An Interview with

ROBERT L. SIMPSON, JR.

OH 187

Conducted by Arthur L. Norberg

on

14 March 1990

Washington, D.C.

Charles Babbage Institute
Center for the History of Information Processing
University of Minnesota, Minneapolis

Robert L. Simpson, Jr. Interview
14 March 1990

Abstract

Simpson discusses his education in computer science and artificial intelligence and work as a program manager in the Information Processing Techniques Office (IPTO) of the Defense Advanced Research Projects Agency (DARPA). Topics include: Simpson's data analysis work while in the Air Force; the work of Ronald B. Ohlander at IPTO; the DARPA contracting process; program manager interaction at DARPA; and the management of DARPA. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DARPA) on the development of computer science in the United States.

ROBERT L. SIMPSON, JR. INTERVIEW

DATE: 14 March 1990 INTERVIEWER: Arthur L. Norberg

LOCATION: Washington, D.C.

NORBERG: Can we begin, Bob, by my asking you to tell me something about your personal background, both in terms of education and in terms of military services?

SIMPSON: If you want to start with education, I got my undergraduate degree from the University of Georgia. I studied mathematics and education.

NORBERG: What period is this?

SIMPSON: This is 1964 to 1968. It was during, as you might recall from your history, the Vietnam era. And I'm not unlike a lot of people in my cohort that were faced with the dilemmas of trying to get an education and trying to figure out what to do in a post, I guess, baccalaureate degree state, in terms of employment and military service. In my particular case, again, I don't think I'm all that unusual, I looked at my options early on in my undergraduate time and saw that I could either do what a number of other people I had seen do, which was graduate from college and go into the Army and go to Vietnam, spend two years, come out of the Army and then continue on with whatever I was going to do. At the time I was under scholarship from the state of Georgia that was providing me support so that I could become a teacher, secondary education. So it was my intention at that time to go ahead and become a teacher. During my sophomore year, it became more and more evident that I didn't particularly care for the track that had me going through, completing my education, going to Vietnam, because that was what was happening literally to everybody, unless you became a conscientious objector and went off and evaded the draft. So I certainly, because I'm where I am now, I think it's obvious that I didn't like the alternative of not fulfilling my service obligations. So once I made that decision, I decided then to look for the best alternative way of satisfying my military obligations. I approached the Air Force at that time about competing for a position in ROTC at the University of Georgia.

NORBERG: This is in your sophomore year?

SIMPSON: In my sophomore year. I did that and competed successfully for entry into the advanced ROTC program, where I spent my junior and senior year, in addition to my other studies. I went ahead and completed my undergraduate degree, and I accomplished all the requirements to get my teaching certificate in the state of Georgia. But then went into the Air Force. I got my commission at the same time and went into the Air Force in 1968. My intention at that time was to satisfy my obligations, which would have been a four year tour in the Air Force. My feeling at the time that a trade-off of doing four years in the Air Force as an officer was better than the trade-off of two years in Vietnam in some role in the Army. So that was the point that got me into the service. I spent my initial four years - quite surprisingly, to me - I enjoyed it. My initial approach to the military service was, "This is something I have to get through." But I found the problems that I was confronted with, because my undergraduate degree is both mathematics and education, I was put into positions of sort of analytic support. It turns out I was doing analysis for the military airlift command trying to figure out how to charge for airlift successfully to recover the Air Force's operating expenses.

NORBERG: Who was being charged, other services?

SIMPSON: Who uses the service. So the Army has to airlift, they have to pay for it, right? It's just like a trucking company. So I found this an interesting analytical problem, and so I got into that aspect.

NORBERG: How was this problem investigated?

SIMPSON: Gathering data, trying to do a reasonable accounting of the costs involved, and then trying to see how to effectively amortize this. This is what's known as industrial funding within the services. It's fairly common practice now in a lot of places. But at the time the Air Force and the rest of DOD was undergoing the McNamara revolution. So there was the new ideas of the five-year defense plan, the new modes of accounting, the major defense categories. So there was a big upheaval, if you will, in the way that the services were having to think about and explain, rationalize, justify the various costs of things. So I was sort of in this turmoil as a brand new second lieutenant. I

was assigned to a headquarters of the military air lift command. So I was a junior officer thrown into a predominantly senior officer group of people that were wrestling with revolutionizing - in their mind - the accounting mechanisms, trying to mechanize some of this so that they could provide the right kind of data to the folks in the air staff and to the rest of the people in OSD.

NORBERG: While you were there were they successful in mechanizing the process?

SIMPSON: Not too surprisingly, if you know the state of computing in the late 1960s, as well as the fact that they were probably maybe ten years behind the state of the art, (maybe not that much) it was quite an ordeal. That was, again, my first big introduction to computer users and the issues of trying to do ADP. I had had some introduction as an undergraduate to computing, but not extensively. The bottom line was that I not only found that I enjoyed the sort of intellectual challenge, the other issues surrounding trying to grapple with these problems as well as the mechanization aspect. I found Air Force camaraderie to be very enjoyable. I found Air Force life to be not as I had previously imagined. So I became quite enthusiastic. The Air Force, on its part, continued to shower me with encouragement. After two years there, I received several accolades from my superiors, and they recommended me for a graduate program. They offered to basically send me to get my master's degree in computer science, if I wanted to. Basically the invitation was you pick your school and get admitted, and we'll send you there. So I did. I picked Georgia Tech. I'm from Atlanta and I went back and in a year got my master's degree. Then I was reassigned to a position in the Air Force that had me dealing with laser weapons testing and nuclear affects analysis. This was in New Mexico. That work I found again another different world that I was not aware of and also challenging from two aspects: the relevance and importance, because at that time laser weapons work was very secret and so the activities surrounded sort of building devices, testing them, understanding their capabilities and limitations were very intellectually stimulating. The nuclear affects work I found not quite so interesting but also challenging in an analytic sense. Other than the usual sort of Air Force duties, my time... This took place in the early 1970s, and it was in Albuquerque, New Mexico. It was a period, also, that made me begin to think that maybe I wanted to spend more than just my initial four years in the Air Force. From there I was reassigned to the Pentagon as a young captain, which once again put me in sort of another environment that was reasonably challenging. Again, as a junior officer I

was put into a position of where I was having to interact with a number of senior people in the Pentagon. It gave me exposure to other, both computing challenges, both in terms of... I was involved in some of the bomber simulations for the B1. This was during the Carter years where the Air Force was arguing why we need the B1, and Jimmy Carter was saying, "Hell no, we don't need it." He, in fact, canceled the program at the time. But my job was to help support the computing on which the simulations were run in order to gather the analytical evidence to make the case.

NORBERG: Was that considered a classified project then?

SIMPSON: The fact that it was done?

NORBERG: Well, no, I mean the project itself, because I don't want to ask you any questions about it if it was classified.

SIMPSON: I'm not going to say anything that's classified.

NORBERG: That's fine. Let me ask what I want to ask, then you can decide. What sort or equipment was being used?

SIMPSON: At that time IBM 360/75, and around the Pentagon at that time I was Mr. IBM 360/75. (Laughter) Bob Simpson was the guy that was responsible for the care and feeding. I did things like I made unconventional improvements to the operating system and to the hardware configuration to enable these very large bomber models to run as quickly and as efficiently as they could.

NORBERG: What does that mean, unconventional?

SIMPSON: Unconventional in the sense that I took a standard 0S360 operating system and made modifications to it that streamlined it and made it more effective. You know an operating system takes some resources. So I made non-

standard modifications to the memory because standard, at that time, IBM supported 360/75 was one million bytes of total memory. Pretty pedestrian by today's standards, but at that time that was a lot of memory. I went about and procured a special engineering change that expanded that to two and a half million bytes of resident memory. So that was at the time a "risky engineering change" that involved a fair amount of work. The only point being that once again, I got an opportunity as a young officer to do things that gave me a fair amount of responsibility and challenge. For that time at the Pentagon, I again continued to receive very favorable responses from my superiors and was promoted early to major which again sort of solidified my view that maybe I was right for the Air Force, because the Air Force was giving me the right kind of positive reinforcement. After four years at the Pentagon - we're now at 1977 - I was selected for a special assignment. This special assignment was to take me to San Antonio, Texas to the Air Force Manpower and Personnel Center. It was there that I was put in charge of the personnel actions for all of the computer scientists in the Air Force, all the computer personnel, that's not [only] computer scientists, that's also ADP people. So my job was to approve acquisition, that is all of the people that came into the Air Force, and into this technical specialty, had to be approved by me. I made all of the assignments, choosing where to place computing officers around the Air Force, the notion here being you want to grow people in terms of technical skills and managerial skills as they progress up their military career. So there was a career management process. I was responsible for selection of people for fast-track programs, for military education, air command staff college, air war college, and special assignments of various sorts. So I did that for three years.

NORBERG: Were you also giving some thought to what the Air Force might need in terms of computing capability, not only in terms of equipment and systems but also in terms of personnel?

SIMPSON: Yes. It was very much the case that during my time there we wrestled with what we... we used terms like "technical obsolescence" of people as they make their way through their military careers until you get to some point, as a senior Air Force manager, where you are some generations out of date with technology. This was a continuingly frustrating issue, especially in the computing area where you're looking at about a five-year half life of technology. So we routinely realized that you bring in young officers as second lieutenants that are reasonably up-to-date in technical terms, you put them into organizations, and by the time they become lieutenant colonels, for example, they

are obsolete, unless there is some program to keep them sort of up-to-date. And some people by their nature will do that. But the Air Force in general did not provide mechanisms to make that easy. It wasn't a natural part of the career progression.

NORBERG: I would assume that would be true of all the services.

SIMPSON: It is true of all the services. So in that sense... well the statement I made is generic to all the services. So I was, as is my wont, pretty vocal with my management about this problem. So I along with other people in technical areas, not just computing but all the other engineering and technical areas, voiced similar concern for the technical obsolescence of the senior office corps of the Air Force. I succeeded in getting a new program initiated that was responsible for selectively choosing Air Force officers for advanced education. The notion here basically is you find people that you think are your future leaders and give them a chance to go back to school and get the next level of technical retrofitting. That program was the "commander's option." A commander had the choice of sort of pointing out the people in his organization that could go back for this special education. As I indicated, I spent the years 1977 through 1980 involved in this sort of personnel management, personnel planning activities. During that time I had been promoted to major, and because I was early promoted to major, one of the things that goes along with that is an obligation to go to advanced military education. So I was selected to go to the armed forces command staff college, down here in Norfolk. So I was on my way there when one of the people in my organization that had been selected for this "commander's option" to go get an advanced degree, instead chose to leave the Air Force. So the organization had a situation where it had a "slot" for an advanced education position, and nobody to fill it. So much to my surprise, my commander said, "Well, Bob, you're one of the people that felt this was a good program, why don't you go do it." So I was sort of hoisted on my own petard; I didn't have much choice. So, in fact, I was reassigned from a military education program to an advanced civilian education program.

NORBERG: What did you think about that at the time?

SIMPSON: Well, to be quite honest, it was a military kiss of death, because I knew the statistics that said: number

one, Ph.D.s in the service, Air Force, don't fare as well in promotions; number two, that going away from the Air Force for a period of time was a negative factor on promotions. Also the fact that getting a Ph.D. and practicing a technical position was a third strike against one. So I was well aware that by taking that choice I was taking a very high risk path. But nonetheless, again, I felt like not only was it a good thing for the Air Force but that it was a good thing for me. Even if it didn't help me in my Air Force career, it might help me in my post-Air Force career. So, at that time, it crossed my mind that that was useful.

NORBERG: Did you give any option to saying no, that you would rather go to the Military Staff College?

SIMPSON: Oh, yes. I considered that seriously. Again, because you have to remember, for these previous twelve years, I had been very successful. The Air Force had been very lavish in its praise. I had been early selected to major. I had been put in special positions, positions that allowed me a fair amount of latitude and responsibility. That was quite rewarding, so I felt like, "Gee, the service was wanting me, so why shouldn't I do things to continue to enhance that." But the bottom line was once again I was told, "Find a school that you can get into the Ph.D. program, and we'll send you there." So because of the shortness of time, I basically approached Georgia Tech for readmission. At the time Georgia Tech said, "Well, you haven't submitted the usual sort of Ph.D. application. We'll admit you as a sort of returning graduate student under probation." So I had to prove my worth. So in 1980 I entered the Ph.D. program at Georgia Tech and from 1980 to 1985 basically did the course work, the research proposal, and research, wrote my thesis.

NORBERG: Let's not be quite so quick. What was the Georgia Tech program like when you returned to it in 1980?

SIMPSON: Very different from the program that I had left in 1971. The program in 1971 was I guess more theoretical, leaning heavily on the information sciences and less on the computer sciences. So when I came back there had been sort of a revolution in that there was a new department chairman. Ray Miller was the new department chairman; I don't know if you know Ray. But anyway he had recently arrived from IBM and was in the midst of converting the department over into an ACM model department. I was sort of one of the... there was a new crop of graduate

students that were sort of brought in with this new model. There were people, for example, that had completed their master's at Stanford and were in for the Ph.D. program. So it was a quite competitive group of graduate students, and I felt it was a nice stimulating environment. So it gave me a chance to get more into mainline computer science. An example being up until then there was no real commitment to do research in artificial intelligence. But a brand new faculty member at that time was there, and that dovetailed with my interest because I was very interested in advanced applications of computers. So I gravitated to a new faculty member who was looking for graduate students, and it also made it attractive because I was a graduate student with my own funding.

NORBERG: Let me be clear about something. I take it then that when you went there, you had not selected a research area that you were going to work on?

SIMPSON: Not at all. I was an "undesignated."

NORBERG: Did you give any thought to other universities?

SIMPSON: Only briefly. Because, remember, it was sort of a last minute decision that I was going to go to graduate school. I talked to the University of Texas because I was in San Antonio at the time. And it occurred to me that I could just sort of live in San Antonio and commute to Austin. So that was one option. Georgia Tech was the next option.

NORBERG: What time of year are we talking about?

SIMPSON: Spring. So it was maybe April. And I was talking about entering in the fall. So this is sort of behind the normal period of applications and acceptances. I believe the only two schools that I really gave a lot of thought to were Texas and Georgia Tech. My family voted overwhelmingly Georgia Tech. My family is all from Atlanta, so both my wife and children wanted to go back there, and our parents and all were more enthusiastic about that. So for me I didn't spend too much time on it; I made the decision. And it was easiest, frankly, to get back into Georgia Tech,

because again I was in a readmitting situation.

NORBERG: Back to the research program. What sort of research program was this new faculty member involved with?

SIMPSON: My thesis advisor was Janet Colodney (?). She's a student of Roger Shank. Do you know Roger?

NORBERG: Yes.

SIMPSON: So she brought the conceptual dependency perspective on AI to Georgia Tech. That was the paradigm of research that I was educated under, and into. Because it was a new area there at Georgia Tech, there was a great deal of struggle building up the infrastructure, the computing support and the undergraduate course work, to teach the intro to AI courses. So we prototyped a lot of those, choosing what area in AI that Georgia Tech might be notable for, because you don't want to plow the same ground everybody else is plowing. So there was a fair amount of struggle trying to define what would be a Ph.D. in AI at Georgia Tech. So I was the first graduate student to plow through this. Everything was being prototyped with me as the example, which I think in retrospect made my task much harder than it needed to be. In some cases, you can come in and there's an existing project that you can fit into that sort of defines what an interesting research topic would be within this. But there was none of that. So we fumbled around for a long time trying to define our identity.

Janet's research was in the area of dynamic memory, which is sort of a cognitive model of human memory processes. As it turned out, my contributions were of looking at the combination of this cognitive model of human memory with normal AI or cognitive science models of human problem-solving, and realizing that the two of these things had sort of evolved as separate parallel things. And it was possible to postulate the union of these two things in sort of a protocol of human problem-solving behavior. That led to a view of problem solving that was - I don't want to say revolutionary - but just different, because it argued that there might, in fact, be active memory processes that interject themselves periodically into the problem solving process. So that it isn't sort of this straightforward algorithmic kind

of process, but it is in fact some reminding that makes suggestions to how to think about a problem, how to formulate it, how to decompose it, how to estimate solution paths. So this architecture came to be known as case-based reasoning. At the time, in the early 1980s, there weren't many people looking at this. That's not too surprising, because, again, the notions of dynamic memory that my thesis advisor had done were not widely understood or appreciated. So there weren't many places that were doing it. The net result is that the thesis work let us argue about some of these new ideas and new ways to think about this. It seemed to receive some appreciation. So it was with that kind of perspective, because in doing this and because I had to fumble through all of these procedures, it gave me an opportunity to be reasonably read in AI. So when I came to DARPA I felt like I was moderately read in AI in general. So I came to the job, not surprisingly, with an already established perspective.

TAPE 1/SIDE 2

NORBERG: Who else was working on case-based reasoning in the middle 1980s?

SIMPSON: At the time, there were only a couple of other places that were doing it. And at the time, even the term case-based reasoning was not that prevalent. I think if you do a literary search you'll find that that term, even today, doesn't show up in a lot of indexes. But at the time, Edwina Rissland, at the University of Massachusetts, was doing work that was very much of the same spirit. She used terms like "reasoning from hypotheticals" as the way to describe. But her process model was very much akin to the spirit. And, of course, other Roger Shank students were also looking at the same sort of memory-based processes in various problem solving contexts. Other than that, perhaps, Mike Liebowitz at Columbia was another one. Literally that just about covers the spectrum. Since that time, there have been more people involved. I am also happy to report that through my arguments on behalf of this way of structuring and thinking about problem-solving I've been able to nurture this area while I've been at DARPA. So now there's a fairly reasonable body of people who identify themselves as being interested in the theoretical and technical problems of building case-based systems.

NORBERG: I'll come back to that in a moment, but getting back to the middle 1980s before you left Georgia Tech the

second time, or is it the third time, third time, isn't it? What was the reception of case-based reasoning at the time, the kind of work that you people were doing? Was there anybody particularly pleased with it? Anybody out there objecting?

SIMPSON: I think it was met with apathy. (Laughter) I think the best way to describe it was it was neither greeted enthusiastically nor condemned. It was just sort of there.

NORBERG: Let me ask the question differently then. Can you make any statement about the kind of reasoning that other people were doing that was getting the attention at the time and how would your work have fit into that?

SIMPSON: I guess in some sense it would be wrong to think of it as contradictory. What I always viewed it as is complementary. Because one of the criticisms, even today, of AI is that often most systems have one way of inferring reasoning through a problem. So if you think about the normal sort of chaining of rules to this, so there's only one process, being this inference chaining of rules, to find a solution or to estimate a solution. Well, my model was that there are several methods to get there, and part of what the reasoning process does is it tries to use multiple techniques depending on the context. So case-based architecture basically says that if you have no cases to reason from, you can use from scratch, a weak heuristic method to approach the solution. But if you have cases they become strong influences on how to estimate solutions, how to estimate partitioning, strategies, or other techniques so that it ends up being complementary methods to let you, for example, reason in a variable way about problems based on how much you know about it from the past. So it's in that sense that people... the mid 1980s were a period of great growth in AI, and I think this was just sort of part of the growth. My view is that, again, the 1980s were a very great period for AI. Thanks to DARPA and other funding agencies lots of money was poured into it, so it was a great expansion of investigation both in breadth and depth. So my view is that case-based reasoning is sort of another one of these branches.

NORBERG: One could also interpret your remarks to mean that there really was no internal criticism in the field either. That since there was sufficient funding around, and since no one was really particularly excited about what you were

doing in terms of it being right or wrong, then there is just a laissez faire attitude that led to each his own.

SIMPSON: Sure. Let the flowers bloom...

NORBERG: ...thousand flowers bloom, sure.

SIMPSON: No, I think that's not an unfair characterization. One could argue that even still, though AI is an emerging area, and too much criticism may at some point may in fact hinder yet unexplored avenues. I mean there is that side that you could argue about. Depending on which perspective you want to take, too much or not enough criticism is a bad thing.

NORBERG: I'll come back to that too, because I have an opinion about that, which I'd like to offer to you and see what your response is when we get to the end. Who was supporting the work at Georgia Tech at the time?

SIMPSON: Remember I was an Air Force officer.

NORBERG: Yes, but who was supporting Colodney?

SIMPSON: My resources were completely provided by the Air Force, so in that sense I provided my own support. Janet was supported primarily by the ARI, the Army Research Institute, the Army Research Office, and NSF. So those were her primary sources of support.

NORBERG: Interesting. Do you remember what the Army's interest in this was?

SIMPSON: They were interested in it as mechanisms for perhaps aiding training. The view that Janet's model of memory was potentially useful from a cognitive science perspective, helping to understand how better to train soldiers in order to affect their retention of various things. So as her research and my research sort of evolved, she

recognized that her model of dynamic memory in the context of the sort of architecture that I was interested in, this problem-solving architecture, gave more context and provided other relevant settings for describing activities. I think it is safe to say that Janet has completely shifted now into a mode of looking at her memory processes in the context of case-based reasoning. For example, one of the interesting areas that she is pursuing now is case-based reasoning in the context of design, so looking very closely at the process that designers go through and thinking about and decomposing and reasoning about designs from the basis of previous designs. So that's sort of a model, and she's investigating both how to build computational agents of that form as well as provide assistance and advice to designers of interactive systems. So that's one area. The other is tutoring as educational aids. Again, the same theme of, if you really believe that this is a reasonable cognitive model of human information processing, then you can use that model to aid the process of transferring information to students. So she is very interested in intelligent tutoring activity.

NORBERG: How successful was she during the years you were there in attracting other graduate students?

SIMPSON: Very successful. I was going to say in my view that was the fastest growing part of the computer science department at that time in terms of graduate students. I think Janet at the time got to the point where she was easily turning away most of the people that were approaching her to do work.

NORBERG: When you were offered the opportunity to engage in this program by the Air Force, were there any conditions specified? Like you need to do it in four years, and you need to be back here, and report to some duty, and so on?

SIMPSON: Oh, yes. The normal Air Force policy is three years for a Ph.D. I accepted the position understanding that policy. Because of the fact that the intervening period between my masters and my Ph.D. was such, basically I had to go back and retake most of the course work. It took me about eighteen months to successfully complete my course work and take my qualifying exams. Once I succeeded in passing my qualifying exams, the program then required me to put together a research proposal and defend it. And it took me another six months to go through and

put my research proposal together and defend it. Once I completed that two year period, that left me a year on my "Air Force tour." I sat down with Janet at the time and looked and we mapped out a strategy for completing my research as my proposal called for. Here's what would need to happen over the next period of time. Janet would not agree to a research program of less than two years in length. Her view was that it was not only inappropriate, it would lead to unsatisfactory research results. And that she would not support my candidacy without that. So I was faced with a dilemma the Air Force only going gave me a year, and Janet was not going to live with anything less than two years. So with a two-year plan that she would agree to, I then approached the Air Force with the research plan and basically wrote a letter to my superiors saying, "Okay, I'm at the two-year point, I'm ready to begin my research. Here is my research plan, and it projects me to pop out with a Ph.D. in two years. I need either your concurrence to add time, in effect waive the requirements, and give me the additional time - or let's stop now. And I have now "completed yet another master's degree", but I've been sort of refurbished in one sense." So I explained to the Air Force that I thought I was a better officer for the experience, but that we shouldn't begin this path if we didn't intend to finish it. Much to my surprise, the Air Force agreed. So I was given an additional two years from that point, which was four years. And I launched into my research, which like most research had its ups and downs, cul de sacs, and other twists and turns. To make a long story short, it turned out that after doing my research, and then starting and doing the writing, it actually was taking longer than I had expected. So I was approaching the end of my two years, and I was able to successfully argue for another three months' delay. Well, I argued twice, I think, for three months' delays. So the net effect was it was four and a half years. Even then, and during that time was when I started actively lobbying for the job here at DARPA. Part of what happened during my research was I discovered DARPA. You see, up until this point, I didn't know DARPA existed. I discovered DARPA when I was here in Washington at the National AI conference. I stopped by the Pentagon to go trolling for jobs, because I had previously been in the Pentagon and I sort of knew my way around. I went talking to various people in the Air Force and OSD. I was in OSD RNAT, I think it was called at the time, and was asking if they knew anybody that was interested in artificial intelligences. I don't know if you remember Edith Martin?

NORBERG: I only know the name.

SIMPSON: Edith was there at the time. And Edith said, "Gee, you oughta go over and talk to these guys at DARPA. Ron Ohlander is the guy in charge of AI at DARPA." So I called up Ron, and told him who I was, said, "Can I come see you?" So Ron agreed, and I came over to DARPA and visited with Ron. I told him that I wanted to be a program manager at DARPA.

NORBERG: Now when exactly was this?

SIMPSON: This was 1984. Ron at the time was in my view, not very interested. (Laughter) He was like all program managers very busy; he had lots going on, and since my projected arrival was, like, over a year away, because this was August of 1984, and at the time, my tour was going to end in December of 1985. No, excuse me... no, that's wrong.

NORBERG: It might have been in 1983.

SIMPSON: Yes, I think this is 1983, and I would have shown up in December of 1984. So he was mildly interested because I couldn't help him any time soon. We corresponded off and on after that. When it came the summer of 1984, June-July of 1984, I once again succeeded in getting a delay to March of 1985. That was when it was time for me to get an assignment from the Air Force. So I lobbied hard with Ron, got an invitation to come up, and I interviewed with Bob Kahn. I seemed to satisfy most of them. Ron, I think, was approaching his foreseeable end of his tour, and Alan Sears had come on board. I think he viewed that it was right for them to try to get that, and so they basically agreed. Then between DARPA requesting me and my lobbying on the Air Force' side to have me assigned, I was successful in getting the position. The Air Force, during that time, had discovered AI, too. So in fact I had several people from the Air Force talk to me about coming to work in various Air Force jobs. In my own survey of what was going on, I was aware of Wright Patterson activities; I was aware of RADC work; I was aware of some interest at systems command headquarters about creating an AI program.

NORBERG: About creating it, but they weren't running one at the time at any one of those places?

SIMPSON: No. So, in fact, I was mildly amused and impressed - I got a phone call from the director of laboratories, General.... I can't remember his name, General Somebody - at that time, who was telling me that I really should stay in the Air Force, and not leave the Air Force and go into DARPA because that wasn't going to be the smartest thing for my career.

NORBERG: What was he proposing you do?

SIMPSON: He was proposing that I either go to Wright - Patterson to head up the AI program at Wright - Patterson or go to RADC.

NORBERG: Why didn't you take either of those options?

SIMPSON: I had had a chance to talk to people at both places, and I had visited here at DARPA. I had some assessment of both the responsibilities, what the job entailed, as well as impact is one way to describe it, what effect I could have. So, in my opinion, there was sort of night and day differences between jobs in the Air Force and DARPA.

NORBERG: Can we talk a little bit then about what you thought you could achieve when you came to DARPA?

SIMPSON: Oh, yes. Number one, it was that sense about... I had been enrolled in research in AI, so I very much wanted to continue to be involved in research. And these Air Force jobs were not necessarily going to be in research. They were going to be in development. Number two, you've got to remember, that was the time the strategic computing program was sort of being started. One of the things that happened when I visited DARPA was I met Craig Fields, I met Bob Kahn, of course, and other people that were involved in strategic computing. Among the things that encouraged me was, "We need people like you that can help us manage, shape, and direct strategic computing." So, I mean, that was my charter. I wanted to come and have some role in trying to make strategic

computing the success that I think it was going to be. So that was the thing I really wanted to do. In terms of effect I viewed myself as a team player on the team at DARPA on this national mission of showing that AI could have a big impact. So that was sort of in my mind what my mission was. That was the way it was when I showed up. I remember Lynn Conway was there at the time, and by the time I arrived of course she had left.

NORBERG: Who else was here when you arrived? Can you just run down the hierarchy for me, in terms of how you remember it? Who was responsible for what?

SIMPSON: In DARPA?

NORBERG: Well, I was thinking of just the IPTO office.

SIMPSON: Well, Bob Kahn, of course, was the director. I guess Paul Losslaben(?) was around at the time. Let's see, I can't remember all these now. Bob Baker was around. Of course, Alan Sears and Ron Ohlander were the people that I interacted most with. Steve Squires was obviously around. Who else? This is terrible; I can't remember all of them.

NORBERG: It's all right; they'll come to mind as you keep talking about the program.

SIMPSON: Also in the intervening period, the split between IPTO and EAO had occurred. So those people, like Jack Thorpe, that I had met originally were no longer there. Craig Fields had moved out of the office and was in EAO. So who else was around? Lynn Conway had already departed. Paul was still there; Dennis Perry had not yet shown up.

NORBERG: How was the program divided then at the time in terms of trying to define your new responsibilities?

SIMPSON: Ron had decided prior to my arrival that Alan Sears should primarily work on speech and, I guess, some of the knowledge-based systems work. So when I showed up Ron wanted me to work in vision and natural language

processing. Those became my first areas, and then Ron encouraged me to structure, organize, and promote new initiatives. So based on his encouragement, I promoted new programs in automated planning and machine learning.

NORBERG: Why those two?

SIMPSON: Well, they were areas that I was interested in. They were areas that it was very easy to make an argument for their importance and utility to DOD.

NORBERG: They were? How? The machine doesn't know.

SIMPSON: Oh, the machine doesn't know? Well, the way that I made the argument, and it is still a valid argument, if you look at all the advanced systems that we were prototyping in strategic computing: autonomous land vehicles, pilot's associates, battle management systems, they all have in their functional core, planners, some thing that manages resources, that worries about deadlines and time constraints, that worries about creating sequences of events to achieve goals. But within the organization of the office and its research programs, there was nothing that was coherent in terms of identifying the basic research issues, identifying the technical technology chunks that could be coalesced and then transferred to applications. As a consequence, I argued for some work in this area. On the side of machine learning, I argued that it was an emerging technical area that needed basic research support as well as some activity in technology enhancement. If you think about knowledge acquisition for knowledge-based systems and also for building systems that were adaptive in their behavior, because they could modify themselves as they got feedback on their effects in the world. That was the rationale. I was moderately successful; today those programs are fairly entrenched, if you will, in the program. So initially coming into the office, the partition was between Alan and me. Alan was speech and knowledge-based systems; I was natural language and vision, image understanding.

NORBERG: Now, can you map that for me onto the strategic computing initiative in terms of the differences between basic research on the one hand, the kinds of basic research you were supporting, and then this 6.2 activity that was

going on in these various other places.

SIMPSON: That's easy. Conceptually, the areas of research were the same: vision in strategic computing, vision in the basic research program. The distinction is sort of an orthogonal one where you look at the distinction - it's a fuzzy boundary - but it's between science and technology, between a focused, targeted recipient of the research results versus a less focused, more diverse recipient of research results. Strategic computing was looking for things that were more focused, or more technology oriented as opposed to the basic program, which was intentionally more generic, more scientific, less directed.

NORBERG: Were you trying to manage both sides of this?

SIMPSON: Oh, yes. And frankly I think that is a very, very smart thing to do: to have a single program manager responsible for either side of this boundary. And in fact I believe it is a smart strategy in the office that you add to that an application. So I think having a program manager who worries about basic issues, technology issues, and application gives you a way to vertically integrate and accelerate technology transfer across those levels. It also provides a mechanism where an individual can make assessments on "What are the important problems that ought to be worked on?," from an applications' perspective, "What things are most likely to have an impact when?," because you can assess the relative progress of getting something up in this pipeline. And you can make sure that it has an impact so that you can demonstrate its utility. So in my view, and frankly that's what is happening now, it is an important linkage. During the early time in strategic computing, in fact, the people running the applications were divorced from those doing the basic research and the technology development. So the last linkage was missing. So, in fact, what I did often times is I established personal and programmatic relationships with applications program managers so that I could work that last boundary more effectively.

TAPE 2/SIDE 1

NORBERG: Now on the end of the previous tape, you just led into a very nice description of the chain from basic

research to applications. I'd like to ask you now to describe for me how you managed that. You just gave sort of a general view of the importance of managing it. But now how, specifically, did you do that? For example, how did new proposals come into the office, what was the evaluation mechanism, where did you get your information to do such evaluations, and so on?

SIMPSON: You have to remember that over the five years, things changed. So in describing it I can either do it by periods, because that was one of the challenges, was how to do what needed to be done and still modify things according to the changes in the Federal Procurement Regulations. The bottom line was that basically a program manager, in carrying out his function, has some overall notion of where he wants to go; both in an application sense, a functionality sense, or in evolving certain technical ideas from basic research into a technology form, or in pushing into regions of technical space that have been unexplored, for example, or need more exploration. So in those general terms program managers in some sense have a model of what they want and what they need. They can convey that either formally through announcements in the *Commerce Business Daily*, announcements at conferences and workshops, or they can announce it informally in discussions with principal investigators or others that come into the office. So those needs and interests are always conveyed, either formally or informally, to members of the technical community. The proposals that come in are of the form either unsolicited, they come in because people that have good ideas are always encouraged to send them to DARPA or they are responses to either formal invitations, like broad agency announcements or other published things, or the informal mechanisms that encourage proposals.

In the earlier part, around the middle 1980s, before the competition in contracting laws were enacted, DARPA operated much more on the informal level. That is, program managers who knew what they needed basically would talk to people to find out, "Is anybody doing any stuff in this area? Could you do something in this area? You can do something in this area? Well, then send me a paper, and tell me what it is that you think you can do in this area." So there was always a dialogue between potential bidders and program managers. That dialogue was an informal way to survey possibilities and to do some level of evaluation. As the competition in contracting laws were enacted and requirements were much more rigorous, let us say, in terms of proof and an audit trail for evaluation, then it became much more formal in the sense of publishing broad agency announcements, encouraging formal submissions,

convening panels of evaluators to provide advice, and showing that procurements were, in fact, legitimate competitions among proposers from all areas.

So during my time at DARPA I've seen it move from a responsibility on the program managers to be not only innovative, aggressive but fair in soliciting, evaluating, and supporting research to a view now that is much more regimented, much more managed, and discourages program managers from being aggressive and innovative in what they do in the sort of informal fashion. I think that, in some sense, that's an unfortunate consequence, but in another sense, it is one way to make the process, if not really more fair, then have the appearance of fairness.

NORBERG: What is the impact, then, on your work, you Mr. Program Manager?

SIMPSON: It adds a great deal to a program manager's work load in administrative activity, because again the requirements for successfully procuring a piece of research is that you have satisfied all of these legal hurdles. So that is the sense in which I said it moved from being an informal process to a formal process. And in my view, if you are going to have program managers who are technically knowledgeable about an area and are charged with achieving some sort of results or who want to make maximum progress in an area, these kind of things are guaranteed to inhibit that. So in my view it is a negative evolution.

NORBERG: Is it possible to get around those problems, though?

SIMPSON: Sure. Remove the requirements.

NORBERG: Well, yes, but that's a trivial solution to the problem.

SIMPSON: No, we always solve problems by changing the environment.

(Laughter)

NORBERG: What I was trying to ask, actually, is what happens in order to get the best research done when you have these constraints?

SIMPSON: You end up sacrificing something. I'd like to think that my colleagues and I never sacrifice the quality of the research. We end up paying the price in terms of amount of work that it requires, added administrative support staff, so mo ney that would normally be able to be spent on research is now spent on support staff. One could argue that there is a cost to the research community added to this. It has delayed onset times from when a proposal is submitted to when the research can begin. So, I mean, the cost is borne in a lot of different ways. It discourages people from coming to DARPA and wanting to be program managers, because they perceive the administrative burden and their flexibility in doing things that are of interest to them. Part of what makes a program manager's job interesting is being able to do something that is intellectually stimulating in your space, that you think will make a difference to the country and to DOD. So it is in that sense that under these constraints you are going to discourage some of your most aggressive entrepreneurial types, because they have sort of less latitude in operating.

NORBERG: Do you have a core of research groups that you depend on to accomplish the objectives of your program? Then, around that, there is this other group of people that we've just been talking about who now are encouraged to participate in the process to increase fairness and to expand activity and so on. Or is this just a general free-for-all of the same kind now that NSF has and does it by peer review, whereas you're asked to make this evaluation?

SIMPSON: I'd say that, again, over the past five years there have been changes, over the past decade, because I can look at the five years before I came to DARPA and observe some of the things. I think there are several factors or several trends that need to be accounted for here. One is that AI as a discipline, as a technology has been exploding during the 1980s. That has meant that, whereas in 1980 there might not be but a core collection of places that had the critical mass of people, good ideas, infrastructure to legitimately bid on and conduct research in the areas of AI or other things that DARPA was doing. I mean, that's not surprising because as a discipline it was quite, it's still quite young and evolving. Again, as a consequence of what happened in the 1980s, not only at DARPA but around the

federal government, the investment went up dramatically. There were, consequently... And more people had been produced from universities, so that the number of people that have good ideas and can contribute in these areas went up. So you had both money available, you had more people available, and consequently the quality could still be maintained and still have more people compete. So from my perspective, AI during this period expanded from a core group of places, the usual - MIT, Stanford, Carnegie - to a number of other places which, I think it's legitimate to say, are second tier. But yet within certain areas the research is every bit as good as research that's done at these sort of traditional centers of excellence in AI. For example, you could support a machine learning piece of research that went on elsewhere. Why? Because the principal investigator and the research group were perhaps people that came from Carnegie-Mellon and someplace else. They understood both the view, the mechanisms for doing the research, the paradigmatic issues, but also knew about DARPA - knew how DARPA wanted to do things, how we emphasized tech transfer, and how we wanted things to have demonstrations in real applications. So there was this education process about how to do business with DARPA, what DARPA was interested in, communicating with program managers, and all of the strange things - the incantations - that were sometimes required, as well as understanding what it means to be a member of the DARPA research community. Because DARPA contractors are in one sense contractors - they get paid by the government to do things, but at the same time they are supporters of the DARPA vision. They promote the DARPA vision. They also help program managers and DARPA articulate and defend the rationale and importance of what we do. So it is in that sense that it is a community. Because, again, there were people leaving Carnegie, Stanford, MIT and forming their own research groups outside, there was a transfer of an appreciation for what DARPA was interested in, how to talk to them, how to understand the mission and the vision, and then how to participate. So that, combined with the competition in contracting, encouraged and provided a mechanism to expand the support. So I'm quite happy that during my five years at DARPA, the number of groups that are "part of the DARPA community" pursuing the vision of AI in development and basic research and technology is quite large compared to five years ago. New groups, new sites.

NORBERG: There are two issues in there, and I don't want to lose either one of them. I'm trying to decide which one to go to first. I think I'll take the easiest one. One of the trends you didn't mention, and I'm wondering whether it is there, is increased resources, or whether there had to be reallocations within the field.

SIMPSON: Oh, sure. Oh, no, no, no, no. All of this would have been very difficult to accomplish if it had been a zero sum game. It was the case that during the middle 1980s through 1987 the resources were generally expanding. I can remember, as a new program manager, I would put together my arguments, my strategy for a new program, machine learning or planning and get a fair hearing, and generally there was a reasonable probability that I would get additional resources to support this. So it was very much the case of an expanding pie, if you will. So my model, at the time, was, in my memory, continue supporting groups that have had a track record of innovation and accomplishment as long as the ideas were good, but also find new people that have new ideas and could be brought into the enterprise.

NORBERG: What sort of arguments were made to increase those resources?

SIMPSON: Of the sort that I gave you earlier. That is, "Here is an area that is obviously of import. Here is why from a functionality it is important as an area. Here are some good ideas on how to do that that are out there in the research community, and here's what a reasonable enterprise - money, investigators - would likely accomplish over a period of five years."

NORBERG: Who did you believe you had to convince about these issues?

SIMPSON: It's like every other program: the office director and the director of DARPA.

NORBERG: So what sort of interaction did you have with Kahn, and then with Amarel, then with Schwartz after that?

SIMPSON: The usual sort of interactions, "Here's some new ideas." The office director, remember, is always, a proactive player in these sort of things, both in encouraging and giving feedback, discouraging, and shaping this sort of things. It doesn't happen in a vacuum. There's discussions among program managers, discussions with the directors. Because directors, you know, motivate certain aspects or demotivate in some instances. It doesn't spring

forth from just program managers. It's sort of this stew that we live in.

NORBERG: I remember from my time at NSF that discussions about resources were always a period of blood-letting, in which if you betrayed something from your side about new ideas you might have, you might suddenly find that somebody else was lobbying against you the next time you walked into the program director's office - had been lobbying against you in the meantime - in order to make sure that that person's resources were increased and not yours. What I am wondering whether there was any of that sort of tension among the program officers, and whether or not that had to be then mitigated by the director of the office.

SIMPSON: Sure. The model at DARPA changes. This is in the sense that leadership in the office is an important intangible, because leadership in the office from a technical standpoint and a, I guess, morale standpoint is one of the things that an office director does in helping manage the tensions that you described. Program managers, by our nature, are competitive creatures. The competition can be on a professional level and still not be personal. If let get out of hand, it can become quite adversarial. Or it can be mutually supportive in the sense that people make natural divisions between their interests and join forces to support mutual activities. I think the forming of alliances, the forming of mutual programs is sort of a natural process that goes on in the office, as long as the office director is encouraging and motivates these kind of things in either one way or the other. I mean, office directors and directors of the agency sort of set major vectors. I am interested in doing this, vice this. So people try to naturally formulate what is a right way to do this. So there are levels of encouraging, interacting, and helping sort of a feedback loop on getting these things refined into something that each level of management is willing to support.

NORBERG: Over the last six years, or five years, though it would be six in the case I'm thinking of, what sort of differences did you observe in terms of management style among Kahn, Amarel, and Schwartz?

SIMPSON: Oh, it was dramatic. Let's see, which level do we want to talk about? From a technical level the technical backgrounds and interests among the three were dramatically different. On the level of management skills, there were very dramatically different management skills that each of these people brought to the game. In terms of leadership

styles... As a military officer I have studied leadership styles, so it is something I'm not immune to. I watched very, very different leadership styles come into play as well as sort of the natural office director responsibilities of charting directions for the office, interacting on a personal level with program managers, and those sorts of things. So other than specifics, it was very evident that we had been in three different corners of the universe.

NORBERG: I'm not looking for negative details, necessarily, unless they are significant. But I would like you to express some examples of the differences among these people and how they affected the office.

SIMPSON: I guess I am somewhat uncomfortable talking "out of school" as it were.

NORBERG: I can understand that.

SIMPSON: But it is pretty well known that I disagreed most severely with Jack Schwartz in his period as director of the office. In fact, he is the first person I have ever worked for in my Air Force career that I absolutely felt that I couldn't reconcile my personal opinions and judgments with the organization and my management. It was the first time I have ever had the occasion, and that was a very uncomfortable thing... In fact, I submitted my resignation from the Air Force as a result of it. The Air Force refused my request, but after six months working with Jack, I decided I could no longer be a good team member and reconcile my differences.

NORBERG: Without going into the personalities here, I'm not really interested in that, but how did that effect your program? Did it effect it in ways that he was suggesting that certain areas couldn't be funded? Or certain groups couldn't be funded?

SIMPSON: Yes.

NORBERG: It was that specific?

SIMPSON: Yes. There is a certain understanding about how... well, history maybe. There is a certain history that was passed down from Bob Kahn to Saul Amarel on how program managers and the director of the office operate, where the boundaries are, these sort of things.

NORBERG: What are the boundaries, or were the boundaries at the time with Kahn and Amarel?

SIMPSON: Well you can think of it in fairly classical management terms as the "what", not "how" - "What ought we to be doing?" My job was to decide "how" and "who." So one could say, for example, "We think that this broad area of activity is important, relevant to this important, so we're going to give it relatively more resources." Or, "We think that this area is one that we want to emphasize and therefore it is going to get a certain emphasis." Who should be selected, and how much of the resources should go to them was normally what a program manager's responsibility was. He was delegated the responsibility for making these sort of technical distinctions, managing, working out interfaces, for example, between these basic research areas and the technology programs, working out interfaces between technology and applications, doing the sort of linkaging and networking with other offices outside of ISTO, or with other agencies. So those sort of things were typically delegated to the program manager. There are a lot details, and it is hard to do. Normally office directors act as sort of final quality control checks, sanity checkers to make sure we aren't doing anything too stupid, as well as facilitators to make these sort of things happen. Typically, the office director is there to help the program manager when he needs some help, encouragement, or support. And it was rare during my years with Bob Kahn and Saul Amarel to ever have one of them disagree seriously with a judgement of mine. They would offer advice on how it might be improved. Never did I have a situation occur where I worked for some period of months to coordinate and develop a relationship between DARPA and another agency, let's say, and then have my office director cancel or otherwise negate the effort. As I think you would appreciate, there is a period of delicate negotiation that goes on, and then you want the approval to come about. Well, there were instances where these sort of things were dismissed as unimportant, discouraged, or otherwise negated. So those sort of things are not only embarrassing professionally, they are quite discouraging.

It is that nature of activities, as well as a general management style that I can characterize as something like the

following. My characterization of program managers and office directors during 1985 through 1987, maybe 1988, was we're a board of directors and the office director is chairman of the board. We collectively decide what the major issues are, the major directions, and propose as a board what we ought to be doing, and interact with each other as mutual teammates on this great enterprise. And the office director helps adjudicate some of the boundaries and resource allocation decisions. Then as a board of directors, we jointly agree, "Yup, this is what we're going to do. We all know our roles and our missions and we charge." That in say 1988-1989 changed dramatically into a role where, "Here is what we are going to do, and here is what I want you to do." It was very much, I call it, "the twenty mule team" style of management. "I am the mule team driver; you are the mules." Without going into more details, that hopefully characterizes the distinction in management style and interaction.

NORBERG: It does. I was thinking about how to classify the first one in terms of the chairman of the board and so on, and thinking of it as "the first among equals." Then "the twenty mule team" characterizes the second one rather well.

SIMPSON: Yes.

NORBERG: Okay, that's fine. I have heard this from others.

SIMPSON: I didn't think it was surprising. (Laughter)

NORBERG: It is not a surprise. Although I've heard it so far only from outsiders who observed this, or heard it from people like yourself.

SIMPSON: Unfortunately, it was very difficult to keep this internal, because it became progressively more difficult to explain, as some have described it, "irrational behavior" on the part of the agency.

NORBERG: I want to go to the second issue that was on my mind a few minutes ago when I said there were two of

them. It will help me not only to explore that issue, but get back to something else that I asked you before that I think is still not clear yet, at least in my mind. You said at a couple of points since we came into this office, something about the DARPA vision and people who are supporters of the DARPA vision. Then a few sentences later you talked about the AI vision. I am interested in learning whether or not ...

TAPE 2/SIDE 2

NORBERG: I think what I was saying is that you had mentioned the DARPA vision, and you had mentioned the AI vision, separated by several sentences. I am wondering if there is a similarity or an identity between those two, or if there is a difference. And if there is a difference, do you see any implications of that difference in terms of who is promoting the images or the visions, and who is the leader and who is follower here? And how has that changed over time, if you know that? So there are three parts to this question.

SIMPSON: Let's see if I can keep them on my stack while I unwind it. I do believe that part of the DARPA vision is the AI vision. I do think they are different, because DARPA's mission and DARPA's vision is much broader. AI is literally a part of the DARPA vision. I think it has been DARPA's historically from let's say the 1960s onward - part of DARPA's vision was AI. AI was a technical ingredient that DARPA viewed as something that both defined DARPA in terms of impact - "Here is an area of technology that we can point to that we've had an impact in" - and at the same time, it has been viewed as one of DARPA's ingredients that it brings to, if you want to think about the style of ingredients that it brings to the recipe for a defense program. So AI was one of those parts of how DARPA perhaps defined itself.

NORBERG: Who is the DARPA here you are talking about, Bob? Is this just limited to the computing people, or are you talking about all areas of DARPA -material science, nuclear testing, and so on?

SIMPSON: Think about this. In material sciences there are projects that apply AI to help in the design of new materials. If you think about aerospace planes, and doing control, design of aerospace structures, AI supporting

design tools, if you think about the controls on the airplanes as having some abilities to search very large combinatorial spaces... So AI is a pervasive piece of DARPA. It is in that sense that I mean, very broadly speaking, directors of DARPA over longer periods of time, back before George Heilmeier, especially during George Heilmeier's period in the 1970s, and then on into Bob Cooper have all understood that AI was one of the things, one of the tricks in DARPA's bag. Therefore AI was part of the DARPA vision. Because of that, they were always comfortable with whatever the office director decided was the appropriate resource level to support AI and the AI vision that I'll come back to. Different office directors, because of their interests and personalities, or whatever, AI as part of the DARPA vision was viewed as being big, small, in relative importance. I think in the case of Bob Cooper it was a very large part of his vision, of the DARPA vision. I think that is what led to the strategic computing program. I think in my time as part of the team during Ray Holiday (?), for example, or others, I think AI as part of the DARPA vision sort of fell down. It wasn't as important. It was known; it was recognized, it wasn't a large portion. Because of Craig Fields' background I think that AI has once again shifted into being a large, if not preeminent part of the DARPA vision. Hopefully that helps you understand my perspective on what DARPA thinks about what I do.

Now, having said that AI is part of the DARPA vision, there is a corresponding AI vision. At various times during DARPA's history, there have been either distributed managers of the AI vision, or there have been single people who have been keepers of the vision, if you will - keepers, promoters, constructors of the vision. It is easy for me to link back historically to my predecessors in DARPA who were the keepers of the AI vision. It is part of my appreciation for what they did, when they did it, and their sacrifices and contributions that make me sensitive, in some sense, to my responsibilities as the current keeper of the flame, if you will. Because it is very much the case that I saw during my five years at DARPA that there were a lot of supporters when times were good, but there were a lot of people that abandoned the ship when the times got tough. I found there wasn't anybody during the toughest of the times that was willing to stand up and take the heat and keep the vision. So there is very much the case that I felt like the lone stranger for a while, the only person that was willing to stand up and argue against what seemed to be the conventional wisdom, which at times I felt like was attempting to blow out the flame. It is very much the case that I felt almost an unwanted responsibility to keep the vision alive, as promoting, encouraging, articulating in some sense why AI is important, why it should be invested in seriously and promoted, and why it delivers new capabilities to

defense that cannot be acquired otherwise.

NORBERG: Now, was this in the last five years?

SIMPSON: Yes.

NORBERG: And this is in other areas of the Defense Department, I assume, not in DARPA?

SIMPSON: Primarily in DARPA.

NORBERG: Where were Sears and Amarel during this process?

SIMPSON: This is, remember, over this five years. When I came to DARPA strategic computing was a new program; AI was one of, if not the preeminent, technical area of investment. So during Amarel's period, during Alan Sears' period, during this, they were all shareholders and promoters of the vision. Amarel left. Alan Sears left. Jack Schwartz came in. The DARPA directors, Ray Holiday and before him... why am I drawing a blank here? You know who I'm talking about. Again, they did not see AI as a major portion of the DARPA vision. The net effect was there were not many people in the agency that were promoters or who could articulate the distinctions, importance... At the same time one could argue that there was balance occurring. I can obviously see that as an issue of resources distributed among many possible areas, there is a legitimate view that there would need to be some rebalancing of resources, and that is fair. By the same token when you look at DARPA's history, as you have done, I think it is easy to see that AI is one of the distinguishing characteristics that separates DARPA from other things. That was, in a sense, why I felt like we were not only disenfranchising part of our heritage, but we were also denying the fact that there was much more to be gained here than we seemed to be appreciating.

NORBERG: If you weren't getting the support within the agency, where did you go for such support, or did you not? In the outside community?

SIMPSON: Basically what I viewed as my responsibility as the keeper of the vision at the time was to articulate as best I could, in as many forums as I could, the urgency for people to step forward and be counted on - in other agencies, among the research community and among my defense colleagues - as well as continuing to fight internally in any opportunity that presented itself on new ideas, new opportunities. As I think you understand, because there is a dynamic here and because of the importance and obvious responsibility of DARPA directors and office directors for making major decisions on priorities and those sorts of things, there are certainly limitations to what a program manager alone can do. In my view, if there was an AI winter, it was literally during that period.

NORBERG: I'd like to make a suggestion to you and see how you respond to this in terms of this historical continuity that you were just referring to -trying to keep the vision. I see two things as characteristic of IPTO programs, I don't want to go broader than that at the moment. Two things occur very early in the establishment of the IPTO program, back in the days when AI was a portion of a lot of other activities...

SIMPSON: Well, it still is.

NORBERG: ...including time-sharing and networking and so on in the 1960s and early 1970s. The two characteristics are a certain convergence between the objectives of the research community and the objectives of the people in DARPA - not just in IPTO, in that sense, because you need to know something know about military requirements and so on in order to be able to understand how this stuff is going to be used. Even though you might take an argument like Licklider has often done in print that it was necessary to convince people that in order to do good command and control you had to have better computing, which seemed to make sense in the early 1960s. So I see that this convergence took place very early in the game, and that that convergence continues at least into the early 1980s. Now I'm trying to see whether you can stretch it for me, or whether you will even deny it. The second thing that I see in the IPTO programs is a certain coherence. The objectives as defined in 1982 for the strategic computing initiative are only a short way removed from the objectives of 1962.

SIMPSON: That is exactly right, you've got it.

NORBERG: Now, that's not to say that the same problems are being worked on. That's just to say that the same objectives are trying to be reached.

SIMPSON: The breadth and depth and sophis tication of the problems and the technologies applied to those problems is dramatically different, but in their core they are exactly the same sort of problems.

NORBERG: Therefore my conclusion, if those characteristics are a correct interpretation of the program over twenty years, is that the DARPA vision, such as it is as applied to IPTO, is something that was generated by a very tight coupling between the research community, on the one hand from whom most of the program managers come anyway as we know, and the program side, regardless of how many constraints there are on funding, evaluation, and so on. Now what I hear you telling me is that in the late 1980s that changed.

SIMPSON: Yes. I think that is true in the sense that, if you think of the research community, the DARPA community, if you want to call it that, and the ISTO, IPTO community, as being this tightly coordinated interacting collection of people who share the best interests of the country from a technical, scientific standpoint but als o from a defense and national responsibility standpoint, these people because they share that same concern could collectively be wise in making research, choosing research directions that were not only achievable, practical, visionary, and could deliver, if you will - since DARPA is a mission agency - deliver to Defense things that would be notable, of consequence. So you are right; I'm absolutely convinced of that. It appeared that, I guess from the standpoint that there was obvious stress going on, there was in fact some risk or some damage to that sort of relationship, that community, in which members felt like the best interests of everyone was being served. So I think that is true, in a more general sense, and in a particular sense in the AI part of the world, which is basically where I confined most of my interests. It was the case that the best interest, if you will, of the AI research community was no longer of consideration. Since part of what I would define to be the AI vision is this relationship between DARPA in caring for, being concerned about the health and well being of the nation's technology base - people, institutions, and research programs - was being

threatened because of resource limitations. I mean, there is a point at which, even if you believe that important redirections are necessary, the pace at which those changes take place needs to be factored in. So my concern in "keeping the AI vision alive" was that, in fact, with some of the directions we were headed in, we would endanger the survival of long-term health of what I still believe is one of our national treasures.

NORBERG: With the expanding activities in AI across so many different areas now, and not all of them confined to computer science and engineering to be sure, would you think that there is any reason to continue to try and keep such a flame alive? After all, it is a thriving enterprise now; it should be able to do well by itself, like a child being sent off from a family situation.

SIMPSON: I think that is oversimplifying the situation.

NORBERG: Fair enough. Tell me why.

SIMPSON: I believe that if you are saying that... Let's take an analogy to physics. Physics has been around and has been an enterprise for a much longer period of time, centuries. One could argue, "Why in the world are we spending money on Physics? It should be able to go on its own."

NORBERG: I wasn't suggesting that. Okay, I see why you mean simplistic.

SIMPSON: So all I am saying is, sure, there is some level of maturity, there is some technology, but I think the ambition which AI has - a very modest ambition of being able to simulate human intelligence - is likely not to occur for some centuries yet. Therefore, there is much work that needs to be done. My model is that maybe we have scratched the first couple, two, three percent of human complexity and reasoning and thinking. From that two or three percent that we've managed to scratch, a whole industry and in some sense new technical capabilities have sprung. If that is what we have achieved in sort of this modest way, if you think about it in the total scope how much money DOD, DARPA, has invested in AI is not all that large. Give me one B1 bomber or one B2 bomber and I'll fund

AI research for the next decade. So the amo unt of money is not all that extravagant, but the impact on the country has been considerable.

NORBERG: What is it that you think those impacts are, which was going to be my next question anyway.

SIMPSON: Oh, easy. Easy. If you look at entrepreneurial start-up companies, they are applying AI technology to all sectors of the economy. I can go down the litany of impact in just expert systems alone on where it has been applied in various sectors of the economy in improving productivity and efficiency in industry.

NORBERG: Can we use the Feigenbaum book as the example of this?

SIMPSON: Exactly. This sort of links back to when I mentioned that one of the things I did was speak out in as many forums as I could to articulate my concerns. One of the things I did was I tried to gather together what facts that I could put my hands on to define the impact of AI on defense, so that I could point to explicit places. It wasn't very hard for me to come up with a list of over eighty AI systems in defense that are either operational or operational prototypes of one sort or the other for which you can see clear, new functionality that didn't exist before for which there is clear dollar savings, in terms of defense resources, or efficiencies of various sorts that easily compensate for whatever was invested.

NORBERG: Two things then. Is that list available, that is, can it be made available?

SIMPSON: Sure. Part of what I've been briefing, sort of the Bob Simpson vision over the past two years, has been the articulation of that information.

NORBERG: Can I get some of that information?

SIMPSON: I was surprised that your literary search didn't turn it up.

NORBERG: Well, it didn't, I regret to say.

SIMPSON: Because the first time that I briefed this was at an SPIE conference.

NORBERG: The one in 1987?

SIMPSON: No. IEEE SPIE Conference on Application of AI in 1988, Orlando, Florida.

NORBERG: They didn't come up with that one. The student did not come up with that one.

SIMPSON: What I did there was for the first time I wrote a paper and also gave briefings for which I sort of continued to improve my list. Part of my articulation of the vision is, "Look, AI is a strategic technology for the United States, and let me tell you, it is also a critical technology for defense. That has been acknowledged as DOD's critical technologies plan. And let me explain to you why it is a critical technology. Let me show you what AI has done in just the past few years." This is where I go through my explanation and modeling of that. That is part of the evidence that I accumulated in this process of satisfying for myself, because I didn't want to continue to promote something that was at odds with my management and which left me in some disconsonance from an intellectual standpoint. So I gathered this data, and it is overwhelming. I presented it to DARPA management, which I think is one reason why AI, Craig Fields _____?_____ ____?____, is now back in, what I think, a position of relatively appropriate...

NORBERG: But usually, when one gives such presentations there are only a few examples offered, because there is simply just not time to go through eighty. So my question back to you again is still is that list available, and can I get it?

SIMPSON: Sure.

NORBERG: The second thing is can you just give me two or three examples now, so that I make sure that I understand what you have?

SIMPSON: Typically when I do this I have my list of eighty, and I show it, "Here is the list," and I ask people to ask me questions about any of them. Then I can talk about specific ones. But I also have my top ten. My top ten examples include, among other things, AALPS, automated air-load planning system. Let me tell you simply what AALPS does. AALPS is an aid to a load master, either an Army or an Air Force load master, whose job it is to figure out, given an arbitrary collection of equipment, how to load that equipment on board an arbitrary number and types of military airlift vehicles - C130s, C141s, C5s. They each have their different geometric constraints on their loading platforms; they have different constraints on how to distribute the weight on the floor; they have different constraints about what needs to be loaded in front of or behind what; how many passengers you can put on versus equipment. There are relationships among the equipment. You want to put a truck and its trailer relative to each other, so the truck can pull the trailer out. You want to not intermingle constraints on inflammable or explosive materials that might interact if the plane shook it up and moved it around. You want to worry about as you load equipment across a whole suite of planes, that critical pieces of equipment aren't all put on one plane, because what if that plane crashes, then the whole mission is scrubbed. You worry about issues like delivery: whether or not you are going to drop this stuff from airplanes and they're going to parachute down; whether you're going to do a low, high-speed pass, they drag it out as they come on; whether you're going to land and off-load it. All of these things make a computationally intractable - it is computationally intractable - in computation terms it is known as NP-Complete, to build a planner to do this. It is impossible to do. So algorithmically there is no way to do it. Ala... it is a prime candidate for AI. Okay? We know it can be done - why? We have human beings who do this, they are called load masters. An Army load master typically, with a group of people in the army, will take up to seven days in order to plan how to load out a division, like was shipped down into Panama. Take up to seven days to figure all this stuff out. Once they figure it all out, the Air Force typically doesn't send them the kind of planes they promised. So at the last minute they have to start shuffling stuff around. Is it any surprise that we have a lot of planes that are underloaded and a lot of equipment that gets left behind? Okay, into this foray comes AALPS. We modeled the

reasoning processes that load masters use to heuristically group things together, aggregate it, so that you reason over larger groups, you heuristically figure out how one collection of equipment will map into one type of plane, and then you duplicate it 'n' number of times - nice heuristic. We are able now - whereas it used to take a group of people, sometimes four or five people, up to seven days to do this, they had to be sort of well-trained experts. Load masters were wise old experts on how to do this. Now a brand new recruit, given training on AALPS, in fifteen minutes can do a load... with no particular training.

NORBERG: Seven days in a group, with the obvious cost savings.

SIMPSON: There are a lot of testing evaluation data that shows that we have improved the productivity, the effectiveness, and efficiency of doing this process dramatically. It is now a standard piece of software and it is available to the folks down in Fort Bragg, and it was used to load out shipments that went to Panama. It is in operational use daily.

NORBERG: Who developed AALPS?

SIMPSON: AALPS has an interesting history. It was conceived originally in 1979 by DARPA as part of the ADS program. So it was one of the projects that looked at how we could use networking and advanced comp uting to help in command and control. During that time at one of the brain-storming sessions people mentioned, "This is one of our hard problems, this loading stuff. Is there anything that could be done?" And people in the office said, "Yes, we think this new stuff called expert systems can help." So SRI was contracted to develop the prototype. It went through the usual rapid prototyping, revolving methodology until the final AALPS system in 1983 was basically declared the right functional system. From 1983 to about 1985, it went through some test and evaluation. Then, when I got to DARPA in 1985, AALPS was turned over to the Army as a project, program, and prototype for which DARPA was no longer interested. I got involved with AALPS because I was called up, being brand new there, by the Army and said, "We need you to send us a letter transferring responsibility for AALPS, or else we can't take it."

NORBERG: Why?

SIMPSON: Strange... some rule. So I quickly, since it was no skin off my nose, wrote them a letter transferring it. Since that time I have been watching it, because it is an instance of a successful piece of AI technology that has moved into test and evaluation development. The Army has been refining it ever since 1985. They ported it from Suns onto PCs, and they have done other sorts of activity as well as giving a very interesting test and evaluation. So I have data showing how it has improved. So one of the things I do is show, "Here's the data on why this has improved this process." And any general in the Army that knows about AI generally knows about AALPS. Because I and my colleagues have made it a point to point out that, "Here is an instance of this new technology called AI, that sort of slowly made its way through the standard development process, it has gone through all the blessings, it is an approved system by the Army." So that it is now AALPS manifests that are generated by this machine are acceptable as manifests by both the Army and the Air Force. So it has been through all the blessings, so it is now one of my eighty ... [TAPE ENDS]

END OF INTERVIEW