

20,000 Gates & 20 MIPS

A Look at 20 Years of
Innovation with Bruce Beebe
and Dave Brewer

Innovation is risky business. It isn't always knowing what you're doing, it isn't always doing the right thing, but it is learning from your mistakes and improving on them. In an interview conducted just days before Amdahl revealed its sixth generation of products, these first generation employees revealed how Amdahl has forged innovations in both the technology business and the "people" business.

How did you first make your way to Amdahl?

Bruce: After going Chapter 11 at a startup called MASCOR, Dave and I were virtually on the street when Gene Amdahl came along with a new opportunity. When you're looking at it from that perspective...it's real easy! We had worked with Gene at IBM®, so we had some acquaintance with one another. It really wasn't a difficult decision to make.

Gene felt that the large end of the market wasn't being aggressively pursued because IBM's technology was always designed to address the midrange, where most of the dollars were. To create a high-end machine, they just tried to stretch that midrange technology. Gene thought if you designed the technology specifically for a high-end machine, you could be much more successful. He first tried to convince IBM of this. When he was unsuccessful, he decided to do it himself—and Amdahl Corporation was born.

Gene's vision was a very small machine, because small meant fast, and very high performance—certainly the success of the 470® was wrapped around the technology behind its 100 gate LSI. The ex-

pectation the first day we heard about the machine was 20,000 gates and 20 MIPS. But as we sorted through it, 12 MIPS became the next goal, and when we actually implemented it the result was about three MIPS. In those days, no one measured these things very accurately. Anyway, we significantly missed our vision, yet we were well ahead of the competition.

"In the early days, everyone was working on the same thing practically within shouting distance and the uniformity of goals was good, the team spirit was good. The challenge for today is to maintain that kind of feeling and that kind of energy with the enormous groups of people that we have."

What about the tools used then? We now make the majority of our development tools....

Dave: And we made our own then. But they ran on IBM 1130s—we did all of the design files on punch cards that we would feed into the machine, and the machine would process them and give out wire lists. And we had a router to produce artwork for the chips and boards.

Bruce: We also had the hardware Logic Simulator. We had this thing about as long as a freight train built out of technology that was one generation older than the LSI. We did that because you could never build out of the LSI and get it right the first time. So we built it out of this

older generation technology that was very changeable.

Each week we'd isolate design problems on the machine and on the weekend we'd implement Engineering Changes (ECs) for the problems we'd found. We'd tear the machine down and change all these wires on the back panels and then we'd put the machine back together at the end of the weekend and try to bring it back up again. We got to the point where we spent Monday-Thursday just trying to get the machine to run again, and then we'd have only one day to work on it before we'd have to take it down for ECs again. The hardware logic simulator was a great concept, but we didn't spend enough time engineering it.

How much storage did the original 470 have?

Bruce: We may have been all the way up to eight megabytes of memory in full configuration—less than PC-class stuff these days.

Thinking about the areas of responsibility such a small group of people had in the early days, how many people would cover comparable design areas now?

Dave: Glenn Grant, Mike Clements, Richard Bishop, and Dick Tobias were the lead designers in each of the major areas such as the I-unit and the S-unit, and they probably had 6-8 folks working for them in each unit. So, when we first started out about 25 people were working on the design.

My guess is, a little later on there were probably 50 people where there are 500 people now. Design Automation then was probably 4-5 people, and now it's 150. The technology group (chip design) was

4-5 people then, and the comparable group now is probably 40-50. And of course, we do many things now that we didn't do then. We didn't have anything called Macrocode then, so that's all new. Diagnostics, that's a big organization now.

Bruce: Things are also complicated much more now because our design groups are working on several different product generations at the same time.

Was the idea of upgradeability a goal from the beginning?

Dave: No. That just kind of happened.

Bruce: I guess the first "upgrade" was the V/8. The V/7 was a new machine, with a redesigned S-unit (storage system). With the V/8 we just went in and honed the cycle time of the V/7. The idea of upgrades wasn't originally a part of our strategy—it fell out of what we did to squeeze more performance out of the same basic design.

Dave: You could almost say we were sloppy with the V/6—we didn't spend a lot of time trying to get the most performance we could out of the technology—we were so far ahead of the competition. It really wasn't necessary to squeeze it harder. And then we went back and squeezed it and went from a 32.5 nanosecond cycle time to, I think, a 24 or 26.

What about factors like manufacturability and user-friendliness? How big a part did they play in the initial design?

Dave: Amdahl was a classic engineering-driven company, which meant that we ignored all of that. We got ourselves in

deep weeds because of this. The bonding of the fine wires to the back of the Multiple Chip Carriers, something we do easily now, was a technique no one had ever done before. We bought some bonders and hired a bunch of people and got all set up to manufacture before we were really ready. The company had a very big setback because of this—in fact, we very nearly went out of business in 1974.

When was Fujitsu first involved?

Bruce: In November, 1971 some initial meetings were held, but those were focused on Fujitsu making a financial investment rather than development. They were trying to understand what we had.

Dave: Bruce and I first went over there in February of 1972, and very soon after they sent a team of about a dozen people over here. They were really bright guys, very good at what they did. I was always impressed with how easily they adapted to us. I can remember playing touch football with them, and of course they had never played football before, but they picked it up very quickly.

Bruce: And the entire Fujitsu contingent learned to drive in the Amdahl parking lot. We had fun working together.

Were they doing the same kind of design work as you were?

Dave: They were sharing offices and working elbow to elbow. They worked long and hard, and then they went off and held their own meetings because they were not only trying to help us get our

design done, they were also trying to figure out how they were going to use the machine once they got it back to Japan.



Bruce: As far as managing them went, we first tried to take a typical Amdahl organization chart and plug the Japanese into appropriate organizations based on their skills. This didn't work, because they are very hierarchical in their management structures. We ended up having one of the lead Japanese being the group leader for 2-3 junior Japanese engineers. The senior level person would work with the Amdahl engineers and any communications to the junior guys would be worked through him.

We allowed the Fujitsu group to organize their way and then we built a bridge between the two organizations. It wasn't a serious problem, we just had to go through that learning experience.

When the 580 came around, the next generation, what was done differently?


Bruce: We started talking about that technology in 1977. The chips were four times bigger—400 gates per chip. It was also a bigger board, 11x11 chips, 121 versus 42. More layers. Yet it was still an extension of the V series.

Dave: To me, the remarkable thing about the 580 was that we didn't do much differently, and that was our mistake. With the 470, we had an inherently more reliable technology than the competition, so with the 580 we said "okay, we'll make a machine about as reliable as the V/6





From a shared office at the IBM Poughkeepsie, NY Development Facility to their current respective positions as Senior Vice President, Product Operations and Vice President, Processor Products, Bruce Beebe (left) and Dave Brewer have ridden on the leading edge of innovation—and they've seen a lot of changes in twenty years with Amdahl.



because that worked well the last time." Then the other guy leapfrogged us.

For all the vendors in those days, the dominating factor in outages was design flaws—you sent them over to customers without getting all the bugs out and ended up debugging them for the first year on the customer's floor. We did it and so did IBM—except that they began to develop better means of testing internally.

So what did we do to change that?

Bruce: Implemented a lot more disciplined testing. Now we're doing much more simulation and design checking.

Is it easier to develop machines now than it was then? Today's machines are much more complex, of course, but we have so many more tools....

Dave: There's one big difference that we have now that we didn't have then. Then we could fail. Now we can't. If we fail now we have 8,500 people and I don't know how many investors riding on this. In those days, we had a handful of venture capitalists who understood the high risks of the business. If they hit one in ten they were happy. So we took all sorts of risks, in retrospect, that we'd never do now.

Does that conservative philosophy carry over into the engineering arena? Are we risk averse in the design work we do now?

Dave: Not at all. In fact, it's the opposite. We have to press the technology very hard, a lot harder than we did then.

For example, a lot of the circuitry that we put in now is for recovery. I don't think we had anything in the V/6 for recovery—maybe a 30 percent recovery at a maximum.

Our conservative attitude today has much more to do with schedule. I don't know what our first schedule was back in 1970—probably shipping in 1973. And we shipped in 1975. We just couldn't afford to do that now.

Bruce: You could examine this graphically if you compared the principles of operation or the design rules that we had then with what exists today. Then, the design rules were probably a half-dozen sheets of paper where there would be a one-inch notebook today. And you'd find this order of magnitude and change in complexity in every area of processor design.

Some people say that mainframes are heading toward extinction, but we're certainly not building our business based on that assumption. Where do you see them going in the next ten years or so?

Dave: Within ten years, I think we'll see some extremely powerful single chip machines. The advantage is that you get a tremendous amount of performance for a very low cost, so the whole balance of the market changes.

Bruce: The requirement to manage huge amounts of data is not going to go away. Therefore, the need for something

very large to manage large amounts of data is still going to be there.

Is it still fun?

Dave: For me, the "fun" issue now has a lot more to do with people than it does with technology. We've gotten to the point where we need to get large groups of people working together well. In the early days, everyone was working on the same thing practically within shouting distance and the communication was good, the uniformity of goals was good, the team spirit was good. The challenge for today is to maintain that kind of feeling and that kind of energy with the enormous groups of people that we have.

Bruce: Human Resource organizations say they are in the people business. Baloney. We're all in the people business, and if we think we're in the technology business and not the people business we're not going to grow the company right. Getting and keeping the right people is how we made it then, and it's how we're going to continue building for the future.