

Don White 3/88

Interviewed by Jamie Pearson at the Mill

DW: In the early days, the modules were basically of two styles. One was the lab module, which we have, and the system module were the ones with the half inch aluminum frame around the board. Of course, all the back panel wiring was either patch cords or soldered wires. And it wasn't until we got to the PDP7 and PDP8 when we reversed the wirings. The system modules were used to build the PDP1. The PDP 2 and 3 were never built except on paper because after we got the first couple PDP1's going there was the question of spreading ourselves too thin. So we stayed with the PDP1 and the ADX system for ITT. So the 2 or 3 never got built. The PDP4, everybody knows about. Number 5 was the predecessor of the PDP8, LINC's and DECmates. The PDP1 was an 18-bit machine.

Back to the module stuff, the 1000 series which was the 5 megahertz stuff was pretty well established. I did do some work on the things for the core testers and memory testers. And Russ Doane concentrated on the 5000's /6000's 10 megahertz stuff. I

concentrated on the 500 kilowatt stuff, the 2000/4000 series.

Russ started the company a week after I did and I'm sure he'll have a lot of history to tell you about PDP-6 modules. By the way, there's a rumor that goes you could probably could drain the mill pond and dig up enough modules to get a PDP-6. Module check-out overlooked the pond, and once in a while they would find modules that would never see the light of day.

That was before wirewrapping. The wire wrap story was the back plane problem; you see, with the PDP1 and the PDP4 every little wire that had to be put in place had to be hand soldered. About the time the first PDP7 was introduced, Gardner Denver had come out with wire wrap technology. This was kind of an option I think from ATT because they had to do some wire wrap stuff; so the PDP7 got built with the first wire wrap technology.

JP: Could you explain all the variations on the PDP8 after the classic 8?

DW: O.K., well the letters were sort of added. They had a use at one point in time. It was not alphabetical and it sort of got out

of hand after a while. But the PDP8 series started out initially with the PDP8 which was a redo and rehash of the PDP5. The big change there was from the soldered connections and the system modules of the PDP5 to the first use of the wire wrap technique and the use of the flip-chip modules. The flip chip modules didn't have a protective handle around them. They had this little thing on the back that was used to pull the module in and out and sort of provided spacing so each module didn't short each other. And the connections, or the connector, through the back plane was gold plated directly on the board rather than having a separate amphenol connector. It had to be riveted. Incidentally, before we get off the PDP5, one of the modules in the PDP5 with which I was very familiar, was the main register module. This had two sets of connectors; one in the front and one in the back. We had some modules that we had a lot of problems with. On the bench, they worked fine: we had this little tester box for every module that came in. We put it in the tester and this tester took signals from something called the signal generator and these signals were applied to the module and we could check it. These particular modules worked fine on the bench but when we got them into the machine they didn't seem to want to work. And it took

us the longest time to figure out why, but we did: they were an eighth of an inch too short. And the back connections didn't want to work. Talk about your mechanical control problems. In any event, when we went into the wire wrap business with the classic PDP8, the modules were still made out of transistors and diodes and resistors and capacitors. At the same time we developed sort of a general series of modules called the R series (they had red handles, "R" meaning "Red"). That was a line with which I was very familiar because at that point in time that was the thing I was doing. The whole R series was my baby.

Then there was an S series that was actually used in the PDP8. The only difference between the R series and the S series was the size of the red resistors used. That was the stock that was used to make the PDP8 because everything was hand layout.

You asked what the letters meant. As I say, the letter was a little bit arbitrary, but the first PDP8 was built with S series modules and wirewrap technique and was a parallel machine. Someone got the idea of building a less expensive PDP8 which was serial, so this was the bearer of the 8S. Some people claim that

S stood for "serial," others claim that it stood for "stupid," and even others claimed it stood for something well, let's not do that word. So the next PDP8 that was built was with integrated circuits, and these were small scale. T.I. and Fairchild, both, started coming up with integrated circuits. So there was a big changeover within Digital. This was the first of the modules that had connections on both sides; in other words, it was a double-sided module, because you needed that to be pins to get the signals on and off where there were circuits involved. So we started coming out with a so-called Mseries model which was Magenta. The chief mechanical box designer, I remember, (his name was Olsen) got into the act late in the game on the PDP8-I, and redesigned the front panel and a few other things. But anyway, the 8-I was the next in the series. "I" I guess stood for integrated circuit. T.I. won out over Fairchild. Fairchild, I think, lost the battle many many times, by the way, because of a series of good engineering and poor management. That in itself is a whole story.

PDP8-L was a low cost version of the 8-I. The 8-I had lots of options. The 8-L was a bare bonde machine that had a few things.

At that point Louie Klotz was very involved in PDP8 check out (there's another name for you to pursue, he's in Salem somewhere now). Louie Klotz and I went to work for Lorrin Gale and the objective was to try and design a large scale integration version of the 8. And so we were working with National Semiconductor, trying to develop something in PMOS that was early in the game, and that kind of fell through after a year and a half. So we went back to the original PDP8 group and designed something called, no I take it back, the next one I think was the PDP8-E right after the 8-L. The idea being that instead of having all of that random back plane wiring that was done by Gardner-Denver, with medium scale integration that was available at the time, we could put all of our logic on a larger PC board and plug it into the back plane which had a bunch of parallel wiring. And we kind of carried that off. It would have been nice if we had better, bigger chips or something, chips that could do more stuff. The 8/E really worked pretty well. Then Louie Klotz and I went off to design this P8 or 8PF, whatever we called it at the time. Then there was this other PDP8-A that came afterwards so that kind of fell through. The large scale integration, it wasn't really ready for it. The technology wasn't right, and so not much

happened from there, so that was a year and a half hiatus. One of our objectives all along had been to build a PDP8 small enough so that we could put it inside of a shoe box. So we had a shoe box size thing that had some switches on the front of it; you may remember the front of the PDP8, it had a whole bunch of switches. Inside the mock-up we had stuffed one of these spring snakes, the kind that that you squish down and put inside of a can, so that anybody that was nosy enough took the lid off and let the snake fly out into his face. Gordon Bell came by one day and wanted to check the insides of this thing. He started take this thing apart and the snake missed his nose by about three inches; that was our biggest achievement with the snake. In any event, after we got through with that IC project it kind of fizzled out for a couple of different reasons. We went back to PDP8 land and designed the PDP8-A. By this time there was enough medium scale integration stuff on the market and we could do a pretty reasonable job of squishing the process on this single head. So we basically used the same box that we had for the PDP8-E and we built some more boards to go with it. So the PDP8-E was kind of a modernized version of the 8-A. And then some people on the west coast had picked up the idea of the PDP8. There was a guy

who went to Stanford and did a doctoral thesis on building a PDP8-A on a chip. The only problem, as far as we were concerned, was that we were going to have to pay five times the price that we would with the medium scale integration stuff, to get the speed, so we weren't very interested. And finally they came in here with their 6100 and told us it was going to cost us \$30 a pop; suddenly we started understanding that maybe this was a worth while thing. In the meantime Tom Stockebrand and his guys had developed this VT50. There was a VT05 before, then the VT50 (several varieties of VT50's) and then came the VT100.

Stockebrand was one of the prime designers of the VT50. Anyway he wanted to take the VT52 and to stick the 6100 in the processor and memory and call it a VT78, which we did. Meanwhile, there was a rather interesting thing that was going on. Somebody had taken a desk and stuck in that desk a PDP8-A and an 8" floppy disk and some software and this was the first of machines that would be WPS (word processing systems). WPS which kept the PDP8's alive along time after they were dead. The PDP8 was of ancient architecture by this time; you had to understand the PDP8's were designed in an era where the flip flops cost a lot of money and took a lot of real estate. I mean, if you've got two flip flops

on a single flip chip module, you're doing pretty well. So flip flops were not at as much of a premium but the software that had to run on this machine was there, and we knew the machine would take a lot of time, in fact several of us commented knew very well that if the guys in PDP11 land had ever bit the bullet and sat down to write WPS, they could have put us out of business. But they never did. So we kept riding their price performance curve right down the tubes; in fact, there are still PDP8s out there, in DECmate III. The VT78 was the first of the PDP8's that had no memory other than semiconductor memory: by this time memory chips were inexpensive and reliable enough that core memory prices were going out of sight. And we had reasonable disks too. The floppy was a good way of loading the program so core memory was pulling out at this point. So the question was, what do we do next? Well one of the things that we decided to limit was....

Q: The core memory was phasing out because of the high cost to manufacture or because semiconductor memory was superior?

A: To both questions, yes. It was a question of how much does

this memory system cost you vs. how much does that memory system cost you assuming that you could make both of them work reliably. If you can't make memory work reliably, forget it. The cost of core was going up. It's very labor intensive. If you've ever looked at core memory, look at all those little wires, three wires going through the core, a very expensive process. And suppose you have one bad core in the bunch? So the cost of core went up, the cost of semi memory came down and so we started using semi memory.

Q: You were talking about the VT78?

A: Well the question was, what do we do after the VT78? The end result was that Harris Semiconductor and I together worked on something called the 6120 which is the central processor in all of the DECmates: the VT278 and the DECmate II and the DECmate III. <And it was interesting that half of the central processor in the Classic 8 is equivalent to the size of the 6120.>

Q: When you were working on the whole 8 family , what do you recall about the reception on the street?

A: I really wasn't in a position to see. Most of my work with the PDP8 was of course with the individual modules, and I can well remember some problems that we had with the individual modules. Like all engineers that have to provide that kind of a service to the PDP8 check out line, the only computers that I ever got to see were sick computers. As far as I'm concerned one of the smartest things they did was, there was a plywood board, the PDP8 was bolted to the plywood board and right across the front was the serial number of the machine and I noticed serial number 49. Now, the next time they called me over I noticed a serial number in the 200's, and it finally occurred to me that in the meantime a lot of stuff had gone out the door, so things were happening.

One of the classic goofs at the time was made by Monsanto. It turns out that all of these IC's that were produced on silicon material. What they do is they grow something called an epitaxial layer on top of the wafer; this is where all the interesting stuff happens. Well, Monsanto blew it, instead of making an epitaxial layer on top of this stuff they managed to produce something that was in reality a very bad diode. They sold these

wafers to Fairchild and I think some others as well, and Fairchild promptly proceeded to make high speed switching diodes. So what we really had was two diodes in series: one of which was very fast, and the other which was very slow. Now this was all fine and Fairchild ??? testers and diodes and it worked well and they shipped it. And this was fine and dandy until we got into a machine and the diodes sat there for a long time in what we called a reverse bias condition. And it happened: the big old diode at the epitaxial area started to come into play. It shut off and the only way to turn the diode on again instead of responding in 10 nanoseconds or so, it took it the better part of a microsecond to turn it on. ????? and one of the places the most critical places was in the core memory and these diodes were used as part of the memory switch and if you sat in one address and did a lot of things which is quite common and you went to access a different address, one of these diodes would refuse to turn on, and that meant the first read that you did to the address. Now that's disasterous. Because anybody who knows anything about core memory technology knows that the first thing you have to do is to read the cores, and in the process of doing this you reset all the cores to zero in that address, and you go

write back in. It's a read-write operation. It's a destructive read out. We finally built a tester box to test those memory diodes and we caught the rest of those bad diodes with a logic test. The technicians were pretty good at finding those kinds of problems. We thought we were the only ones that had this kind of trouble, but one day Fairchild sent a letter of apology out and in doing so, they managed to stuff the letter addressed to Computer Control Corporation into an envelope addressed to Digital Equipment Corporation. We got the letter of apology from Fairchild addressed to another company. We said, well I guess were not the only one seeing this problem. I used to periodically tug the chain of the semiconductor manufacturer; his name was Vince Sabella. He came in one day and I remember showing him a problem that he had with one of his semiconductors. It turned out the way the transistor was made it had three leads staked around a ceramic donut, 3/16 of an inch in diameter ??? and a silicon dome at the top it was cemented on. ??? the leads were not staked to the donut properly. So, if you pushed on the thing just right, lead was worked out. Pop the top open, we refer to it as a flip top box. Building a transistor really was the black art of the mid 60's. And when things went wrong they

went wrong in a big way.

Q: Did you work on the 11's too?

A: Somewhat, not a whole heck of a lot because at that point it was 8's going this way and 11's going this way; there was a little bit of a parallel path.

Q: I know that in terms of word length the two machines were different, but how would you explain the difference between the 8 family and the 11 family?

A: The 11 was a totally different departure physically. We could take advantage in 8-land of things the 11 people were doing and couldn't quite pull off because the machine was more complicated. With four more bits (we had 12 and they had 16) sure, you got 33 percent more bits but you also have THAT squared of problems. You can do a lot more things with those bits than they did and so their architecture grew where ours didn't. At that point, you see, the PDP8 was an established machine. We were riding the price curve. We knew what the performance had to be and we made

at least had to do that ????? we may have added a ???? but we'd never take anything away.

The 11 got going because they had to do something with the hole that existed between the PDP8 and the PDP10 and the PDP4/7/9/15 series was kind of a limited architecture, there was no question about it. They could only go so far with that and those were pretty big machines; they were almost as expensive to build as the PDP-10. The PDP11 of course expanded upward and supplanted the 10. But we elected to try and follow the PDP8 down to the point where it became a competitive commodity machine. Our advantage always was we could do things faster and for less bucks. The PDP11 can't, that has always been true. The problem with the 11 was that we had these extra bits and we had this extra memory floating around so we'd get sloppy do things in a sort of general way. Classic case of point: somebody made an FPP12 that hooked onto a PDP12 and it ran floating point. I think it was Richie Lary, and some of his friends, who ran a Fortran 4 compiler. At the same time some people were writing Fortran 4 compilers for the 11. Richie, being the kind of guy he is, (Richie's a great man, he's certainly one that you've already

heard from and you'll hear alot more from him). You could never find Richie, because he would come in at 10:00 every night and he'd leave at 7 in the morning, a night owl, always operating at night. So if you wanted to get in touch with Richie you would leave a message, scotch taped right across his doorway -- literally -- so he had to run into it. But the next morning sitting on your desk would be an answer. He wrote the compiler. At the same time, people were doing it for the 11. And when it got done, they had a drag race and the FPP12 ran circles around the 11 to the point where it was a total embarassment. But the PDP8 series always came back. We elected to drive that thing, we tried with some of the PDP11 turf a few points in time and we were told to keep our hands off.

<Regarding terminals and printers,> Stockebrand would certainly have a lot to say about the DEC tape, LINC tape and terminals stuff.

<These days> Lou Klotz is somewhere in Manufacturing, I don't know where he is. He started out as a PDP check-out technician. One of the things that a technician would always do in the

morning with the PDP8 was he got the key lock and he'd spread those terminals apart and then he'd turn on the switch. Well, there was an interesting trick: this guy or several guys would get down amongst the rats nest of wires and they would slide on one of these little AG flash bulbs it was just big enough to spread the terminals and connect them to one of the power supplies. They stuck two or three of those things in there and a guy would come over and turn on the switch and a blinding flash of light would hit him in his face and he'd dive from the switch to shut off the power. You'd think for sure he'd blown himself up.

The PDP8 was the first hard luck; it was absolutely and totally stupid. In order to get it to be able to read paper tape, you'd have to toggle in about a dozen instructions. Anybody who played with an 8 knew what those instructions were: 6032, 6035, 357, 7106, 7886, 7510, you got so very good at toggling those things in. The most disheartening thing was to fire up this chunky teletype, the ASR33. When you get through reading this tape and it all worked well, the accumulator lights read 0: that was the check sign. And it had better be 0 -- and if it wasn't you had

to start from scratch! Reread in the binary loader, retoggle,
reread the tape and the whole check sign. We had some sad
experiences with some of the ASR33's.