

Revised - 9/17/03

LASEWICZ: This is an oral history interview with Dr. Pat Selinger, IBM Fellow, on July 30, 2003, conducted by IBM Corporate Archivist, Paul Lasewicz. Welcome and thank you for coming.

SELINGER: You're very welcome.

LASEWICZ: I think we'll start off by talking about your background first. Can you describe your current position and title and what you do?

SELINGER: I'm an IBM Fellow, which is the senior most technical recognition. And then in addition to that I'm a vice president of data management architecture and technology.

So, that's a position where I really have overall responsibility for the architecture and direction that data management takes and its product direction and investment.

LASEWICZ: What schooling and education have you had?

SELINGER: Harvard University Bachelors, Masters and Ph.D. in the subject called applied mathematics, which is their word for computer science. You apply mathematics to biology or to economics and in my case to computers.

LASEWICZ: What courses interested you? Why did you choose mathematics?

SELINGER: I was a math major but got into programming classes and really loved that much more. Applied mathematics is really the way to get into more of the compiler or operating systems kinds of education classes.

Those were the ones that interested me. I had no classes on databases whatsoever. I didn't even really know what made a database at the time that I came out and joined IBM.

LASEWICZ: When did you first become aware of them and what attracted you to them?

SELINGER: I wrote a paper in programming languages that was given at a conference in San Diego in November.

I was in school in Boston, it was sleeting, it was snowing, it was miserable. Three buses ran by me just sort of standing with my suitcase on the way to the airport, throwing waves of slush all over me, drenched me.

And three subway rides later, I finally got on the airplane. Got off in San Diego, it was sunny, it was warm, I had a hotel room that was in the top floor of the hotel looking out over the beach, sunset, and it was phenomenal.

I said, I am going to work in a place in California. Why am I spending my time in Boston when I could be working in California? So I looked for places to apply in California.

And with the universities and various research institutions, IBM had a lab there. They said, well, if you're going to go to IBM in San Jose you've got to do databases. Okay. It'll be worth it. So, that's what I did...I applied....to other places too, but principally what I intended was to work in California.

LASEWICZ: You said that you were math major. When did you first become interested in mathematics and why?

SELINGER: That's an interesting story because I've always loved reading and enjoyed English and, in fact, was better at English and Social Studies and things like that than I was in math and physics.

But I always liked the mathematical things more. In high school and even the first year in college, I had in my mind that although I liked math I should do English because I'm really better at it. But, 40 papers in my freshman year convinced me that maybe there was something to this mathematics interest.

So, I ended up moving into pure math, which at Harvard is very, very pure, very theoretical. Wonderful intellectual exercises to think about things in very abstract ways and to have very structured disciplined thinking.

I ended up taking a logic class that was 900 people strong, taught by a very, very famous professor who had a very soft voice. And he was in the front, I ended up having to come from a far away building to get to the class. So I was always in the last row.

I couldn't hear him at all. And [I] decided, I have got to find a class to take that is closer to my last one. Well, there's a programming class that was available and I switched out of the logic class that I couldn't follow anyway, tried the programming class and the rest is history. I just loved it. So, again, not so much career planning but good taste or good luck or something.

LASEWICZ: Could you give us a summary of your work history? Was IBM your first job and if not what did you start with before then?

SELINGER: When I graduated from school with my Ph.D., yes, IBM was my first job and again, I fell in love with California and the people that I met on my interview day.

They were just fascinating to talk to and I thought these are good people to work with.

And I think I would advise anybody who's doing that kind of a selection among choices of jobs to pick who you want your colleagues to be, who you want to work with a little bit more than the subject matter because that's going to change. Particularly in a technology field.

LASEWICZ: So, in your years at IBM, you've been in databases.

SELINGER: I started out, I was very, very lucky, to start out in the early days of the System R project. And I brought with me a background in compilers and in operating systems. And it turns out database is a funny synergy of those areas plus a lot of complexity added because you're storing data.

It was a wonderful mix for me because I brought skills to the project that they didn't have so I had something to offer from the beginning. And I learned an awful lot. So I had great fun.

And so I started in database and kept going in that. I took a management job after three years. Another five years after that, I took a third line management job. So, I went

from first line management to third line management without ever being a second line manager.

And that third line position became a fourth line manager and then about six months after I did that, I became a functional manager and took over the computer science department at IBM Research in San Jose, CA. I managed 22 projects, 150 people or so. And got a taste of what going up the management chain was going to be like.

And at that point in time I really started missing technology, being able to get my hands on something concrete and specific.

Managing the computer science department was more like taking 22 classes for the 22 projects that I was supervising. So, for me, it made sense to find a way to stay close to the technology.

And what was happening at the time is that the Research division was doing these things called joint programs where the development division and the research team that looks into that same area joined forces and [they] work on things that are very relevant to the development team and yet very fascinating to the research team.

As the functional manager of the computer science department in Almaden, I put together a proposal to do one of these

joint programs in the database area. And then as I got more involved in putting that proposal together, I wanted that job. So, I formed it in the shape that I really wanted to do it. And for the next 12 years, I ran that program. I knew that I could build bridges between the research team and the development team. That came to kind of a finale for me when we took the research team, went to Janet Perna, who was at that time in charge of producing a database product for the UNIX environment.

I went to Janet and said, we are late to the marketplace, we are behind in technology here, we're not going to be relevant as a company or as a research team unless we have something out there that's very competitive, very fast.

So, what we want to do is we want to stop doing research and we want to go help the development team. We're just going to join forces, put together our technology and we will bring out a product that will be better than anything anybody else has.

So, that's what we did. For a period of about three to five years, the research team didn't do research. We learned how to do FVT [Functional Verification Test] and unit test and specifications and design reviews and error handling and checking in code into nightly builds and all of the things that people in the development team knew how to do.

And they taught us that and there was kind of a mutual respect that really forced a strong bond. And I had so much fun doing this joint development that I decided I needed to join the development team. I really couldn't go back to doing research. I needed to have my hands on the stuff that ships.

LASEWICZ: What were the clear benefits that you saw coming out of that, that synergy between research and development that might not have occurred otherwise?

SELINGER: We brought a second generation of technology to the relational database asset in IBM --and an addition to the world's relational database technology.

So, we had clearly leadership research results, new ideas. What we did is we took all of the work that we did with the System R, which was the original database project and that became the foundation for DB2 in the mainframe.

Then the research team stepped back and said, okay, now that we've done this once, let's do it the way we think it ought to be done. And that second generation was a project called Starburst.

We took Starburst, the refresh of technology, with its new ideas and put them into DB2, the next generation product for distributed UNIX, Windows, and LINUX.

That's the technology we put in there. The Starburst technology gave us a structure inside the code that made it possible for us to add more code and more function faster than any other database management system.

That helped us really close the gap for systems like Oracle and Microsoft [who] had been around for quite a number of years.

LASEWICZ: Can you give me the rough timeframe when all this was occurring?

SELINGER: Well, the System R project went from 1975 to 1979. That technology went over to the product divisions and shipped in 1981 and 1983 and the limited edition early ship vehicles for VM and for VS at the time. That became the first generation of DB2 as a product.

The functional manager job that I took was from 1983 to 1986. 1986 I formed the Database Technology Institute (DBTI). And that was what I did from '86 to '97. DBTI's technology was the work that formed this pipeline, this

technology pipeline between the research team and the development team.

And because people on the development team had all of the know how, as well as the actual technical, mechanical privileges to go put things into the builds that are shipping out the door, there was very fast delivery turnover from you when think of it to when it ships.

Our collaboration using Starburst technology to build the distributed platform DB created a foundation on which you could build things fast and then you had the ability to have fast access to an outlet for that technology through the development work that the research team was doing.

In 1997, we shipped the second version of this combined research and development technology. And that's when I moved over to development and took a job with title of director of DB2 integration.

The particular thrust of that job was to unify the DB2 technologies. Having a second generation of technology here and a first generation there, I really wanted to melt these together so that we could add in the strengths of this new technology into the core mission critical DB2 product that was offered on the mainframe. And so I spent about three years doing that.

Now I'm focusing more on incubating new directions and making sure that we're driving and investing in the things we have to go into to ensure growth for the future.

LASEWICZ: Would you classify this as a third generation?

SELINGER: I think so, yes. And, in fact, we're taking a technology called XML and adding that to our relational systems. And for us, the people who are involved in it, some of us were involved back in the System R original relational system and we all see this as another chance to set a new foundation, a new direction.

The basic idea here is that relational databases took care of a section of the world's data that is very structured, very organized. Inventory, billing, payroll checks, banking, that kind of thing. And, you know, everything is very simple. There are numbers, there are names, there are things like that, addresses. And then there's all the rest of the data. And that's something like 85 percent of the data and its documents, pictures, audio, income tax forms, it is pictures of my car accident, it's e-mail, it's genome data and you just go down that list and say, I've got to manage that too.

And, in fact, to have a clear picture, to really serve my customer when they call up and find an insurance support customer care type person, I need to see the customer's e-mail, I need to see somebody's notes about the last phone call they had with the customer, I need to see the transaction records, did they pay their bills.

What kind of insurance policies do they have with me? Is it cars, is it houses, life insurance and so forth. And I need to put that whole picture together really fast while I'm still on the phone with the customer so that I can talk about their specific case and make the customer feel like I really know what their situation is.

And so it's that other 85 percent of data that is not structured and in relational systems and being able to pull that together and to analyze it quickly; that is where we have to do some work now. That there's more invention needed in that area.

The old structured data that I've spent the first 27 years of my career on, well that was just the start.

LASEWICZ: You mentioned the synergies between research and development.

SELINGER: Right.

LASEWICZ: And melding the two perspectives. And I'm wondering, I guess you can throw professional organizations into that as well. Do you belong to professional organizations? And if so, what benefits have you derived from them? What value do they bring to you?

SELINGER: I'm a member of the Association for Computing Machinery, ACM. And a subgroup of that called a special interest group on the management of data. So, that's where the database researchers [are], that's the conferences they go to, that's the periodicals they subscribe to and so forth.

That's the list for me. I'm not a member of IEEE or SWE or any of these other organizations. I look at professional organizations as a way to get access to conferences and to associations to find out what's going on academically and in other industrial research labs and in other company's development teams.

All of those constituencies are represented in the database conferences that I go to. And that gives me a bigger picture, a view of what other people are looking at, as well as an outlet, a peer review outlet for the kinds of things that we're thinking about.

LASEWICZ: Okay, I'd like to go back to some of your formative experiences just a little bit. Were there any teachers or family that you found particularly supportive of your decision to get into technology or math that influenced that choice?

SELINGER: My dad was an electrical engineer. And my early memories are sitting next to him learning how to use his slide rule. So, my interest in math was always encouraged.

I think I had math exposure from early days just because of my dad having that particular career path. And he ended up choosing a management direction over a period of time but I absorbed from him a mathematical interest in what's a tangent, what's a cosine, what does this stuff mean and how do you calculate it.

My dad also didn't finish college before World War II so he returned back to school when I was maybe in third or fourth grade and continued on to finish his degree when I was in ninth grade.

So, it was an educationally good atmosphere at home. We got to study whatever he was studying. When he was taking music appreciation, we listened to music a lot, and so on.

So, I think the emphasis on learning being a valued thing and worth time and night school and weekends and so forth had a strong influence on me. His investment by example and his sharing his interests in electrical engineering with me I think was highly valuable.

And there were also science teachers who were just fascinating. So that contributed as well to my always liking and enjoying science and mathematics.

LASEWICZ: How about once you started your career, have you had mentors, have you been a mentor, what role do you see mentoring playing in an organization?

SELINGER: I think it's hugely important. I was lucky enough to have a superb person as a second line manager on the System R project. Frank King.

He was phenomenal in terms of his business understanding, his technology understanding, his ability to manage a software building team, [his ability] to make things happen from a research side, and to make them happen corporate wide.

Without his efforts, I am absolutely convinced that the relational model might have still stayed as a interesting

thing for IBM to have done a research project on and never would have turned into a product.

So, he was a phenomenal example and I learned a huge amount about how to manage and how to motivate people and that kind of thing from him. And at the same time, personally I spent just a lot of time talking with him and he was very encouraging.

If I were struggling with some area, he would sit down and talk me through what did I need [and tell me], "It's okay to feel nervous when you're doing presentations". All of the things that we need somebody to say, "It's okay, we all feel that way. Keep going. It'll be all right.

And practice in front of a mirror if you have to." You know, those kinds of things. It just made a huge difference to me, who had a year or so into the company to have a mentor at the second line level who then went on to be the functional manager, managing the entire computer science department.

Frank King and I kept up that relationship and I continued to talk with him about my ideas and vision for where we ought to be investing and where we're maybe falling off the road a little bit and needed to do some corrections.

And I hope I was helpful to him. I'm absolutely certain he was invaluable to me. He was a tremendous mentor for me.

And I continue to have, usually in my management chain, mentors throughout my career and very often, more than one. Janet Perna from a very early time while I was still in research was someone who was always there, always willing to listen.

And she would listen to what I had to say. She would help when there were things that were not quite right in research. You know, funding is getting smaller and those kinds of [issues]... and give some advice about what to do, how to handle it.

And Don Haderle, who is now my boss, but for... oh, gosh, I've known him since 1983, 1982 probably. Don is a great manager. He motivates people and yet has authority at the same time.

There are people who are very nice, but not very good managers because it's hard for them to make decisions. And there are people who make decisions but very autocratically -- but you don't like them very much.

It's much more rare to have people that have the ability to do both of those things. Both Janet and Don have those

qualities, as did my first mentor, Frank King. They all have been tremendous examples to me.

So, having benefited a lot from the mentoring system, I am giving back. I have about 30 or so people that I mentor. My assistant keeps a list and we meet every four months and [discuss] how it is going.

And some of them need [to meet] every two weeks. I have meetings with them over lunch, breakfast. For the ones who are not local, when I go to Toronto to see the development team there, I will have breakfast or dinners or lunches with them.

It's funny, people laugh when I book my calendar for a trip to Toronto, all the meals get booked up first and then the rest of the day. So people laugh that I eat my way through these business trips.

But it's easier to talk to somebody when you've got something to do with your hands and you're relaxing and just having a nice cup of coffee or whatever.

Some of the people that I mentor, I only met maybe two or three years after I started mentoring and it was all phone calls. Most of the people I mentor are in some kind of software. Some are in services.

LASEWICZ: Is 30 an unusual number?

SELINGER: It's pretty high. But people just keep coming and I can't... what I try and do is [say] okay, I'll have lunch with them and then I'll recommend somebody.

But they're really so good, I don't want to pass them off to someone else. I want to see how their story turns out. I want to give this person encouragement. They really have something there and I want to make sure it gets nurtured. So, that's how I end up with 30.

LASEWICZ: Okay, to change the topic a little bit. You got your Ph.D. Is it still relevant? How do you use your technical training in your position today?

SELINGER: Unfortunately, for my own personal satisfaction, I think, I don't get to program anymore. I wish I did and I find little projects at home to do to kind of make up for that.

So the programming part of my training it's relevant because I have the knowledge, not relevant in terms of being able to write programs and compile them everyday.

drive the future of this technology, not only for IBM but for the world.

I affect the lives of millions of people as we decide what goes into the next product releases and versions. We set some new directions and we decide how our investments are going to be shaped.

And that really makes a huge difference to customers in terms of their efficiency, to end users, to somebody at a bank who's trying to cash a check using an ATM machine. That's technology that we've done. So we touch a lot of people. And so having that ability to have impact is very satisfying.

The second thing that's more personally more satisfying to me everyday is the fact that I work with really smart people. And I just love the team that surrounds the whole data area.

And that's now growing from just data to WebSphere, to Tivoli, to Lotus. I'm getting to know more and more of those people as it becomes more and more important to our customers to really have a whole stack of IBM Middleware working together.

So, that's a difference. I think it's particularly true in software. Within data management and within software, I see a lot of women getting very engaged and they are finding the combination of the technology and the teamwork that you have to do to build software very appealing.

LASEWICZ: Has this change in the demographics in software had a noticeable effect in the work environment? You said you really didn't notice it before. Did you notice anything after?

SELINGER: I can't say that I really... this is one of the questions that I thought was hard because I don't have any stories about things used to be mean and now they are soft and friendly. It wasn't like that. I don't have anything useful to say there.

LASEWICZ: That's good.

SELINGER: Yes, I think so.

LASEWICZ: Can you talk about how you balanced your career and personal life? Were there any corporate programs or practices that helped you keep a semblance of balance?

SELINGER: Well, I took considerable advantage of flextime. As my kids were growing up, I was very involved

in their sports activities. So, weekends were spent going to soccer game after soccer game for my two boys. And as my older son became varsity for high school, now their games were 3:30 in the afternoon. And I would just go. I just went to every one of his games. He sat on the bench a lot, but I was always there. And I'm happy that I had the flexibility to do that.

I took off from 3:30 to 6:00 to go to my son's soccer game, then there's dinner and homework and so on. And about 9:00pm my day resumed and I finished off what I needed to get done that day.

So, being able to do that through the combination of flexible colleagues, the flextime that IBM would offer me and the ability to have home terminals and have e-mail and do a lot of work by e-mail.

I worked with the Toronto team very much during those years and they were three hours earlier than San Jose anyway, so they were going home at 3:30 when I was leaving for the soccer game. So, being in a multinational, multi-time zone company was also a big benefit to me.

LASEWICZ: Is there anything you would have done differently looking back?

SELINGER: I don't think so. I had a wonderful babysitter who was the only babysitter my kids have ever had. She took care of both of my kids. My oldest son is now 22 so from the time that he was eight weeks old, her home was his home as well as my home.

I would go over there for lunch. When her kids got married, we went to the weddings and now we've very close. They were my family while I was in California because the rest of my family is in the Midwest and further east.

So, that was a wonderful support network for me. If I hadn't had that, I would have had a much more difficult time I think. No regrets. No, I wouldn't do anything differently. I just was very lucky.

LASEWICZ: That's nice to be able to say that.

SELINGER: Yes.

LASEWICZ: What do you consider to be your most important contributions to science and technology?

SELINGER: Many people would probably say it was the access path selection paper on cost based optimization for relational databases. That really was the meat of what made the high level of the SQL language perform.

In SQL you could specify programmers who live in San Jose and who are in department 50 and the optimizer would map that down to a real way of walking through the records to get that data efficiently.

Optimization is the ability to do that, to decide if I'm looking for San Jose programmers who are in department 50, or if I'm looking for department 50 people and then check to see if they're San Jose programmers. Cost-based optimization means being able to choose between those two ways of accessing the data using a cost metric in terms of efficiency of access.

Cost-based optimization made relational work. It made relational databases perform. And it made it practical to code in this higher level, SQL, which gave people the application productivity being able to develop applications faster, which made relational databases valuable.

So, there are a lot of people who would say that that's my biggest contribution. The paper is still quoted. It's one of the most referenced papers in all of the database literature. It's taught in every college class that does database education. Database 101 always has that paper in it.

But I personally think that my biggest contribution was setting up that close connection between research and development. This is something that Xerox Park never could quite do. People kind of think back and say, well, researchers and developers they never "got it". So that they had beautiful research things but the company never could sell them or ship them.

This has been a traditional problem between research institutions and development institutions. And the Database Technology Institute that I formed solved that problem.

I set up a mechanism where there was a reward structure for everybody to cooperate. DBTI got people personally connected, which is what you have to do to get things moving fast.

You can't just make every decision go up to the chain of the managers and then come back down again. What you want to do is to have people with common interests meet and spend time together and just cooperate together all the time to get something done.

If you do that, then you have 1,000 connections going and not just one little pipeline up at the top of the management chain. And that's what we established. And it's

heartwarming to me to see how even after I've left that job and come to development, that close relationship continues.

And those relationships continue to thrive and influence what's important. We've got researchers deeply involved in customer critical situations because they're the best people to understand the technology. So, if the researchers are the best person to go help a customer, they'll go help. I spend time doing advanced technology investments with some leading-edge customers.

It doesn't even occur to me to worry about phoning up three researchers and two developers and saying, I need you to come out to customer X because they've got a wonderful warehouse problem for us to work on. And they come.

I don't feel like I have to go phone somebody and say, is it okay if I take one of your people for a day to go do this. It just all happens.

LASEWICZ: Do you see any influence between the institute that you were a part of and I guess starting in the late eighties and the directions that research is taking since then?

I mean, you do have that researcher-for-hire program going on now and after Mr. Gerstner came on board, there was a

closer tie between research and the products in terms of trying to shorten the time when research, innovations get into products. Do you see a connection between the institute and that?

SELINGER: I spent a huge amount of time in the mid-nineties telling people who were forming those kinds of organizations or informal organizations what I did in the Database Technology Institute because they were all trying to model their organization after what I had done in DBTI.

I had some rules of thumb and so forth. But mostly it comes down to the fact that there's no substitute for rubbing elbows. And that's easy to say and hard to get to happen sometimes. It really comes down to individuals.

I'd like to think that I set a direction that people wanted to model and wanted to copy. And I spend an awful lot of time helping develop some of these other organizations, to at least share my thoughts about what I did.

In those other organizations, you may have a different circumstance. You clearly have different personalities involved. But this is what worked for me. I tell them to think about using these kinds of techniques, these kinds of mechanisms as you set your directions, as you try to get the two teams involved, et cetera.

LASEWICZ: You've obviously achieved a great deal of success during your career. What are some of things that you think have contributed to your ability to be successful?

SELINGER: I think my ability to learn fast has been very helpful. And as we talked about before, the ability to get a picture of a working system in my head with someone just sitting down and drawing me a picture and I can start to animate parts of it and be able to say, oh, yes, I can understand how these gears fit together and that pumps the widget and I can see how a system works.

That is not anything that you can be taught. But I think you can hone the skill. You can refine it. You can practice it. So, I think that I had an innate capability there that helped me a lot.

But I've never thought that I was the smartest person in a room or anything like that. When I didn't feel that I am learning as fast as I could or getting as much done as I wanted, it was just kind of turn up the heat, turn up the volume and work harder, spend more time, focus more. So, I might not have been the smartest person, but I was probably one of the hardest working, most dedicated.

I think that the passion, the enjoyment that I get from my work has contributed to my success. You're much more eager to get up in the morning and get into things when you like what you're doing.

It is also important to like the people you're working with. I think that has made a big difference and that has helped me be very successful. And I was lucky to start my career, at that point when we were just starting to build relational systems. That was good timing on my part. Accidentally.

But, I think after such a lucky start you kind of make your own luck and you build on successes and you leverage what worked before and you try and reach out and take the best ideas and engage other people in the problem you're trying to solve. And I think that makes people very successful.

LASEWICZ: Well, that's obviously worked for you. In terms of somebody just starting off in the business in the field of technology, is there any advice that you'd be able to offer them from your position?

SELINGER: I think there are two things: One is that people who are really talented will get pushed into a perhaps earlier than necessary decision about going into management. If you're talented and you're capable, you lead

naturally. And most people spot that and say, well, then I want to make you the manager of this group.

I would say that you don't really have to do that as early you get encouraged to do it. I moved very quickly up the management chain and it was good for me for a while, kept me very busy.

But, I made a lot more mistakes than I would have had to make if I'd waited a little bit longer. I didn't really have to make that decision that quickly and I might have been better off had I taken things a little bit more slowly.

So, for a technology person who's thinking about management, stick with the technology a little longer. You can do both for a while, maybe as a first line manager for example. And continue to invest in the technology side of your career.

Build a strong base, a strong foundation, a strong set of knowledge and then that can be a base for going on up through management. Or for going in the direction of more of a distinguished engineer, IBM Fellow kind of thing. And it keeps your options open.

Lots of opportunities will come to people who are talented. This isn't the only chance that someone's ever going to have

to decide to be a manager or not. So, take it easy. Do what you love. Follow your heart.

Don't take a job because it's a promotion or a rising up in management level and status or something like that. Do what you love to do because that's what you're going to be happiest at doing.

LASEWICZ: We covered a lot of ground. But is there anything else you'd like to add?

SELINGER: I don't think so. I think that's really covered what I think about when I talk to people about my career or what I would recommend for them. I think those things are important.

I'm glad we spent the time to talk about mentoring. As my career goes on longer, I spend more and more of my time doing that because I think you've got to keep pulling the generation behind you. We can't just be the experts here and then we all fall off the end of the earth. We've got to keep our successes going. I've got my retirement to collect!

LASEWICZ: Well, thank you very much

SELINGER: It was my pleasure.

[END OF INTERVIEW]

