

**Stretch Reunion
Interview with Fran Allen
September 28, 2002**

**Interviewed by Dag Spicer
Computer History Museum**

IBM 7030 (“STRETCH”) REUNION TRANSCRIPTS
General Notes

On September 28-30, 2002, a unique group of computer professionals met in Poughkeepsie, New York, to celebrate the IBM 7030 (aka “Stretch”) computer. This computer, first shipped in 1961 and over five years in the making, is one of the most remarkable computer products ever designed. With dozens of new architectural concepts that revolutionized the industry as well as the nascent field of computer science, Stretch embodied the very best of IBM—the best people, the best technology, the most demanding customers.

This transcript is a verbatim transcript of interviews conducted during the course of the Reunion. The Computer History Museum, home to the world’s largest single collection of computer artifacts, is proud to offer this series of transcripts as part of its ongoing mission to preserve and present the artifacts and stories of the information age.

Every effort has been made to check the accuracy of this transcript. All interviewees were asked to verify the relevant transcript. When they replied with changes or comments, this is indicated in the footer of each document’s pages by the phrase “Checked by Interviewee.” Note that most of the subjects did not respond to CHM’s request to proofread their comments.

If you have any questions or feedback relating to this transcript, please contact Dag Spicer, spicer@computerhistory.org.

Computer History Museum
<http://www.computerhistory.org>

© 2002 -2004 Computer History Museum

DAG SPICER: It's September 28th 2002, we're in East Fishkill, New York at the Casperkill Country Club, the old IBM country club, with Fran Allen Allen. Thanks so much for being here and talking with us about Stretch. Can you tell us how you got hired on to Stretch and Harvest and what attracted you to the project?

Fran Allen: Oh, I was already working at [IBM] Research. Research was actually up here in Poughkeepsie at the time and I had transferred in advanced guard down to a place called Lamastate and they gathered a group of us up shortly after I'd gone down there and sent us back to Poughkeepsie to work on Stretch / Harvest, and I was interviewed with Sullivan Campbell who was one of the people that was very involved in and a wonderful man and he . . . it didn't take me very long to be pretty excited about working on Harvest and on the compiler and on the project.

DAG SPICER: Can you tell us a bit about Alpha and how that came about?

Fran Allen: Alpha is a very high-level language that was designed collaboratively between IBM and the National Security Agency in for codebreaking. And so the cryptologist at the

Computer History Museum
<http://www.computerhistory.org>

© 2002 -2004 Computer History Museum

National Security Agency then would be able to write their codebreaking programs and run it on Harvest, which was a very, very complicated machine and so the language was able to bridge between the problem and this complicated machine which was really designed to solve the problem, and it could do it in a way which was very straightforward for the applications programmers at NSF--NSA, I'm sorry, NSA, I'm sorry <laughs>.

DAG SPICER: Some people may not know the distinction between Stretch and Harvest. What made Harvest Harvest versus Stretch?

Fran Allen: Well Harvest was an add-on to Stretch, so Stretch really was the host for the Harvest machine. The Harvest machine was physically bigger and had more transistors and took up more floor space than Stretch did, so it was a was a monster machine in addition to having the Harvest part it was also a "Tractor" tape system, and which could stream data feeding the streaming capabilities of the Harvest machine. And perhaps the most unique thing about the system was that the I/O through the tractor tape system and the Harvest streaming unit was perfectly balanced, so instructions on the streaming unit, [it] might be just one instruction, could go for hours and-as the data just flowed in from the tapes and back-and the answers flowed back out and cycled through.

DAG SPICER: I noticed you brought a tractor tape cartridge which I'm glad you did because I've never seen one. I took about twenty pictures of it because I think they're fairly rare. I think the Cryptology Museum at the NSA maybe has one, but otherwise I've never seen one. So the group you were in, your immediate supervisor was Sullivan Campbell was it?

Fran Allen: No the immediate--he was the one that was . . . I don't exactly know what role he had but he was one of the people that was overseeing the whole Harvest system. And one of the people that we honored tonight because of his seminal role and in the Stretch / Harvest system, putting together the team.

DAG SPICER: When you gave a lecture for the Computer Museum a couple of years ago, I think, on Alpha and so on, the streaming data aspect of Harvest reminded me of the bombes that they used at Bletchley Park. They used paper tape that was going, I think at sixty miles an hour or some incredible speed. If you can't answer I understand, but was this an electronic implementation of that same kind of algorithm?

Fran Allen: I'm not familiar enough with what they were doing at Bletchley Park, but of course they were working on some of the same problems. And it's very possible that some of the

antecedents of the ideas that we were using came from there.
But none of us were aware of that.

DAG SPICER: Of the actual application space?

Fran Allen: That's right, and in fact, even though I had a high level clearance in working with the NSA, I had a Q clearance, I knew nothing about the methods for codebreaking <laughs>.

DAG SPICER: One thing I ask everyone is how on earth did Stretch come about, this incredible qualitative leap in computer architecture, not just a slightly faster machine but a radically different way of solving performance problems. Some people say it's just fantastic management, other people emphasize the very young people that they hired, and others cite IBM's corporate history just seemed to be ready for that. Do you have any philosophical thoughts on that?

Fran Allen: It's basically that we didn't know better than to try and do it, and we did it <laughs>, because if one thinks when the Stretch system was actually starting to be built, which was essentially in 1956 and we delivered it in '61. In that very short period of time everything was built, and including hardware, software, without many tools, and much of it just invented as we went. And Harvest was going on at the same time,

and this was an add-on to Stretch so there was a lot of relationships between the teams and, you know, the two teams were the same, that was started in terms of actually building it in around '57 and that was delivered to NSA in '62. So we had these huge gigantic projects, you know, related to each other and delivered in a very short period of time. We would never these days tackle anything that big and that new I think, on that kind of schedule. But because we were young, because we could invent what we needed and not have to just—we just did it <laughs>.

DAG SPICER: It seems to have really bootstrapped IBM, it let it leapfrog in so many different ways.

Fran Allen: Right, and as one of the people at NSA have written since then about the project, about the Harvest project, but it applies to all of it, is he said this had a great impact on engineering but more than anything else it probably impacted the engineers and their futures.

DAG SPICER: So it was much like the "Fairchildren," the people who used to work at Fairchild who then went out and started all of Silicon Valley. So the same kind of thing happened?

Fran Allen: Right.

DAG SPICER: I wanted to ask about two people, John Cocke who couldn't make it, and Steve Dunwell. I don't know if you interacted with them, I think John Cocke you spent a lot of time with?

Fran Allen: Yeah I've interacted a lot with John Cocke.

DAG SPICER: We'd like to hear your thoughts about him.

Fran Allen: Well, John joined IBM in '56 and so this was right at the beginning of the Stretch project and he was one of the most inventive people on the project. He just kind of circled around and, you know, he was young like the rest of us of course, but we all were intrigued by his ideas and we shared ours and it was just kind of he was part of this general milieu of sharing ideas and working hard and having fun. He was a very special person and I believe that he has probably had the most influence of any person on the computing field as in the shape it's in today.

DAG SPICER: He seems to have played the role of a honeybee or something, cross-pollinating ideas.

Fran Allen: Yeah that's quite a good analogy, right, he would circulate with ideas and he would visit lots and lots of places,

Computer History Museum
<http://www.computerhistory.org>

© 2002 -2004 Computer History Museum

you know, people and offices and wherever they were, but when he was working on--with each of the people was one particular aspect of--of a set of ideas. And so he would walk into somebody's office and pick up a conversation about an idea that he'd started a month ago, and pick up exactly in the sentence where he left off and so he was a true example of being able to parallel process but had total input and output were also going all of the time, so he understood what was going on with the people he was talking with and he was a wonderful free spirit but not a free spirit in the sense of just dropping ideas, they--he always came back to them, he worked on the same problems for most of his career, they took different shapes, he saw more insights, and he added new solutions, but problems that he started working on when he first started he was still working on much later <laughs>.

DAG SPICER: One thing I'm interested in--and I think you're uniquely qualified to comment on this--is the role of optimizing compilers and RISC hardware and for him did the idea of a software optimizing compiler start first and simplified hardware drop out of that or vice versa?

Fran Allen: Well, what had happened is that after our experience on Stretch and Harvest was we came out of that knowing that we needed to focus on performance, we'd focused on a lot of

functionality in that project and we had focused also on some very ambitious compiler structures. But then the next project which we had was ACS, Advanced Computing System, which never came out but it was and there we focused upon getting the performance and getting a set of algorithms and techniques in place so that we would be able to get very, very high performance, and on that project we started the compiler first, and we had a compiler running before the machine was designed, we drew the compiler output drove the machine design through a simulator, and that idea then was carried forward into the next big project that John was working on and that was the RISC project where again, they started the compiler first and then just made the tradeoffs between what could be done and compiled by a compiler and what could be done in the hardware so that he really honed in that project the idea of having the compiler in the-and the architecture designed at exactly the same time.

DAG SPICER: This was the 801?

Fran Allen: That was the 801, right. So but it was a single stream of work over quite a few years.

DAG SPICER: Stretch's influence on your career: what happened after Harvest?

Fran Allen: After the Stretch Harvest then I went to join the ACS project and because I was very interested in the optimization and knew where we had failed, had some problems. We hadn't failed on the Harvest side but we had failed on the Stretch side with some of the performance and from a compiling point of view. And so that's what I did next and when I finished the Harvest work and I wrote kind of a final report on Harvest and so forth and then I had a choice of going on to 360 because this was in 1962, '63 I guess, and there was a lot of opportunities for hard work on the 360 at that point because it hadn't come out, it didn't come out until '65, but I-there was an opportunity-the thing that attracted me most was John Cocke called me up and said let's come up-back to research and look-we'll look at another optimizing compiler. And that led eventually to the ACS project, so.

DAG SPICER: And then after ACS?

Fran Allen: After ACS I worked on applying the techniques to new ways of building compilers and did some work on experimental compiling systems and then after that and working also with the PL Point 8 people but I wanted to try and use the techniques in a little different way than redoing it again on a machine design. And then after that I went into parallelism, right.

DAG SPICER: Is there anything you want to say about the parallelism, parallelizing compilers?

Fran Allen: Yeah that was another great challenge and that turned out to be great fun because we could build the parallelizing work on top of the existing work that we had and existing codes that we had, and in fact we ended up building parallel compilers for [System] 390's, for-we had a experimental machine in research on our-called Research Parallel Processing 3, "RP3," and we had several others that we worked on and we built systems for. And there had been some great work before then done at the University of Illinois with Dave Cook's work, and so we hired some of his students and and that was again a fun project, we built a great team and got a lot of good results there.

DAG SPICER: I know one thing that we think about in terms of getting our [the museum's] content on the web is designing a kind of museum compiler where people can interrogate our various computer historical assets, you know, on the fly and literally generate a timeline of Stretch for example and just pull all the pieces together.

Fran Allen: Oh that's interesting, yeah that'd interest me.

DAG SPICER: Is there anything you want to leave us with about Stretch and when I look at the people here today it really is an incredible group of just brilliant, hardworking. Everyone seems to have been really happy and exited and passionate about the project. You don't see that a lot in people these days.

Fran Allen: Well I think the thing that characterized that and characterized most of these projects that I have enjoyed so much but is the challenge and these were projects with great challenges and it wasn't just tweaking something, it was really setting very high goals and some of the goals would be schedule goals--and some of them would be performance goals and some of them would be capability goals--then working as hard as we could to achieve them and it's partially invention in order to satisfy the goals, but some of the--some of the other--the rest of the time it's sort of putting aside invention and doing the engineering which--and they're equally interesting, both the invention and the engineering to make something real.

DAG SPICER: Lastly any special people that really influenced you or mentored you or were an inspiration to you on Stretch or after or before Stretch?

Fran Allen: Oh gosh all the people I've worked with have been great, but of course John Cocke, he was the person that I always

found exciting to be with, and I-and his ideas would just sparkle, you know, and this I found that just just thrilling. And then Sullivan Campbell was a wonderful man that was very, very encouraging and created an environment where it was just fun to work, I think that's an important part of-of work.

DAG SPICER: He's still a mystery to me even after a lot of people have said a lot of peripheral things about him.

Fran Allen: He's just a warm great storyteller. Again he's somebody that not only could talk and tell good stories but he also would read what was going on, he could walk into a room where there was a huge amount of tension going on, a big discussion and start talking and everything would kind of quiet down and the problem would go away or get solved before he had finished talking <laughs>, and I've seen him do it, I saw him do it at NSA one time. He came down and he got-was coming-came down, flew down late for the meeting, got in the meeting and we were yeah IBM and NSA was-was-we were going at it on some technical problem and it-so he had no idea when he worked in the room that this was what he was going into, and he never asked "What's the problem?" he just started to talk and the problem got solved <laughs>. Yeah, but that's the kind of person he was, and both of those guys, John Cocke and Sullivan Campbell

could read rooms and read people and knew how to keep going straight ahead, but involve everyone that was with them. Right.

DAG SPICER: Anything else you want to say?

Fran Allen: No I don't think so, this has been a great a great event and I think I'm so glad that you people were here, because I think it's important to capture these events.

DAG SPICER: We have a soft spot for Stretch, actually, I do anyway.

Fran Allen: Yes and you pulled together some marvelous material on Stretch, you know, netted it out in terms-in the technical terms and very succinctly in one of the one of the pamphlets from the Computer History Museum. Yeah that was nice. Congratulations yeah.

DAG SPICER: Great.

-----END-----