipulator to see how well that would perform. Because that's when I was working with, quite closely, with AEA. That was, I would say, late '80s. With a guy called Patrick Fischer, who was my Ph.D. student. He worked with AEA. They had a particular vision of trying to productionize the remote decommissioning of various nuclear contaminated artifacts within the UK, mostly associated with the UK atomic energy program in the 1940s and 1950s. Sort of things like early, early bond manufacture and so forth. So this left a number of legacy components, which you have similar legacy components in the U.S., where you have poorly defined geometry, you have stuff inside ponds, you have machine tools that you want to take apart. And so we were interested in trying to work out whether we could use a nuclear hardened commercial robot and use that device as a remote teleoperator to perform the sort of tasks that would normally be performed by a human being wearing a hazard suit. So if you can imagine somebody with a big saw wearing something looked like a Michelin man there being pumped in. They would be in a radioactive environment where they could only work for, say, one or two hours trying to cut these things up and put them inside stainless steel cans ready for long-term storage. So we had a vision that we could use nuclear hardened commercial robots such as the big Puma. And we were looking at using things such as huge jigsaws and big drills. Cutters which would cut through plate steel very quickly. So they wanted to be – well, we had a vision that we could couple a human being to these machines. So I had a look at the literature on human-machine interaction, and I became very surprised at the nature of the literature. And this is another thing which I came up with, and a number of people that I felt thought the same. In that what people tended to do was they tended to build a machine which did something and then they reported on what it did. So the literature of the time seemed to be full of, "Look, Ma, no hands," type papers. And with very little effort going into the underlying science behind what they were trying to do. So there wasn't much in terms of trying to develop experiments which would determine, say, the science of manipulation. It was mostly about building something to perform a specific task, and then you'd report on how well it performed that task. With little thought about, "Well, can you actually reduce the problem to a much simpler problem and ask yourself questions about the underlying science and the physics of the system? So that got me interested in the design of remote manipulators. And one question I felt had not been answered was how would you design a good teleoperator and what limited the performance of teleoperators? And that people were designing these things for 50 years. They'd been sort of trying to build these devices. And so I asked the people in the industry. I said, "Well, okay. If I gave you two manipulators, one designed like this, one designed like that, how would you determine which one was better? What performance measures would you use?" And they couldn't answer. And then I then asked the question, "If I wanted to design a manipulator with a particular specification, how would you achieve that specification and how would you determine whether that specification was achievable? Again, that could not be answered. So I developed a model of the manipulation process in which I removed the manipulator and just replaced it with an abstract object performing manipulation. Then start ask questions of that model regarding what would be the consequences of achieving the design aim of perfect transparency for this device? And it quickly became apparent simply from application of

Newton's laws, that any perfect manipulator would be something which would perfectly reflect momentum. And if you perfectly reflect momentum from a large manipulator, if you imagine what we had, we had a manipulator carrying a 50 kilogram saw which would then thump into a piece of steel. Now, the sort of forces involved with one of these very large robots driving into a lump of steel were extremely large. If you got perfect force feedback for this device, and you're holding onto a very low mass, very high bandwidth device and it hit something, it'd just fly off. Then if you imagine you got the forward coupling in which the robot itself is actually following your demands, the robot itself would follow your rebound and itself would rebound. And so what we found was that this huge increase in kinetic energy, by virtue of the fact that you are actually faithfully reproducing the momentum that is the force, the changes in force, so the changes in momentum that the device is experiencing. So that got me thinking about what type of information that the human being needs in order to perform manipulation. So it got me interested in the structure of how people feel force, and we came up with the notion of trying to excite the kinesthetic field at very low frequencies as well as exciting the sense of touch at high frequencies and cutting out the middle frequencies. And we chose how much to cut out the middle frequencies and the bandwidth of that notch by the limit space, by the momentum transfer that you get between the device and the human being. Such that you keep the momentum and transfer such that the velocity of rebound of the human is going to be the same as would've been the velocity of rebound of the device. So you don't get this increase in kinetic energy by virtue of the perfect forward coupling of your position. So the human being would then feel a much reduced mid-band force. Consequence of that is we developed this strategy and my ideal, well, we put together a video which was designed to put the noses out of joint of certain people. At this time there was a thing going on about repair of satellites using the space shuttle. And they were training astronauts to put in things, in orbit units, sort of big sort of boxes about this big. And I went and visited JPL, in which they were developing this system for inserting these boxes. And they were developing various haptic feedback devices for enabling astronauts to perform this task. And they had these boxes about six inches on a side, which they would insert, or try to insert inside satellites. And they had this ground facility for training people and testing out various ways of doing it. And they're finding it quite difficult. So when I got back to Oxford, I tried to set up an equivalent task. I don't know if you know about LEGO, which is, these are tiny little blocks. The difficulty of remote manipulation increases as you reduce the size of the object and the fineness of how you want to control the device. So what we did was we set up a little LEGO house remotely, and then instead of a highly trained astronaut, we got the son of one of the academics, who was, think he was nine at the time. And so we took a video of this little guy standing on a box to reach the haptic interface device, and he built a small LEGO house remotely, using force feedback on a Puma 560 arm. And he did it stably. And the amount of training that he had was about 10 minutes. So we managed to get this little child to perform a very, very difficult remote task, which others were trying to perform spending much more money than we were spending with much more highly trained people. I thought it was a good demonstrator of what you could do if you actually matched the information flow that you required to the dynamics and the information requirements of the human hand. So that got into the area of basically matching information flow to the human. So the concepts such as impedance matching in electrical engineering. And here we were trying to basically do

information matching by making sure that information that was encoded in the momentum was matched to the dynamics of the hand and excited the right receptors in the fingers at the same time. So that was one area that I was interested in. The other area I got interested in was, as I said, the control of compliance systems. So we built a rather large compliant robot to model how a compliant arm would work. So this is a huge direct-drive robot, which was novel at the time. There were very few direct-drive robots around. And this was a direct-drive compliant arm that carried out a number of research questions with this device. One was how do you compensate for compliance to get accurate positioning? And the mechanism we used was to use an excess skeleton around the compliant arm, to pick up the kinematic, rigid kinematics, in parallel with the compliant kinematics, and use that to derive the true end point. And using that, we were able to demonstrate really high speed control of this dynamic arm by incorporating feedback from strength senses on the compliant structure. So we became interested in how do you use information from strain sensors and so forth, for the control of vibration? At the same time, there was also an interest in how do you go about achieving high speed computation? So you got to remember, at that time, a really fast computer was 10 megahertz, which seems ridiculous now. But then 10 megahertz was a fast computer. And so we had a very early multi-processor computer based on something which people now wouldn't know what it is. It's called a transputer, which was a computer chip with integrated communication such that it facilitated process of communication. And so we became interested in how you go about using parallel processing to compute things such as inverse kinematics, inverse dynamics, and so forth. So that got us interested in inverse dynamics. And so I had a number of people looking at implementation of inverse dynamics for vibrating systems, compliance systems. So we managed to demonstrate remote drilling using compliant arms. That then went on to generate interest in where does the information come from when you perform remote force reflection tasks? And it turns out that the information flow is actually below the capacity of the encoders to send the information. So if you can imagine that you're performing a rigid contact task, that you've got these encoders and joints, when you're in contact with a rigid surface, the joints themselves are not moving sufficiently to excite the encoders. So the rate of information flow that you get from the joints is very, very small in terms of bits per second. So Shannon's Theory tells us about how much information you need in order to perform a task, in terms of the bandwidth and the information rate. And the information rate going on inside these joints was almost zero, because you're getting the occasional tick from the encoders. And these encoders were being used in velocity feedback systems. So we think, "Well, how on earth is this working?" So we sort of did some analysis of this, and it turns out that it was, in fact, the non-linear friction of the drives, which was stabilizing these machines when they're in contact with rigid surfaces. And we demonstrated this by disabling the velocity feedback and the thing was still stable. And it turned out that contrary to what it said in the literature – the literature said that stable contact with rigid surfaces is very difficult. And stable contact with very soft surfaces is very easy. What we found was stable contact with soft surface is easy, but also stable contact with rigid surface is easy. What's difficult are intermediate surfaces where the information rate is insufficient to stabilize the device, and the device is not going so slow, such that the nonlinear friction effects are not stabilizing the feedback loops. So we found that, for example steel, and steel was stable. A robot on a sponge was stable, but if you got something which would sort of spring, it was

unstable. So that surprised me in terms – it's very important when you're designing feedback systems to think about the information flow. So that brought me on to the dynamics of force reflection, force control. So how do you go about controlling a manipulator when it's performing a contact task? Now, a contact task, as I said, was considered to be very difficult if a very stiff surface. And what people reported in the literature was very high phase lags in force controlled server mechanisms. So if you identified the transfer function, say, of a robot performing a rigid task, rigid contact task, you got phase lags reported of several thousand degrees. Which was humongous phase delays in these feedback loops. And people say, "Well, this is terrible. No wonder they're unstable." So I got to thinking, "Well, where on earth are these phase delays coming from?" And the accepted view in the published literature of that time is it's coming from the compliance in the arm. So I did an estimate of roughly what the arm would have to look like to give you this level of time delay in the loop. It turned out if the time delay was due to the compliance in the arm, you'd need an arm which was roughly, say, 100 meters long, made out of 1 millimeter steel wire, where you basically twisting this wire, watching the waves propagate. And these arms were actually made out of Monaco Aluminum structure, and they were, you know, as rigid as this table. <bangs table> They're like this. There's no way this arm had this type of time delay in its mechanism. So that got us interested in the mechanism for where these time delays come from. And what we found was that the delay was coming from a non-minimum phase effect generated through parallel paths in the transmission. So if you take a gearbox and you think about a traditional gearbox, you have a cog which rotates and another cog which rotates against it. You transmit torque. That works fine when the gears are moving, so on and so forth, keep moving. If the gears are essentially stable, you get stick-slip motion between the teeth. During this stick-slip motion, there are now two mechanisms for transmitting torque. Firstly, there is the motion of the stick-slip between the teeth, and secondly, there is motion by virtue of the bearings that support the gear wheels. And so what we postulated was that effect. We had two parallel paths of transmitting torque. One was stick-slip through the teeth, and the other was the mechanisms, the basically rocking mechanisms, between the gear wheels. And so we went back to our identification experiments and we actually managed for the Puma arm to identify where these phase delays were coming from. And you would get these sudden 180-degree phase changes in the high-frequency behavior of these arms. And you could identify which particular gear teeth were sticking. You could say, "Ah, that one's coming from that gear tooth sticking on that one. And then that one frees up and this one starts vibrating." And we could actually correlate all of the phase delays with the vibrational modes of the internal structure of the gearbox, or the gear train, which transmitted the torque inside this arm. So that was the postulation. So we then built a mechanism to try and demonstrate it experimentally. So we built a gear drive, a singled gear drive with a tiny arm in the output. And so we drove it and we got exactly the same behavior. We got this very large phase delay. But this gear device was special in that we actually had two motors and two gearboxes, such that we could drive the output differentially. So the other way of driving this, instead of the drive, just driving the arm backwards and forwards, get this huge delay, is you could actually drive the gearboxes differentially. So one moving that way, one moving the other way, and the output was the difference in velocity between the two, such that you never got stick-stick in the gear teeth. If you did that, and then drove, tried to drive the little

lever, which was the pretend robot link, it became almost frictionless. If you got hold of it and you could spin it and “vzzzh,” just keep on going. And if you did a dynamic experiment, all of the phase delays disappeared. So experimentally we were able to show that if you removed the stick-slip behavior such that you removed the capacity of the gears to perform these multi-modal vibrational patterns, you could get rid of all of the phase delay. And so we looked at whether or not the appropriate way to design arms which are performing force reflection tasks, would be to drive the arms differentially. So that started another area of interest in which I became interested in working with the Joint European Torus, in which they were interested in having this arm which could go in and do maintenance once they injected tritium into the device. Don't know if you're aware what JET is. It's the European fusion research reactor, which is about 10 miles down the road from here. Probably closed down by now. But in those days it was being prepared to be fired up to do fusion research using a tokamak. And one of the issues was it was going to be ferociously radioactive once it'd been fired up and used with tritium. So they were devising ways of inputting large manipulators with force reflection. So the same problems with the CEGB. So same pattern again. And we were looking at ways in which you could actually use drives to remove the friction from the drives. So you have two problems. One is the phase delay and the other is if you're trying to use force reflection to a human, you really want to get rid of the friction in the drive train for the input device. There are a number of ways of doing that. You can either use a very low friction device, use fancy feedback. The third way of doing it, which is what we were looking at, was using differential drives, such that you'd get rid of the friction altogether. So the friction is then just, because you're driving both directions all the time in both ways, there's no coulomb friction. There's only the viscous friction left. So it becomes a very linear system. So we looked at that, but it didn't come to much, because it turns out that if you have differential drives you need very large power flows between the two drives in order to control very small signals. So there's an issue in terms of control authority of differential drives, which needs further investigation, which we never pursued any further. But that was the main practical constraint on differential drives.

Peter Asaro: Did you do much investigation into neurological proprioception when you were thinking about the issues?

Ron Daniel: A little bit. I worked with a local medical neurological group called Mary Marlborough Lodge, which looked at mechanisms for helping or assisting particular type of disabled people. And one area that we became interested in was a particular type of disability in which people tended to freeze up. So I had a Ph.D. student who looked at ways to enable these people to communicate with a computer. So what we designed was, in fact, in about 1990. Perhaps a bit earlier than 1990. Was a wrist with a set of accelerometers, a wristband, which we put a number of accelerometers in the wristband. And using this wristband we could compute where the arm was. So this was a very early form of gesture control. And what we found was that these kids that couldn't perform tasks like touch a computer could actually perform gestures. So if you said to the – so if you had 12-year-olds that really couldn't use things like a computer, you say to them, “Could you –“ you know, one of our favorite ones was Superman. And they'd

do that. And that gesture could then be used to control things. And another one they liked was stroking a cat. They could stroke a cat and do Superman and use these type of gestures, or hammering. And we could use these as ways of communicating with the computer via these accelerometers. And so we developed a filter which was based on some predictive control work the other colleagues were doing, which would, from the accelerations, estimate what the excitations were in the muscles driving the accelerations. So you'd back compute the excitations in the muscles from the accelerations that we measured to compute the intentionality of what the person was trying to communicate. And we demonstrated that you could actually analyze and communicate with a computer using these gestures. So it's slightly different to the more modern stuff in which you're trying to recognize things like hand shapes and so forth. This was more to do with working out the muscle groups that were firing in order to generate the acceleration such that we could then predict what it was that the person was trying to do and use that as a signal for driving a computer.

Peter Asaro: I was just curious if you knew whether in the human arm, whether you're looking at velocities versus momentums or a little of both or... In terms of the actual neurological signal?

Ron Daniel: We didn't look at it that way. In terms of the feedback controller I did, in terms of human-machine interaction, I was more interested in matching the dynamics, sort of the – so if you imagine my model of the human receptor system was of a set of band-pass filters. And I was trying to work out where the information was most important, such that I could maximize the information flow whilst matching the momentum transfer. So in theory, you'd want all the information available from the environment to be transmitted to the human being. But if you've got a mismatch in a device where you have a very large device being controlled transparently by a very small hand, you can't transfer all the information. Because if you did, you'd get this momentum transfer. And if you get momentum transfer, the human control system just can't cope with this. And so you have to be able to keep the momentum transfer within manageable bounds whilst transmitting as much information as possible. So you have to choose where the information is going to be flowing. Whereas the work we did with the communications was these people had very, very large number of difficulties in moving, but we postulated that they could control muscle groups, but not very well. So our postulate was that there was some problem with their feedback mechanism. But feed forward mechanism was still intact, so they could perform tasks. It's just that they couldn't correct them. So what would happen is they would become very rigid when performing a closed-loop precise task. So if you replace a very precise task with a very imprecise task where the motion is very much open loop, then they could perform these motions, and then we could back compute what the muscle firing pattern should've been based upon that acceleration, rather than trying to measure the myoelectric signals directly. So we were interested in knowing whether we could back compute these gestures from the accelerations, so... And proved to be quite interesting. We got some way of doing that.

Peter Asaro: In your work on the atomic robot manipulators, were you influenced at all by Red Whittaker's work, who designed the systems, the robots, that cleaned up Three Mile Island? Were you aware of that work?

Ron Daniel: Not really. That's a different type of problem. This was very much driven by the needs of tightly constrained static devices where you're not trying to navigate through things, you're not trying to find your way through things. We assumed that most of the stuff had been mapped out and we were particularly interested in design of the machine and its impact upon the human being as they perform the task. It was more about the stability of the human-machine interface, rather than sort of AI type problems of moving around. I've never really worked much in AI. It's mostly dynamics, mechanism design, control theory, information flow, not taking the data and constructing a model of the world and then moving about in that world.

Peter Asaro: So who, in terms of theory or other researchers, have been influential on how you pursue problems?

Ron Daniel: Oh. Think, ah, difficult to identify. You know, I like the work, the main people that influenced work I was doing, people like Shannon <laughs> and Bode, Nyquist. So I come from a control theoretic background and looking at how do you talk about information flow within a controlled contest? It's a big gap in control in that people talk about control systems in terms of dynamics. They rarely talk about control systems when you're really talking about information flow. So a simple little paper I produced many, many years ago to demonstrate this was one of the things people when they design a robot drive system is they want to design a simple proportional plus derivative feedback loop around some motor. And classically you would see sort of folksy ways of doing this in which you turn up the velocity feedback until it becomes noisy. Then you turn it back a certain amount, and then you wind up the position going to a certain amount until you maximize the performance of the sever mechanism. Well, this was nice and folksy and worked well for devices in the 1960s, say, when the position devices were things like potentiometers. But you still got the same behavior when you had encoders. So a question that I posed once was, "Where does the noise come from in an encoder?" And in fact, it's not noise that you get in an encoder. What happens is as the velocity of an encoder goes to a certain limiting value, which depend upon the amount of feedback you got on the system, you actually have a restriction on the rate of information flow from the encoder. Because you only get information by virtue of the line crossings. So if you then look at the Shannon limit of what you can achieve based upon the information flow coming from the encoder, you find that a lot of the control predictions of what you can achieve on this encoder are nonsense, because the bandwidths that the control mechanism is predicting is way beyond the rate that the information flow can actually support. And what happens is that when certain boundaries are passed, the encoder and the control system interact, the dynamics interact, and the system becomes chaotic. So if you try and extract more information out of a control system than is available from the sensors, there is a tendency for the system to inject noise in the form of chaos.

Peter Asaro: Sounds like something Ashby would say. <laughs>

Ron Daniel: That's what the physical system does. And you start getting these strange cyclics. You start getting this frequency doubling and so forth as the device starts limit cycling between the encoder lines. If you got too high a gain for the speed that you're going, for the bandwidth you got available, for the bit rate you got available from the encoder. Then you start getting all sorts of strange values coming out of it. And it wasn't noise that the people were experiencing. They just had too much gain for the bit rates that they had available. That is they're trying to exceed the channel capacity of the control system. So control system not only has a bandwidth. If it's gotten a sensor on it and it's digital in any way, it's got a channel capacity. And very few people actually consider the channel capacity of a control system. And that puts a significant constraint on what's achievable. And I've come across that with many colleagues that have tried to achieve impossible things. For example, I had somebody trying to achieve velocity control of a turbine set and they said, well, I can't get this thing to work. And I said it'll never work because you don't have enough bits available. Your number of lines on our own code is too small such that the bitrate that you're getting from your encoder in our feedback loop is too small to support the bandwidth that you're trying to achieve in terms of your disturbance rejection. If you're trying to get more information – in disturbance rejection you're basically trying to reconstruct a disturbance using a control system and reject it. If it in reconstructing that disturbance you need more information than is available from your sensor system you're never going to reconstruct it no matter what fancy control system you use. You just don't have the information rate to support that task and they just don't consider that. It's not part of the control theorist's lexicon. They don't think about the channel capacity of their sensor. Any sensor has a channel capacity. And you cannot perform a task, a dynamic task which requires more information than is available from your sensor. If your sensor is providing a 1000 bits per second, then that's it. That's your channel capacity. If your task requires 10,000 bits per second it doesn't matter what fancy software you've got you'll never achieve your task. You don't have the channel capacity.

Peter Asaro: Who have been some of the people you have collaborated with over the years especially in the robotics area?

Ron Daniel: Whoa. You're going back a long way now. CEGB, the Atomic Energy Authority, U.K. Robotics. I spent some time developing software for companies selling machine tools. I was using some robotics work that I did on parallel mechanism for a company called Geodetics which was trying to develop a parallel mechanism. It's basically a robot for doing machining. They were interested in trying to improve the accuracy and speed and controllability of those devices. Various people in Europe, all over the place. I'm trying to think through.

Peter Asaro: Any roboticists that we might recognize?

Ron Daniel: I spent a lot of time working with Ross McCarrie [ph?] who is now in the Australia who now works in mining machinery. Most of the people I tended to talk to were within Oxford. So the people like Andrew Blake, Hugh Durrant-White, John Leonard, Dave Forsyth, Margaret Fleck. Mostly Mike Brady's students. We used to spend quite a lot of time chatting to Mike and to the other guys that came into the lab. Quite a few people have been through Oxford and moved on and have become quite well known people like Hugh Durrant-White, for example. I've got some of my students now come back to Oxford. Paul Newman. I remember Paul I supervised his undergraduate project.

Peter Asaro: What was the project?

Ron Daniel: It was the design of an AGV, believe it or not, autonomous guided vehicle. And as an undergraduate project what I got him to do was to look at whether we could reverse engineer these electric screwdrivers to make a very cheap autonomous vehicle just as a platform. He had one or two of his own ideas and he developed that. A lot of students come through that. I remember Dave Forsyth. I remember John Leonard working on AGVs and his interaction with traction batteries.

Peter Asaro: What's a traction battery?

Ron Daniel: The batteries in an AGV these are big batteries. And we used to have these autonomous guided vehicles in the basement of the Jenkin Building which is an old Victorian brick building. And the ATV lab was in the basement with these huge support beams. You had to walk around in the dark underneath the beams as you went through. And they had these little vehicles that trundled around with little Polaroid sensors. You'd walk in there and these things would be wobbling around with cameras poking at in all directions. This is a completely alien world to me because I remember what I worked on sort of big chunky manipulators with people hanging on to them and going unstable, so hanging on stirring away all trying to perform a drilling task. And they had this vehicle sort of trundling about the basement with all of these little things going like this. Lots of ticking noises with the Polaroid's going off. They used to recharge the traction batteries in the basement which I thought was a bit dodgy because there's lots of hydrogen giving off of these big 48 volt batteries, big stacks of batteries bubbling away inside the basement is not a good idea. We used to have little safety arguments about that. And how do you go about unbolting the traction, the 48 volt traction battery. You don't use a non-insulated spanner when you do that. You have to ask John about that. So these were students. I remember those days. And I remember our very early Sun network, Sun workstations. We had these – well, they were based on 68,000 and 10s.

Peter Asaro: Those were the spark stations?

Ron Daniel: No, way back. I remember we put in the very first thick wired Ethernet, local area network in our lab. People used to come around and look at this coaxial cable, big thick yellow coaxial cable at a diameter which we had about – we must have had about four computers on this network which that was quite something. And we had these Sun workstations with 68,000 and 10s in them and that was 10 megahertz, powerful machines. And we had a server. <laughs> A machine with a server, a network. We had our very own version of a network system with a server.

Peter Asaro: Any funny robot stories?

Ron Daniel: I remember when Paul Newman discovered the Web. <laughs> Yeah, I remember the very first when the Web became available he came in and said, “You can do searches and find things. Look at this.” And he was absolutely full of that. And I remember the arrival of email on your desk. It was before we had to go into the sort of – in order to use the email for something you went into the central VAX server room and logged on to a terminal and then you could send email via text.

Peter Asaro: With the Pine?

Ron Daniel: Before Pine, yeah. Pine modern. <laughs> I’m saying early eighties. It was very, very early stuff. I remember when I printed my Ph.D. thesis I had a very advanced printer because I had a double daisy wheel which meant that not only could I type Roman numerals it could do mathematical symbols. I think I was one of the first people to print my Ph.D. thesis using a printer that could print mathematical symbols as well as other sort of ordinary letters. That was quite something when we did that. These huge things were flying about, big wire moving this great big double daisy wheel around.

Peter Asaro: You’ve also done some work at a startup company in robotics?

Ron Daniel: Yeah. One of the things that – we’ve been through some developments in the department. My robot lab was closed down and dismantled while we rebuilt the building. While that was going on I started to look at because Mike Brady had been doing it for so long I started to look at computer vision. I developed my own stereo code and I looked at ways at taking stereo and merging stereo, all new geometric stereo with photometric stereo using Kalman filters and factor <inaudible> filter. We developed this technology where we can take a fairly low resolution stereo reconstruction on an object and merge that with a very high resolution normal map generated by photometry. And that gives us quite a precise 3D small field 2.5D reconstruction of objects with very, very high resolution, sort of 100 micron resolution. And I used that idea to start a spin out company and I now have a spin out company which measures

wind volume. We point the 3D camera at a – do people want to know about the gory details of diabetes? One of the issues of diabetes is that you lose two things, two problems. One is you lose the sense of touch. And secondly, you have a restricted blood flow. A very common experience of people with diabetes is they might sort of stick their foot on a rose thorn, they don't feel it. The rose thorn goes bad. The blood flow is insufficient for it, for the wound to heal and they develop quite a nasty ulcer which can be quite a large hole in say your foot or your leg. It's not uncommon for people to develop these large ulcers on their legs and feet. And an issue with these things is how do you choose the dressing regime? In order to know whether or not you want to put an expensive dressing on these wounds you need to know whether they're healing or not and what the rate of healing is. And the problem is that it's difficult to tell what the rate of healing is using an area measurement. If you have an ordinary camera you photograph the wound. All you can really measure is the area of the wound. Our company sells a camera, a 3D camera. You point and click, you process it, you get a 3D model which you can then measure things like wound volume, exact area of the wound. If you want, you can fly in and inspect things because you've got such high resolution. And from that you can work out whether or not the wound is healing and whether or not you need to change the way in which you're treating the wound, whether to make it more aggressive, more expensive treatment and so forth. So that's what that is setting. It's also used for things like measuring scars and so forth, so it will measure the shape of a scar or shape of a wrinkle.

Peter Asaro: And did you do this with your students?

Ron Daniel: I did that on my own basically because I had sort of small office on the seventh floor of the lab. Seventh floor of the department, no lab. I had to sort of basically use paper bound research methods for three or four years while the lab was being rebuilt.

Peter Asaro: When was this?

Ron Daniel: This was late '90s, early 2000s. We now have a brand spanking new lab up there. And when I came to move my stuff into the designated space in the basement, I found that the lift was too small to get the stuff in there. At that time I was then there working with somebody called Mark Batcheck on using robots within wind tunnels. So we had this idea of having a very high performance force effecting robot inside what's called a free flow wind turbine – wind tunnel. If you imagine a big fan which blows a jet of air into free space and you stick something in it we then measured the forces being applied to the model under test and then we'd move the model under test such that we could simulate what the device would do if it was flying in free space. We could actually fly aircraft models or air foils inside a wind tunnel dynamically. And it was of interest to people that were interested in unstable – controlling the unstable regime, whether your control system is going to work. And that got on to a subject called harder in the loop control in which basically you're trying to insert a mechanism inside a control loop. But

part of the device that you're trying to perform experiments on is simulated in software. Other parts of the device are actually the real device. Typically, you would be interested in looking at the performance of some horribly nonlinear component which is not well understood which would be attached to a large object whose dynamics you are well understood. A typical example would be say an aileron on a wing where you want to look at the dynamics of the aileron but you don't want to have the whole aircraft. So you simulate the mass and the dynamics of the aircraft. And you actually put the small component in the wind tunnel and you couple it up via actuators and sensors to your software system which simulates the large object which you want to put in the loop. So interested in how you go about performing such tasks and getting reliable predictions from it. I did that for a number of years. I'm slowly moving away from robotics and more towards control systems and devices.

Peter Asaro: Is that in one of your main projects today?

Ron Daniel: What I'm working on today is what I've continued with my research on stereo. And what I'm interested in in stereo is similar to what I'm interested in in robotics, that is what determines the ultimate performance of a stereo system? And the sort of stuff that I do is basically highly calibrated, very accurately calibrated stereo systems. I'm interested in knowing what statistical tests you can perform on algorithms such that you can reliably predict what the stereo system is going to do in the field on real images rather than the test sets that people tend to produce. There are a number of test sets for testing stereo algorithms. Most of the test images are not – I would say they were not particularly typical of the sort of images that I would see. And they would not be typical of natural images and they have different statistics. For example, supposing you were interested in doing navigation say the stuff that Paul Newman does using stereo. Then testing a stereo algorithm based upon many of the standard image test tests would not tell you much about how your stereo algorithm would work in the wild because the images from the wild have many challenging features that are not present in many of test sets. I've got a Ph.D. student at the moment looking at how do you generate synthetic images which have the statistical characteristics that you would wish to test your algorithm so you can match your test images to the specific tasks that you want your stereo images to perform. There's no such thing as the sort of general image. And many of the stereo which is absolutely crazy for my application so to be called stereo image would have a bright orange teddy bear against a bright green background, against a bright blue something or other with very clear surfaces, no fuzzy stuff, whereas if you're looking at real image it tends to be various shades of gray. Lighting tends not to be too good. You have problems in things like repeating texture and all sorts of issues that these test images tend to avoid.

Peter Asaro: Could you tell us how robotics started at Oxford?

Ron Daniel: Well, if you go back into ancient history in robotics in the U.K. there was something called the Robotics Initiative which was run by somebody called Peter Davey via what is now called the EPSRC but that is the government funding agency for research within the U.K. for robotics type research. This was a program where academics was supposed to work with industry to develop robotics applications of interest to industry. And Peter Davey was in Oxford and he had a project which was robotic vision guided arc welding. And so with a very early, I think, it was an Acer [ph?] robot run by I think it was... some sort of small Vax driving this thing. Basically it was a computer about this big with a disk drive in it which was one of these big sort of like a sort of sauce pan upside down, sauce pan disk drives which you had to pump like that and could store something like ten megs on one of these disk drives. It was a big disk drive for those days. And this was running real time control of basically an old CCD camera, very old nowadays which performed all of the visual processing in a delay line. So you ran the image string through a delay line. Tap the delay line and you performed convolution using a tapped delay line because the computers just weren't fast enough to do convolution. So we had to work out ways of doing convolution. That was done using the tapped delay line. They were in a research group doing this arc welding which started in around about 1980, 1981. I came to join basically with that research group but working with the CEGB on in reactor tele-manipulators in 1982. In fact, February 1, 1982 I came to Oxford as a research fellow at Sir Edmond Hall working with the CEGB. And Peter Davey was running a research group in the department of engineering science on arc welding. He spun out a company called Meta Machines which then went on to sell these arc welding machines to the automotive industry. And before they did that I sat down with Peter Davey and talked about the future of information engineering and software engineering as an engineering subject inside Oxford. And we sat down and spoke and a strategy came up to try and get a chair in information engineering at Oxford. And I was sent off – on one of my conference trips to visit somebody called Mike Brady at MIT to discuss with him whether or not he'd be interested in coming to Oxford and that this had a lot of support from the vice chancellors and so forth. He was approached and headhunted for the chair of information engineering. And from the nucleus we started the research group in robotics. Initially it was myself and Mike Brady and Dave Forsyth as his Ph.D. student and two Sun workstations. So that was the initial robotics research group. But it wasn't called a robotics research group because we didn't have research groups in Oxford and that was not a term used. And I had a Ph.D. student at the time that joined in 1982 called Vaughan Mitchell and he said to me, "Why don't we call ourselves the robotics research group?" There weren't any research groups in the department so we sat down and we designed a little sort of letterhead which we could reproduce in a laser printer because then we had a laser printer, one of the earliest laser printers that I saw. So this is this novelty to actually have a letterhead. So we do this little drawing of a little robot drawing out the Oxford skyline and we called ourselves the robotics research group, the three of us. And that was very much frowned upon in the department because there weren't any research groups. People didn't call themselves research groups. Very modern, the idea of having research groups. Now everybody has research groups. Everybody is in a research group. We were the very first people to actually call ourselves a research group. Mike joined the robotics research group and created the nucleus. He approached a number of people like Andrew Blake. Kevin joined us. Dave Murray, Hugh Durrant-White. I'm trying to

remember the names of all of the people that have come through. Guy Scott came through. Dear me. A long time ago now, this would be 1986 or '85. It might 1985, I think, when Mike arrived. I've been running the group – by then I had been successful in getting grant income. It was a bit embarrassing. What I did is I went to see the EPSRC and said well can I use my then to be awarded research grant to do the controller compliant manipulators to put Sun equipment for the lab? And they agreed and so I managed to purchase something called a data cube. I don't know if you know what a data cube was, again, very early computing device that could do 64 multiplies at once. It could do convolutions and this thing was for doing semi real time image processing. And the data cube, having a data cube was a big deal. I bought a data cube on my project to look at the control and comply arms. This was the early work, the early stuff used by Mike, Dave Forsyth and so forth whose early work was on, I think, it was color constancy or something. He was doing this research and I had a number of Ph.D. students coming through. Then the research group just exploded and we got money for doing autonomous guided vehicle control. I'm trying to remember the names of everybody that came through. Who else was there? Jim Allen. He was a research fellow. There's too many. We had quite a lot visitors come through, quite a few come through. Most international researchers came through. And then Andrew Zisserman came.

Peter Asaro: Were there more from MIT?

Ron Daniel: Mostly visitors came through. Not more from MIT as faculty members.

Peter Asaro: <inaudible>.

Ron Daniel: I don't recall any from MIT coming through. A few went off to join like John Leonard. I remember John Leonard turning up and his early work on something that became known as slam. This is this simultaneous navigation. He worked with Hugh Durrant-White on this stuff called slam. They had these little AGVs that would build maps and navigate around. I used to have a little bit of fun at times trying to compete with him with undergraduates. For a joke one year I put little – we have these final year projects that one undergraduate performs. For amusement I got one undergraduate to build an AGV based on a penman plotter. I don't know if you know what a penman plotter is. It was little turtle that could be programmed using turtle programming which would move around like a little robot and you would be able to put pen on the paper and draw. So we had this penman plotter. And we had a visitor from Japan called Masayuki Inaba who brought something called a multi window vision system with him which was a video system which could extract a small window, a number of small windows and perform computationally intensive processes on a very sort of focused component of the image. And so I got this undergraduate using this multi window vision system with a little cheap camera sitting on top of this penman plotter. And I had a table top model of a factory with a little cardboard models of machine tools. And we had this penman plotter navigating its way and

doing slam using this multi window vision system as an undergraduate project, which was quite good fun to demonstrate somebody doing that. So a lot of our fun projects we tend to do with undergraduates. I had a couple of students come across from Princeton as the undergraduate exchange students. And I got them to build a little walking insect similar to ones that are built in the States based upon radio controller surveyors..But the interesting thing they day was we got a line camera and we managed to using etching techniques from the physics department to attach fiber optics to each single element of this line camera. And these fiber optics were taken off from the line camera and poked into small holes drilled into a ping pong ball. So we had a ping pong ball which with about 30 receptive fields on it which acted like a compound eye. These students built this compound eye. Stuck it on top of this little insect and we built a little micro controller which we looked at this compound eye and made decisions about whether or not it was safe. And this thing was sort of navigating around the world completely built by two undergraduates very little help from me. Quite amazing what they achieved. And we had ways of communicating with it so you'd send it a signal and the way it sent back what it found was by using bee dances. So what we got the insect to do was to do the bee dance. We had various sort of wiggles and hopping up and down and raising a leg which would mean certain things. So we developed a little vocabulary for this little walking insect which was similar to what a bee would use because I'm a beekeeper so I have an interest in bees. We had this thing doing bee dances. And so if you got near it it would be frightened and say I'm frightened and scuttle away and try and find a dark hole to hide in. It had all of these performances based upon what was called subsumption architecture. They had a very, very, very elementary version of the subsumption architecture leaked into this little segmented eye in which this insect could act autonomously in the world and sort of trundle its way around. So that little insect went back off to Princeton when they returned to Princeton as final year students. That was quite good fun. So a lot of what we do is basically fun with students. They love building the robot systems.

Peter Asaro: You mentioned the visitor from Japan. Were there other visitors that you've had over the years?

Ron Daniel: We've had many people come through all of the time. Masayuki Inaba was somebody that stayed for about three months. We got the funds for doing that.

Peter Asaro: And at the time the robotics research group was starting up here, what else was going on around the United Kingdom in terms of robotics?

Ron Daniel: Yes, that was interesting. Yes. I remember something called – there was a thing that Mike Brady was involved in called the U.K. Robotics Research Center in which the larger research groups wrote research proposals to be considered as a research center of excellence in robotics research. I remember Edinburgh University writing one. Oxford wrote one with AA. And then we all sat back with bated breath waiting for the announcement of who had won this

competition. And it turned out to be Salford University which was surprising. The U.K. Center for Robotics Research was placed at Salford University in a sort of business park in Salford. And we must have had one of the biggest research groups in Europe in Oxford. We had huge depth and strength in computer vision. We had Hugh Durrant-White doing autonomous navigation. Mike was doing his AGV stuff. I was doing stuff with tele-operation. Andrew Blake doing computer vision. Andrew Zisserman. We had many, many people here doing all sorts of things. And then Mike moved off into computer vision and medical imaging later after that. He sort of left the robotics field and became more interested in medical imaging. That's after he sort of visited France for a while.

Peter Asaro: And in terms of the relationship between British robotics and greater European and American robotics, would you say there was more influence from roboticists in continental Europe or in the United States?

Ron Daniel: I think most of the influence in Oxford came from the U.S. I think the U.S. had the strongest influence. There was a lot going on Europe. There's a lot of influence on the vision side in Europe. But most of the robotics, pure sort of stuff that influenced me was mostly Japan and the U.S.A. It wasn't Europe. There was a small amount of work in the Netherlands which was good. Most of the really good stuff was in the States in Japan in robotics. But the computer vision is different. That has a lot of highly mathematical computer vision stuff that's come out of France and so forth, various other research like that.

Peter Asaro: Which was the Netherlands group?

Ron Daniel: Delft University. I'm trying to remember the names of people. So long ago now. The name is on the tip of my tongue. I can't bring it to the front at the moment. But you know what it's like I'm getting old. University of Leuven.

Peter Asaro: Yes, what's his name? It's difficult to pronounce. Bruyninckx?

Ron Daniel: Not him, before him. His supervisor. They were particularly interested in describing force control processes. So languages for describing assembly processes and describing force control algorithms.

Peter Asaro: <inaudible>

Ron Daniel: No. The name is gone. I know him very well.

Peter Asaro: Cincinnati Micron machine <inaudible>.

Ron Daniel: Right, yeah. But they were mostly influenced by the work by Matt Mason. Matt Mason wrote a very influential thesis called “Pushing and Shoving” which was about moving objects around and that was extremely influential on the force control community. People became very interested in force servoing. And there was another guy who wrote a really good book called “Control of Machines with Friction” who went and started a spin out company and disappeared from – Armstrong I think his name was. He wrote a really, really good thesis on machines with friction, which also was very influential. But, again, that’s the States. Matt Mason and Armstrong. Oussama Khatib another influencer. Lots of people.

Peter Asaro: In the grasping area, Ken Salisbury.

Ron Daniel: Yeah, I did very little work on grasping. There’s the Salisbury Hand. I looked at some of the kinematics of grasping. I had somebody work with me that did work in that area but I moved away from – I didn’t work on the control of hands. The only work I did on hands was I did some work with a guy from South Hampton which was a hand called the Nightingale Hand, Peter Kyberd and then moved off to Canada, I think. So he worked with me on the design of intelligent hands. So we looked at ways of linking the sense of touch to strategies to grasp so the hand would automatically grasp things if it sensed they were slipping. So that was the late nineties I should think. But I didn’t do much work in that area.

Peter Asaro: What do you see as the biggest problems that robotics is facing, the most difficult challenges?

Ron Daniel: It’s difficult for me to say because I mostly worked on tele-operation and then moved across and I’m now tending to work in harder in the loop simulation. And I’m now working on applying these type of techniques into mechanisms for use inside jet engines. I’ve moved away from robotics and tended not to work in that area for a long, long time. The sort of things that I found frustrating about robotics was the lack of self-discipline regarding experimental design. That is you couldn’t easily utilize the results that are being published because many of the experiments were irreproducible. What was frustrating is somebody said I’ve built this machine and here’s its step response. What can one tell? Part of the reaction was that I was involved in early moves for something called experimental robotics which was a series of conferences in which you try to concentrate on experiment. I think experimental design for developing the science of robotics, I think, is a key challenge. What experiments do you need to perform? How do you design an experiment? And how do you ask questions which are well posed scientifically rather than being involved in what I describe as phenomenological research in which you report on what some mechanism did. I think it requires more thinking about what experiments do you need – let’s take what I’m doing at the moment stereo. A typical stereo

paper will consist of I've developed this algorithm. I've tested it on this test set. Gee, it works well. Isn't that neat. Whereas, if you were to say to look at the physics literature they don't do that. What they would do is they would describe very carefully how they would design an experiment to isolate the issues that they're trying to investigate. For example, in stereo you might say I'm interested in designing an experiment to test the impact of the noise on the algorithms. How do you describe this? What experiments do I need to perform? What are my controls? What conclusions can I draw? In what way is this applicable to – how extensible is this? How can I extract the maximum out of information about the problem that I am trying to solve? I think a major challenge to roboticists is experimental design. So I've noticed it with things like autonomous guided vehicles. It's very entertaining in public. The public may be very interested in seeing say a video of a vehicle trundling around some environment. But how much effort has gone into the design or the experiments to test that system? And how does that stack up against, say, what someone would do to design the experiments at CERN? You're looking for the Higgs Boson. They didn't just sort of glue together some accelerator and turn it on and say ooh, what happened. They spent years thinking about the design of the experiment. Any physicist will spend a lot of effort designing the experiment. What's the experimental design? Even in the softer sciences such as psychology experimental design is crucial. I don't know if you agree with me at all on that.

Peter Asaro: Absolutely.

Ron Daniel: To me I think that's the big challenge for robotics, the robotics community is to come up with a science of robotics. So this my little – the little thing I did on tele-operation which was to look at well okay let's have an abstract tele-operator, what does that tell us about the tele-operation system? And that very simple thought experiment just came up with the idea of momentum transfer. Once you've got that, you can do some experiments to illustrate this or the work I did on phase delays, postulate a mechanism that's causing this, build a piece of experimental equipment which illustrate this thing. Have a hypothesis, make some predictions, do an experiment. Is your hypothesis supported? What statistical bounds can you put on the parameters that you're going to put in? Shannon's Theorem what does that tell us about control systems? What do you predict will happen if you try and exceed the bounds on the channel capacity of this sensor? How would you perform an experiment to try that out?

Peter Asaro: What's your advice to young people who are interested in a career in robots?

Ron Daniel: Be prepared to go and visit a number of places and be prepared to work in a number of labs. I think one of the mistakes I made was I stayed in Oxford for too long, I think. I didn't get enough experience of other people, working in other people's labs. If you're interested in robots go and find some really neat projects going and be prepared to move to go and work in those projects, I think. Don't sort of sit at home. It's quite easy to stay, for

example, in the U.K. If you're a U.K. student and you're interested in robotics be prepared to go to the States. Study in the States. Go find where the good people are.

Peter Asaro: Great. Is there anything you'd like to add?

Ron Daniel: Not really. I think I've said my bit.

Peter Asaro: Thank you very much.