

The *Analytical Engine*

JOURNAL OF THE COMPUTER HISTORY ASSOCIATION OF CALIFORNIA



Hewlett-Packard 3000

HP Archives

Volume 3.1

Fall Double Issue

November 1995

The Analytical Engine

JOURNAL OF THE COMPUTER HISTORY ASSOCIATION OF CALIFORNIA

Editorial: X-VICTORY!!

We scarcely dared to hope, but we dared to work, and work brought forth what hope alone never could have. The magnificent SDS 930 came home to California....by a whisker.

On August thirty-first we received e-mail from the Federal agency that owned the 930, saying that for administrative reasons, a firm acceptance of permanent loan would be required within a month. Well, move it or lose it. The CHAC had been working on this since late last year, and every request we had made in e-mail, all the pep talks we'd given at conferences, every stern warning we'd published in the ENGINE, had netted....a few hundred dollars; and we couldn't hope to load, and ship, and unload the 930 for less than a few thousand. Could we raise, in a month, what we had not raised in half a year?

We have three phone lines here, and 1,600 names in our database. The only answer was obvious, arduous, time-consuming, expensive, and a bit scary — but it was the only answer. Telephone calls, confirmed with e-mail and backed up by stamped return envelopes, brought in over \$7,000 in pledges by the end of September. It was more than we ever thought we could raise — and yet it was barely enough, because as we progressed from mover to mover and described the job, the estimate of costs almost tripled. Yet energy, community and generosity prevailed.

At 8:30 on Wednesday morning, October eleventh, a moving van pulled onto a blacktopped lot in the South Bay and finished its fourteen-hundred-mile trip from the Space Environment Center in Colorado. We joyously took delivery of fourteen racks, two Teletypes, a console, a VDT, and approximately forty boxes of tapes, docs and spares, to a total of 10,300 pounds — which only (yikes!) filled about a third of the truck. A picked crew of CHAC stalwarts spent two hours removing

ductape, non-historical stickers, deteriorated insulation, etc., and eased the whole computer into a locked, ventilated steel container. We're planning a Tiger Team Party in November to clean, dust, lubricate, and otherwise prepare for long-term storage with zero deterioration. Although the amazing thing is, how clean and pretty the 930 is even now....

This has to be counted a stunning success — meaning that it stunned *us*. It's true that the CHAC was in an advantageous position to raise money, since we had a good cause, a serious deadline, and a great list. Even so, in the two weeks that lifted us from bare hope to the edge of triumph, we felt dizzying relief. The words "emergency" and "computer history" don't often occur in the same sentence; we had no idea what might happen when they did.

The donations were as encouraging in their pattern as in their total; they ranged from \$20 and \$25 to literally thousands of dollars. It proves at a stroke that support for the history of computing exists at every level, whether corporate or individual, and across a whole spectrum of professions.

In February the ENGINE will publish a photo feature on the 930 with a Heroes' List of donors who will consent to have their names published. Meanwhile, to everybody who donated, advised, drove, e-proselytized, or got their hands dirty, our most heartfelt thanks and regards.

LATE BUT GREAT (we hope)

Through clenched teeth we admit that, even by normal standards, this issue of the ENGINE is late. The reasons are straightforward: The fundraising campaign for the 930 Rescue, and the writing and installation of our Web page, together required about as much time and energy as one ENGINE. We didn't have it in reserve and magazine production had to slip. Three points follow from this:

- 1) For the first time, as we're publishing one issue, we have almost all the material in hand for the next. November may be late, but February won't be *as* late.
- 2) The CHAC can only increase its resources of time and energy by enlisting volunteers. We are eternally grateful for those we have; we can always use more. Please contact us if you can spare even a couple of hours a month.
- 3) We promised thicker ENGINES and we're delivering. This issue starts a provocative interview with Lee Felsenstein and winds up the HP's EARLY COMPUTERS series with an article by Chris Edler on the development, release, redesign and re-introduction of the 3000 series; the letters and queries are here too, along with a review of Max Maxfield's superb introductory text on microelectronics, and a lot of current news that we only had time to summarize. More people seem to be interested in writing, they are *not* running out of computer history to write about, and we'll continue to publish as many pages per year as our cash on hand permits. Enjoy!

CONVIVIAL CYBERNETIC DEVICES: From Vacuum Tube Flip-Flops to the Singing Altair

An Interview with Lee Felsenstein (Part 1)

[Lee Felsenstein — sysadmin of the pioneering wide-area network Community Memory, contractor for Processor Technology, hard-working moderator of the Homebrew Computer Club, and designer of the Pennywhistle modem and the Osborne One — has lived a life constantly interwoven with the history and philosophy of computing. Through it all he's pursued Ivan Illich's ideal of "conviviality," insisting that technology attains its highest purpose only when it becomes understandable, approachable, repairable and usable by ordinary people.]

This first part of a projected three-part interview takes us from 1959, and a computer club without a computer, to 1975, Homebrew, and Steve Dompier's history-making implementation of "Fool on the Hill."]

CENTRAL HIGH

KC: Let's start with flip-flops in high school....

LF: Yes, at Central High School in Philadelphia; it was the college prep public high school — actually the third high school in the nation, founded in 1838. Until 1983 it was a boys' school but it's since gone co-ed under pressure of a court decision. But in 1959 my brother Joe was a senior and I a freshman, and he founded the Central High School Computer Club. He had been looking in books from the forties that he found in the library, with vague circuits in them of gates and flip-flops, and he wanted to build a computer — I think because he wanted to program one. Of course computers in those days had a tremendous cachet, they were UNIVACs, the "giant brains," they would occupy half a floor of a building. He would show me these books and say, "Can you build this, can you build this?" — and while I had my doubts, my relationship with him was such that it wasn't legitimate for me to say no, or so I thought; it was just my position as a younger brother, to always be tagging along. So he founded the club and installed me as the chief engineer, and we would take 6SN7 tubes — dual triodes — and mount them in sockets in

cigar boxes. I had a basement workshop at that point, I was familiar with the tools, and I even had a rather crude audio oscilloscope. And we tried to make flip-flops so we could make logic elements, counters and registers and so forth. Probably they would have worked quite well, except that I was completely ignorant of so many fine points. For example, we were trying to create pulses for counting by flicking wires against Fahnestock clips on the cigar boxes — and I can't think of any better way to create a random stream of pulses than that! But that point never entered my mind and I didn't really have any mentors who would point this out. We certainly could have used help from an engineer, preferably in the computer business, but we never really considered going out and asking for it. It had only been two years since Sputnik, and the school was dominated by a competitive aspect that ruled us all. I mean, some of the boys were keeping running grade point averages up to three decimal points in the backs of their notebooks, that sort of thing. We were all supposed to go claw and scramble over each other and get into the best colleges. Actually, that kind of thing had been happening at Central High School for generations.

In any case, the Central High School Computer Club managed to produce a couple of two-bit counters that probably worked about fifty per cent of the time, using our techniques. I concluded from that whole experience that computers were not for me. I knew I had a future in electronics — I wanted to be an electronics technician at that point — and that gradually evolved into, well, I'm going to college so I might as well study electrical engineering. I had gotten into electricity and electronics at the age of 12, when I read in a chemistry book that if you wanted to be a chemist you had to learn a lot of mathematics, and I was having trouble with percentages at that time. Only much too late in the game did I learn that electrical engineering required the most math of any engineering. So much for career planning! My experience with the Central High School Computer Club made computers simply too frightening for me, from a hardware standpoint, because they had to get it right every time, every clock pulse. I had no particular interest in the software, so that turned me away from computers for a while. (My brother, incidentally, went on to become a professor of genetics, specializing in mathematics of genetics, population genetics and so forth.)

BERKELEY, FSU, AND AMPEX

My next experience even close to computers was in 1965, when I worked with what was called the Free Student Union, which was the successor to the Free Speech Movement at Berkeley. I learned how to punch cards on an [IBM] 026 card punch. You could walk right into the card punch room and use it without any questions asked, so I punched and sorted the membership lists and we used them for local meetings that nobody came to.

At the end of 1967 I had dropped out of school, I was in a psychological depression, and in the beginning of '68 I went off to work as a junior engineer at the Special Products Division of Ampex Corporation; then at the end of 1969 I dropped out of that job, and then dropped back in when my money ran out.

Ampex is almost gone now, but it was a big company in those days and a division might have two or three thousand people — the Magnetic Tape Division as an example. The Special Products Division, which was a hundred and fifty people in a little building on the Redwood City campus, took responsibility from beginning to end for one-of-a-kind products. We might just take an item from stock and paint it purple, or we might design something from the ground up. We'd gotten into the design of a large-scale audio-visual instructional system, called a Pyramid system, controlled by a [Data General] Nova minicomputer. About three of these were built. Well, in 1970 I was put on this project — that I didn't want to be on, because of that computer controlling it — and I was rather quickly given responsibility for catching the bugs in designs that other people had prototyped, and then teaching other people about how they worked and to support the ongoing effort; it's called maintenance of product design, MPD. I started with the tone detectors, which were nice comfortable analog things; but I moved on to something I had done previously, which was design of audio circuits for high speed tape duplicators — IS oscillators at 1 megahertz, that sort of thing. In effect it was hi-fi design at 40 times frequency step-up, which was very interesting. But now I couldn't avoid getting into computers....

KC: Right.

LF: And first of all I was sent to class to learn BASIC at the Service Bureau Corporation, SBC, which was the setting for some very important realizations. The instructors at that place were young guys in three-piece suits who liked to show off how much they knew. And they would say, "You see how the system sort of hiccupped there and it's running more slowly? That's because the computer in Los Angeles has been turned off and it's been transferred over to the computer in Kansas City." This was my first experience with networked computers, and I understood that geography could be irrelevant on a computer network, providing it extended to the various places. They also showed me how files could be made public with levels of accessibility, by pre-pending asterisks to filenames; three asterisks on a filename and anybody in the system could get the file — in effect, it had been published. So I saw a hierarchy of distributed publication possible there.

KC: *Now, what operating system was this?*

LF: I don't know what operating system it was. I presume that SBC had won some kind of anti-trust suit with IBM and somehow gotten an IBM network out of it, probably running on [System/] 360's, but I don't know the details beyond that and you may have to look them up if anybody's interested. Almost certainly these were IBM computers, we used Selectric terminals, 2741's. All I knew was, I was learning BASIC. But I still had some important realizations there, and several of them related to media and publishing.

I had spent much of my time in Berkeley working for the underground press, on the grounds that this was a certain kind of community media — media with less hierarchy. There was still a hierarchy, by definition, because that's built into any publication system, whether print or broadcast, where the information is replicated identically and sent out from a single point. In fact I had given up on the underground press because it was subject to many of the same tendencies as the mainstream press, including centralization, and the temptation to make a spectacle of itself in order to increase its circulation. That wasn't the game I was interested in playing, or assisting. I wanted to help define, and help create, media that contributed to the development of a local community. And now I was presented with networking technology that could

be a foundation for communities of interest — communities that could exist in defiance of geography. We could, at least in theory, untie the knot that tied a single person to a single community. This was a vision that I wanted to elaborate, and continued to pursue.

Meanwhile, at Ampex I also learned enough machine-language programming on the Nova to do little drivers and so forth. By the end of 1970 I did the programming to demonstrate the control structure between, if I recall, 12 key pads and 12 buffer tape recorders. The system worked through a series of master recorders, each of which had 8 tracks on quarter-inch tape and could run bi-directionally at high speeds. A matrix switch with reed relays sent output to a buffer recorder at each student position, or carrel. Each carrel had earphones, a video monitor, and a 12-button key pad whose interface I had designed. You would sit down at the position and punch up the number of a program you wanted. That would set up a transfer between the track on the relevant master unit and your buffer unit. So the master would energize and do a high-speed tape duplication, 40 times speed, across to your local machine. In a minute and a half [of duplication] you'd have an hour's worth of program copied over. I think the transfer was carried out backwards, so that when the duplication finished, your buffer machine was at the beginning and would start playing, and you would hear the program material on the earphones. Buried in the program, meanwhile, were periodic bursts of 55 Hz tone which a tone detector picked out and turned into a slow serial bit string; and this would command the computer to do various things — not only tape motion commands, but also primarily transfers of pictures from a moving-head video disk recorder with a 14-inch platter to a set of buffer tracks on another video disk, so you'd get a still picture transferred from a repertoire of about 300 still pictures. That picture might present you with a question, and the tape would be stopped, until you responded on the keypad. The keypad, like the rest of the system, was interfaced to the computer, and depending upon what your answer was and the instructional program that was running, the tape would either continue with the right answer, or rewind to another track and play back an explanation of why you had been wrong. It's interesting that, by 1990, I was able to build all this functionality into a unit that attaches to your belt,

works off a compact disk, has a private-eye display and headphones, and audio synthesis and speech synthesis and audio playback, and a little pointer that you hold in your hand — an isometric pointer; so that's a closing of a circle, in effect. But the machine language programming that I had done, that allowed the buttons on the keypads to command the tape recorders to do the transfer and then start playing, saved my job [at Ampex] during a reduction in force that was going on.

So I gradually became more familiar with computers, through designing interfaces. I also did a design analysis program in BASIC for the active filter components in the tone detector. Those things are easy enough to design, certainly out of cookbooks; the hard part is to build them so that the filters' performance remains acceptable while components range over their tolerances. I had been given a design with ideal values, and as part of MPD I was supposed to figure out how to build it with actual parts so that it would work every time without adjustments.

KC: In other words, it'll work this way if it's perfect, but what about when it's not perfect?

LF: That's right. Engineering is mostly about how far off the ideal you can stray and still be safe. So I constructed a test suite with a lot of nested loops that would vary each component through its possible range — that's on the inside loop — then go out and vary the next component one tick, go back in and vary the first one through its range. There was a nested loop for every component. Running this in BASIC on the Nova, I was able to demonstrate that no possible set of combinations for that design would work over its entire range. Meanwhile, of course the units that had already gone in the field — the prototypes, in effect — were giving all kinds of trouble. So they sent out an engineer who was clever enough, and simply didn't want to hear about anything else; and he sent back a cardboard extension to the board wired with potentiometers and one-shots and some other parts. He had done almost everything that I wasn't allowed to do, and I was told "Okay, just put that in the design and build it." Naturally we did wind up with adjustments, and that was an excellent lesson in the relationship between engineering and management.

KC: They realized that there had to be in-line adjustments built onto the board to compensate for the drift of the components.

LF: Basically, yes. The problem, though, was that this didn't work particularly well either. What he had done, in effect, was tune to the performance of the master tape recorder — the AG-440 master. We proved with a recording oscillograph that feeding the [55 Hz tone] bursts into the master tape recorder caused the circuitry in there to go into an oscillation, so that the baseline of the signal would wobble drastically, and the audio jumped around at a very low frequency.

So this guy had been totally empirical, and tuned his circuit to that bouncing behavior. He just scoped it and caught the level. This worked until we changed the master tape recorder for the next system, and suddenly it stopped working. I could tell that this was getting us into trouble and I began to advocate using automatic gain control — control of the signal's variations. I kept an automatic gain control in development as long as I could. I had to stand up at review meetings and be told "Kill that, freeze the design the way it is and go on." Eventually I did as I was told.

A couple of years later I went back to see what had happened, and discovered that an engineer named Al Alcorn had been given the circuit to fix up, apparently by putting in automatic gain control and increasing the size of it from one board to two — in short, the things that I had wanted and tried to do.

KC: This is the Al Alcorn —

LF: — who co-founded Atari. And in fact, a year or two later I was sitting in front of his desk applying for a job. Al was looking for manufacturing engineers at that time. He didn't remember much about the tone detector, but it was certainly an embarrassing and instructive moment. That event and a lot of others have kept me from ever really feeling good about the wisdom of management.

At the end of 1971, I had been doing psychotherapy for a while and things were getting better. I took educational leave to finish up my degree. Since I was on the co-op work/study program, I was able to convince myself that my combination

of class work and experience had put me five years ahead of my classmates; I had eaten up four of those years, from the end of '67 to the end of '71, so it was time to get my degree — which I did, in June of 1972.

The Pyramid System project was moved into the Videofile building, where Alcorn was, and I spent a few weeks commuting down to Videofile. Then I left, and Ampex basically collapsed. Everybody got laid off.

WIDE-AREA GOES PUBLIC

Meanwhile, in 1971, a guy I had known in my activism days came running up to me and said, "There's some people over in San Francisco that have a computer, and they're going to use it for non-profit activity." I decided to investigate. This was a group of four computer science students who had left Berkeley during the Cambodia crisis, and essentially dropped out of the system as it was; they had come to rest at a "warehouse project," or warehouse community, called Project One in San Francisco — where a group of people basically formed an unincorporated non-profit association and leased a warehouse. Within this Project One, these four guys created Computer Group One, and were trying to deliver computing power to people and organizations who lacked access — to non-profits, to radical groups, to social-action groups of every kind.

Through an inspired act of hustling carried out mostly by Pam Hardt, they managed to get an obsolete time-sharing computer — an XDS 940 — donated together with money to set it up and run it. This machine was serial number 4 and was donated by Transamerica Leasing; it had been down at SRI and apparently run a robot known as Shakey. We understood that name very well, because the 940 had a DMA channel which occupied a whole cabinet, and we went deep into it and found the terminators for the bus were wired to the wrong voltage — to four volts instead of eight, which put them square in the transition region. That would make anything connected to it "shaky." I think this 940 had also been used by Doug Engelbart, so it had had an interesting life.

They had also begged a time-shared BASIC and were starting out with time-sharing, which was

ambitious to say the least. I signed on basically as chief engineer. I was to be trained in system maintenance by a guy named Al Montoya who lived at Project One — it was living and working space both, kind of like four stories of lofts — but Al disappeared the day it was installed and didn't turn up for months. When I complained about his leaving all he said was "Here I am." Meanwhile I had absolutely no tutelage and no source of any; for one thing those computers weren't all that common, there were only 58 ever produced including conversions of 930's and 9300's. The language we used for our information retrieval system, called QSPL, was only available on 940's or on IBM 360's which wasn't necessarily a time-sharing implementation.

KC: Now, at that point the computer had been moved from SRI to San Francisco?

LF: It had been retrieved. Transamerica Leasing had leased it to SRI, who obviously kept it for the requisite three years. At the end of the lease Transamerica didn't have another taker for it and chose to, in effect, donate it or long-term loan it to this non-profit, Resource One. So it had to be delivered in a couple of trucks; it was a row of about a dozen 19-inch relay rack cabinets, each one about two feet wide — 19 inches is the internal dimension — and so 24 feet of cabinetry. This machine also required 23 tons of air conditioning. We had to run our own power lines from the main power busses downstairs to keep the thing happy, you didn't just plug it into the wall. It had a fairly big line printer, a console that sat on a table, and a Model 35 Teletype for the console terminal. The multiple serial interface handled 8 lines, I think, and I had to design and build a rack of modems where we could collect, if I recall, four modem lines. It was usable by ASCII terminals, mostly Teletypes which we'd begged from Tymshare after their lease contracts had expired. They'd been used pretty heavily and when we got them they were kind of sticky.

I think that, to today's computer user, the most amazing thing would be the disk file we bought for it. This was a double width cabinet, so it was about four feet wide, and it had two pull-out drawers each of which would hold a 2314-style disk drive with a stack of 14-inch platters that were removable; you took a plastic dome cover and slipped it down over the pack, and you could detach the pack and pull it out. We got one with double density, so it was 58 megabytes. To fill just one of the two pull-outs cost us twenty thousand dollars.

KC: You say like a [IBM] 2314, but not in fact a 2314?

LF: It was manufactured by Control Data, I think. We worked with people from Systems Concepts, which was a mile away in San Francisco, and Fred Wright — I think — designed and supervised the construction of an interface which emulated the Data General Nova back panel, and let us plug in a Nova disk controller that we had bought. I did the mechanical design of the whole thing, and that and the interface got it working. The system also had a swapping drum, hundreds of fixed heads arranged around a big conical drum, which would spin up; centrifugal actuators would fling out and the drum would hop up into engagement with the heads, because it was tapered towards the top. But this drum was flaky — the surface was deteriorating, and nobody knew how to fix it. We had a whole bunch of problems! The Ampex TM-4 tape drives had gas thyratrons driving little solenoids that actuated pinch rollers, and if you missed releasing one of them you could snap your tape, because the tape was going in one direction and the other pinch roller would come in pulling in the opposite direction. Mostly it would just squeak and stay there, but it could snap the tape.

KC: These were capstan drives, not vacuum-column drives?

LF: No, they had tension arms to buffer the speed variations. The tape wound back and forth along these little roller guides in the arms and interleaving stationary guides. So the arms followed these arcs and their positions were sensed, and when they got too far out, the tape would be speeded up. So no vacuum columns. I muddled through somehow, and certainly it was all a good education in antique electronics. But I never really got trained

strategically, to answer questions like: What is the structure? What is going on inside this computer? Those answers would have been distinctly useful, but they had to wait till later.

ROGIRS, A PUBLIC DATABASE

In 1972 we had the thing installed, and we embarked on writing an information retrieval system. Richard Greenblatt came through town at just that time and gave us a combination lecture and pep talk with a focus of, "Let's write an information retrieval system in one day." And he sketched out a free-form keyworded information retrieval system whose data could be arranged according to new index words on the fly. The easy way to do an information retrieval system is to determine what words you're going to use as indexes, put up a list, and have the user enter an item by choosing from the list. The hard thing is to make that list expandable in the middle of everything, so Greenblatt lectured on how to do this. I had brought a friend, a systems programmer named Efrem Lipkin, who had an appropriate background, and he took on the direction of the project. The final code was called ROGIRS, the Resource One Generalized Information Retrieval System.

We were going to use this for the benefit of the Bay Area switchboards — volunteer information and referral agencies. Anybody with a card file box and a telephone could set up as a switchboard on whatever topic or in whatever area, publish a free notice in the underground press, and be in business. The trouble was that, for the great majority of switchboards, the filing system existed only in one person's head. When that person got burned out and left, which was only a matter of time, another person would have to come in and start the filing system from scratch. It was quite inefficient.

DIFFICULTY AT THE BEGINNING

Now, Resource One had been attending meetings of switchboards in San Francisco, because they had taken over the corporate shell of something called the San Francisco Switchboard, the other half of which had become the Haight-Ashbury Switchboard. We were going to give the switchboards a common filing system and enable them to share resources, thus the name Resource One. We took the better part of a year to get that retrieval system written and debugged, at which point we went back to the switchboards, and discovered that all our contacts had burned out and left and nobody remembered us — although one person in the meeting apparently remembered hearing about us. So for practical purposes we walked in cold, all smiles, proposing that each switchboard pay \$150 a month to rent a teletype and a modem, so that they could then laboriously key in their files, and only then could their users get anything out of the system. Nobody knew how to deal with this, none of the switchboards had \$150 a month to spend, and we were not really treated seriously.

So we then had a reasonably powerful, empty information retrieval system and went about looking for things to do with it. We talked to many people, including some librarians at the Bay Area Reference Council — which is the library of libraries; if your library doesn't have a book, the Bay Area Reference Council is the organization through which it gets the book from another library. The librarian understood what we wanted to do, but said "You have something a lot like a set of shelves with no books on it. Why don't you get some books on it and see what it looks like?" Efreem was inspired by this and had the idea to set up a public terminal in Berkeley, where he lived.

At that time, about June 1973, I had burned out from the stress of living and working at Project One. Efreem and I realized that, if we could establish a combination office and apartment in Berkeley, I could live there and disengage myself from the pressures of communal life, and Efreem wouldn't have to commute to San Francisco. We proposed this to Resource One and convinced them to finance it.

ON LINE AT LEOPOLD'S

By August of 1973 we were ready to put the public terminal up. We had a possible location in a record store that was being run by the [UC Berkeley] Student Union, called the Leopold Stokowski Memorial Service Pavilion — later it became Leopold's Records. We went to the Student Senate and said "We'd like to put a computer terminal for an electronic bulletin board there," and they said "That sounds great, why don't you do it?" I built a foam plastic enclosure for a Teletype 33, with a clear vinyl flap that attached with Velcro, to muffle the sound of the print hammers. It had two ports in the front covered with overlapping vinyl flaps, like a cat door, so you could stick your hands in and use the keyboard. Now, Berkeley didn't have local phone service to San Francisco but some Oakland exchanges did, so we used an Oakland prefix number installed at the store — apparently it was a residential number and I still don't know exactly how this worked — to make one call per day; that is, in the morning we would dial up the phone and plug the handset into the modem, and have a free connection to the 940 in San Francisco all day. We were using an Omnitech modem, or something like that, which was good for 300 baud under the most favorable circumstances. I remember that the modem cost \$300, which would be the equivalent of \$600 to \$900 today.

PENNYWHISTLE MODEM

Then I started exploring the possibility of building a modem ourselves, and I was encouraged by Fred Wright and the folks at Systems Concepts, who said, "Well, essentially all you'll have to do...." These were MIT guys, who gave me that word "essentially," which really meant "and someone else will do all the hard work," like the MPD engineers. But I didn't know that; I just went ahead, tried it out, and hit upon something.

The first part of building a modem is to convert digital data to tones, and that's easy. The hard part comes when you have to detect the two tones and discriminate between them. In essence we have a phase-locked loop oscillator circuit, which produces a voltage that varies with the frequency of what it's locked to. Against this we need a reference that says, when the voltage crosses this line, the bit changes from a one to a zero. The hard part was setting that reference and keeping it set, because it would be perturbed by component tolerance, by temperature, and by any number of other factors.

KC: You were trying to draw a line down the middle of a graph, and if that line was any less than straight, you risked misinterpreting bits.

LF: This is, again, "how far off the ideal you can stray and still be safe." You can draw as straight a line as you want, but the signal in the graph is going to meander somewhat, jiggling back and forth all the time. You're dealing not only with short-term oscillation but with long-term drift. What I realized was that the modems of that day — and mostly of this day too — were all specified down to zero baud. That would be preferable for something like a burglar alarm, where you have a rented phone line, a modem and a switch, and if that switch ever opens you want to know about it at the other end. But that's not what we're using it for. We're using it for data communication, and that implies a minimum rate of transmission below which, who cares? So that became the first breakthrough.

The second insight came from the work I had been doing to survive. At this point I was consulting for my former bosses at Ampex, who had formed a little consulting company after the layoffs, and I was learning a lot about serial data transmission. What I learned among other things was that for asynchronous data transmission, in the absence of an error condition, the signal goes back to a one level between every character. (If it doesn't do that it's a break condition, which means "The line may have broken, look out.") I set up a floating reference, such that the "one" condition would always re-set the reference, and whatever voltage that "one" wanted to be was fine for the circuit; it would tend to drift slowly down towards "zero,"

but then the data signal would come in and bat it back up to a "one." It maintained very good referencing, unless you went into a line break situation and after a few seconds your data would turn to zero — which was appropriate, because a break is a break. Rather than referencing the signal to the external standard of a potentiometer, I was tying it to the last known "one" condition that had come down the line, which was much more accurate and more flexible. My goal was to run it off a cassette tape recorder, because that was a popular and inexpensive storage standard of the day, but the potential variations in speed and pitch meant you couldn't do that with regular modems. You could do that with this modem. I developed this in 1973, and we used it for the Berkeley terminal to the 940, and after further development it appeared as the Pennywhistle modem — the commercial kit modem — in 1976.

KC: Okay, the setup of the XDS 940 at the center of this ad hoc network of modems and teletypes was what became Community Memory. Now — are we to the Hazeltine VDT yet?

LF: Just about. We opened the public terminal at Leopold's somewhere between the 8th and 14th of August, 1973, I think it was the 8th. I planted an article in the Berkeley Barb that week, describing it and giving a political rationale for that model of distributed information, which still holds true with me. But in January of 1974 we decided to move the terminal to the Whole Earth Access store on Shattuck Avenue, which was then in its first year of operation, basically as a catalogue store for hippies and for communes. And we leased a Hazeltine 1500 CRT terminal capable of operation at a princely 300 baud — actually, the terminal was capable of much higher baud rates, but the modem wasn't. We connected it to the Pennywhistle prototype, which I don't think was called that yet, it was just "the modem."

Almost immediately, something happened that raised pivotal issues of public access and control — I wasn't present when this happened, but I was told about it. We had a service contract [for the terminal] and the service technician was working on it, and dropped the circuit board for the keyboard. The chips were ceramic packs in two clamshell halves, sort of baked together. The top of one of the chips popped off and you could see the ac-

tual die in there with the leads bonded to wires, and somebody who was looking over the tech's shoulder asked, "Isn't that going to be a problem?" And he said, "Oh no, it works." That experience soured my belief in contract maintenance, and we began talking about various ways to make a system like this develop and grow and survive in a public access environment. Efrem was in favor of armor-ing the equipment, to keep everybody out.

CONVIVIALITY

I took the opposite tack. My father had recently sent me a book called *Tools for Conviviality*, by Ivan Illich and published by Harper & Row. Illich was a former Jesuit rising star who got off the official track and began writing books about de-schooling society, that was in 1970, and went on to establish some little center that he worked from in Cuernavaca, Mexico. He had a perspective that admitted technology and yet was very much outside the industrial model of society. He described radio as a "convivial," as opposed to an "industrial" technology, and proceeded to describe basically the way I had learned radio, but from the standpoint of its penetration into the jungles of Central America. Two years after the introduction of radio in Central America, some people knew how to fix it. These people had always been there. They hadn't always known how to fix a radio, but the technology itself was sufficiently inviting and accessible to them that it catalyzed their inherent tendencies to learn. In other words, if you tried to mess around with it, it didn't just burn out right away. The tube might overheat, but it would survive and give you some warning that you had done something wrong. The possible set of interactions, between the person who was trying to discover the secrets of the technology and the technology itself, was quite different from the standard industrial interactive model, which could be summed up as "If you do the wrong thing, this will break, and God help you." So radio could and did, in effect, survive in that environment because it "grew up" a cohort of people around it who knew how to maintain and sustain it. And this showed me the direction to go in. You could do the same thing with computers as far as I was concerned.

KC: Although possibly the computer technology of the period was less forgiving than tube radio technology.

LF: Certainly; remember, I had backed away from it as a kid because it was clear it had to work at every clock cycle. Such reliability was unheard of to me, in light of factors like noise in circuits.

STRATEGY SESSIONS

Well, I convened a discussion group around this whole concept, and announced it in Community Memory. Bob Marsh, whom I had known in college, turned up, and so did Ray Bruman — I think there were about five people in all. So we had several discussions about a computer appropriate for a public access environment, how it would be built, what it would be like; and my proposition, following Illich, was that a computer could only survive if it grew a computer club around itself. What is this computer like? How will it work? And we decided that the central aspect of the computer would be memory — random access memory chips were just then becoming available. Also, as part of the Resource One effort, I had heard from Don Lancaster who had written the September 1973 article in *Radio Electronics* about the TV Typewriter.

KC: Which he then expanded into the TV Typewriter Cookbook.

LF: Yes. In correspondence with him, and in a phone call from Resource One, I was trying to find out how to use TV Typewriters as terminals, because several people had come by and mentioned his article and said "You know, maybe you could use this."

KC: Okay, now. Just for my sake — in what way did this Hazeltine terminal not resemble a TV Typewriter, or was it simply too expensive?

LF: \$1500 or \$1600 was far too expensive. The promise of a TV Typewriter was that you could take a TV set and connect it to this little box, which you could build yourself with a couple of hundred dollars — not even a couple of hundred, a hundred dollars worth of parts — and you would have a usable terminal. Technically, the structures were the same. It was an excellent exercise for someone who wanted to learn digital electronics, not just digital but general electronics, and *Radio Electronics* got ten thousand paid responses for the two-dollar plans in the article whereas they might have been expecting twenty or thirty [responses]. So ten thousand people right away wanted to build it, and that was positive proof to me that we were watching the emergence of a convivial technology.

Having said that, I emphasize that the TV Typewriter was a very difficult thing to construct. There must have been quite a discrepancy between the number of people who ordered the plans and the number who actually built something. It used sequential memory, shift register memory, and was extremely tricky to build and debug because it used lots of little analog delay pulses. Bob Marsh had built one and he tried to expand it to improve its performance, and he never really got it working — but that's how he learned electronics! and that must have happened to quite a few people. Also, the TV Typewriter was not entirely usable as a terminal, because it was a paged display system; when you got to the end of one page, and that last character went up on the screen, you had only one-thirtieth of a second before the screen blanked and you saw the next [character]. So you were actually likely to be outrun by it.

KC: It didn't have any buffering for back scroll?

LF: No back scroll, that's right. I asked Don, "Why did you do it this way if it's supposed to be a computer terminal?" and he said, "It isn't supposed to be a computer terminal, people just want to put up characters on their TV sets." And he was right. So I had to think about why they wanted to, why "convivial technology" depended on doing something so simple. I concluded that people felt an urge to gain control over technologies that they knew were really important and were affecting, or would affect, their lives: computer technology, somewhat metaphorically in this case, and televi-

sion. The TV Typewriter combined the two, and it was perfect. Lancaster was a significant figure as a result, so far as I was concerned. He had created a synergy between convivial technology and computer technology *and* he had shown that there was a tremendous demand for this.

THE TOM SWIFT TERMINAL

So I was picking into the design and working out some details, and I called Lancaster back, and he said, "I've got a new version that I'm working on and it uses random access memory." Now that idea planted a seed with me and I said, "Aha, here's a route into the convivial computer." Computers need random access memory, and once you installed RAM in the computer, you could use the same memory to run this terminal. I believe that that was, in effect, the genesis of the architecture of the personal computer, and if any academics want to make their careers by arguing this point, they're welcome to see me and I'll help them. My idea was to begin with a memory card and add a card that was built to put data into the memory, either from a keyboard or from a UART, and a third card to get information out of the memory and put it to the screen. Connecting the three cards implied a bus, so I defined a 44-pin bus structure that used Vector connectors, which we could get quite cheaply, and which was good for DMA transfers. That three-card set constituted a terminal. I defined what the terminal would do when it got various control characters, like, what happens when it gets a carriage return? — well, it sets the character counter to zero, but the line counter doesn't change. When it gets a line feed, the line counter would increment, and all these would be pointing to the memory. I put together a specification that I called "The Tom Swift Terminal, or a Convivial Cybernetic Device," and I mimeographed it and sold it for 25 cents. That was in the fall of 1974.

But the common memory, and the bus structure, made this device potentially a lot more than a terminal. You could plug in a replacement front-end board, or maybe another display board, or an input editor board with more smarts and possibly even a microprocessor. Along the way, if you needed more memory, you could plug in more memory. Just as the boards plugged into the bus, I wanted to make sure you could plug the busses

together to expand them. I never really finished designing the whole specification, but the goal was that the builder, the user, would control the rate at which the device grew upwards into an intelligent terminal and then into a computer *per se*. You'd just keep plugging! And that would be the realization of the convivial ideal.

RESOURCE ONE DOWNS THE 940

In January of 1975, Resource One decided to shut down the 940. We had gotten permission to put [terminals] into the Co-op Markets, which would have meant a significant expansion, and we didn't trust that 940 to support a larger system, given that in certain ways it was marginal already. We weren't sure that the drum storage unit was going to survive, although apparently it did for many years after that. We weren't willing to risk the farm. We also didn't want to continue development along the dead-end line of QSPL, since what was that going to run on? Finally, I had worked on the [Data General] Nova and we were aware of the [DEC] PDP-11, both of which made it clear that minicomputers were the way to go. The Community Memory Project decided to split off from Resource One and go its own way, attempt to secure a minicomputer and move the software over to that, and come back with a system that was genuinely replicable. I think that, overall, it was a wise decision for the long term; for the short term, actually, it was not, but we weren't operating from a tactical perspective.

BEHOLD, IT STREAMS ACROSS THE FIRMAMENT

As you know, January 1975 also saw the publication of the Altair issue of the *Popular Electronics*. I was consulting and publishing out of my studio apartment, which was somehow full of a desk, a filing cabinet, a little work bench and a rack of shelves which stood over my heater grate. That was LGC Engineering — LGC stood for “Loving Grace Cybernetics,” which was inspired by Richard Brautigan's poem “All Watched Over By Machines of Loving Grace” and had been tested out a little bit, it was hanging up on the wall at Efrem's place. It was intended that the Community Memory project would call itself Loving Grace Cybernetics, and when it incorporated in 1977 I

believe it did use that name for a short while. We figured that LGC Engineering would be the future hardware arm of this.

THE FOUNDING OF PROC TECH

Meanwhile, through these meetings in the fall of '74, I got re-acquainted with Bob Marsh and he was looking for something to do. Apparently his in-laws were helping to support his family, so he wasn't in desperate straits, but he was unemployed and casting about for something to build. He convinced me to go in on a garage with him, at 2465 Fourth Street in Berkeley, which cost \$150 a month as I recall for 1,100 square feet. I set up my bench downstairs and he took this little loft office that was already in there, and put an air conditioner in it.

He and a friend of his, Gary Ingram, had access to a supply of center-cut walnut boards that were exactly eight inches high and about a quarter-inch thick — this had come from the Wisconsin woods through a third party. They were trying to decide on products to build with this walnut, and considering an electronic clock that would have a walnut case, trying to sell that through Bill Godbout or by mail order. That never went anywhere. They were also looking toward some improvement on the TV Typewriter, whether it was a redesign or just a kit to improve the existing one — that's when Don Lancaster told me “People just want to put characters up on TV.”

Right then the Altair article came out. Bob showed it to me and said “Look at this picture on the front. It's clear it's a phony. Look at that lump on the side, that's not real.” And the pictures of the boards that illustrated the article certainly didn't match what we could see on the front. So he said, “This thing has nothing in it, it's an empty box.” From what we could tell, there was nothing much more to an Altair than the microprocessor interfaced to a series of connectors. Bob wanted to start building peripherals for it, so he and Gary formed Processor Technology Corporation.

I remained a consultant and never became an employee or part of the corporation, but I took on contract work drawing schematics — I knew how to draw a schematic properly after doing some drafting in 1967 — and writing the assembly manuals. I wrote the caution that basically said “If you are not an experienced kit builder, find someone who is and get them to help you,” in effect warning that unless you apply convivial techniques to this you’re not going to make it. I sub-contracted Jude Milhon — who was Efrem’s girlfriend, and actually I met Efrem through her — to draw some little animated drawings of, say, a ceramic capacitor lifting a top hat and making a dance step, which demonstrated how you had to form the leads so that the capacitor would fit between the chips diagonally and still plug in.

Now, I did not do a lot of the engineering. On the 3P+S [Processor Technology port card] circuit, I figured out how to jumper the divider counters to get the various baud rates, although it’s not easy to explain how that worked. Other than that, I may have made very slight modifications to some of the designs. But I was doing enough then to pick up a certain amount of money.

KC: We must be about to the first meeting of the Homebrew Computer Club — if I remember it happening during the first week of March ‘75 in Gordon French’s garage.

LF: That’s correct — I think it was March fifth. I can’t recall whether Bob was doing any of this [Proc Tech] work prior to that, I don’t think so, but he and I kept talking about the Altair.

PEOPLE’S COMPUTER COMPANY

Meanwhile, ever since Resource One days, on Wednesday nights I had been going down to pot-luck dinners at the Community Computer Center in Menlo Park, that was run by Bob Albrecht.

KC: Was that the same thing as People’s Computer Company?

LF: Yes, I think sometimes they also called it People’s Computer Center, PCC. On Menalto Avenue, on the Menlo Park-Palo Alto border, they had set up a couple of minicomputers running

time-shared BASIC service to be a game parlor, in effect, for kids. What Bob Albrecht cared most about was kids and computers. He had written a book called *My Computer Understands Me When I Speak BASIC*, and some games for kids, and he published a newsletter that was also called *People’s Computer Company*. There’s a picture of Bob on the inside back cover of one of the earliest Whole Earth Catalogues, sitting teaching a bunch of kids to use a four-function calculator — which in those days was a big thing. He’s got this brush cut, looks like a porcupine. The connection to the Whole Earth Catalogue lent a certain prestige because we all thought of Stewart Brand as the guy who knew exactly where it was at. So Bob set up the Center, and a guy named Fred Moore sort of attached himself to it; he would sign people up when they came in the door and was building a list, and he wanted to organize some hardware classes, although he knew next to nothing about hardware. The people at the Center tolerated him.

Meanwhile, one night Gordon French went to visit the Kaylor Electric Vehicle Shop a few doors down, and happened in on this place; so he ended up on Fred’s list.

THE FIRST HOMEBREW MEETING

When the Altair was announced, Fred convinced Gordon to make his garage available, and then put out a call to the list. Thirty people responded — including Bob Marsh and me, because we drove down in a pick-up truck that I’d borrowed from a friend — and we all stood around looking at this unit.

That may have been the moment at which the personal computer became a convivial technology. You see, I had this particular Altair before the meeting; it was serial number 8 or something, it had been sent to People’s Computer Company as a review copy, and they gave it to me. I took it over to Efrem’s place and asked him what to make of it. Frankly, he considered it useless, and in a way he was right, since there was nothing to it but switches and lights. He kept it as a sculpture in his living room, on the same table with his guinea-pig cage, with its lights flashing to keep the guinea pigs company. I retrieved it and returned it to PCC, and it turned up as the centerpiece of the first Homebrew Computer Club meeting — by which

time, apparently, it had a sort of critical mass around it. Thirty people stood and looked at it and started to tell each other what they knew about it. Steve Dompier was there and reported on the trip he had made to Albuquerque to try to get his Altair kit from MITS.

KC: Right. As I recall, he had sent MITS a tremendous amount of money and said "Send me one of everything," and back in due course, or undue course, came a letter from MITS that said "We don't have one of everything."

LF: Exactly. I don't know whether he had to go down there to find that out, but he went down, and I don't think he was able to bring his kit back with him — he got it a little later. But he reported on what he had seen there, on what was in process, and he said it was a very small operation — which came as a surprise to a lot of people.

So the people in the room, including Steve Wozniak and Roger Melen, began to understand that maybe as a group we knew as much as these [MITS] guys, and that possibly the Altair wasn't even an action item. It changed our focus and we said, "First of all let's be in touch with each other, and sign this list, and tell each other what we're doing." The written record of that was reprinted and distributed to everybody as the first newsletter. I talked about Community Memory and the Tom Swift Terminal. Wozniak talked about the Break-out game he had done, and discussed the video terminal he was working on. Marty Spergel, who used to put together kits of parts for the *Popular Electronics* articles and sell them, actually gave away an 8008 chip at that meeting, which was a very nice gesture.

"WE WEREN'T NECESSARILY IMPRESSED"

Meanwhile, those of us who opened up the Altair box and looked inside weren't necessarily impressed. We discovered, for example, that out of four possible places for motherboards, only one was filled with a group of four sockets, and of the four sockets only two were occupied — one by the processor card, and one by the memory card. The memory card in turn had sockets for eight 256x4 static RAM chips, so the maximum capacity of the card was 1K bytes; but MITS only supplied two chips, so the stock Altair had 256 bytes of RAM,

which was not enough to support the kind of programming that most people wanted to do. The processor card was nothing more than the 8080 processor driving the lines on the bus through buffers. They had installed separate data-in and data-out buses, which was really inefficient, because you never did data transfers both ways at the same time. That told me that whoever designed this didn't know what he was doing — didn't really understand what a computer bus was.

As we worked with [the Altair] we found problems that only confirmed that impression. For example, the Phase 1 and Phase 2 clocks were sitting right next to each other on the bus. Transmission line effects therefore caused coupling between the two, and the last place you want coupling is on a clock line. The designer also used a positive-polarity Phase 2 clock — the critical one — and he should have used a negative-polarity clock, because TTL is much more effective on its pull-down than on its pull-up, and you have much more noise margin going from 5 volts down to 2.4 volts than you do when you're going from zero volts up to .8 volt. Many existing techniques could have been used to fix these problems, but none of them were, so it was a flaky design.

Also, the Altair needed everything. It needed I/O, it needed memory that worked. MITS offered a 1K dynamic memory card, not when the first machines were shipped but I think shortly thereafter, and they worked okay. Naturally, then as now, everybody wanted more memory and there was tremendous pent-up demand for a 4K board, so that people could run things like paper-tape BASIC. MITS began producing 4K dynamic memory boards from a different design, and they loaded all kinds of disk capacitors on them, but that's not all it takes to make it work. MITS never really got the 4K boards to work, but they sold them anyway. Almost all of them went out to customers and straight back to Albuquerque.

Bob Marsh didn't take long to figure out that he could build a board from the same 1K x 1 chips which I was using on the Tom Swift Terminal, [Intel] 2102's. Eight of the 1K-bit chips made 1K bytes, and then 4 ranks of those made 4K; and Bob did a good enough delay generator circuit, that board was produced as the 4KRA, I wrote the manuals, and Proc Tech was off and running.

At the same time the Homebrew Club — which wasn't called that yet — was growing amazingly. It met every two weeks. The first meeting was thirty people in Gordon French's garage; the second meeting was in the auditorium of the AI lab, John McCarthy's lab, at Stanford; and the third meeting was at the Peninsula School. At the third meeting there might have been a hundred and fifty people, although don't hold me to that, but I'm convinced there were more than a hundred, because they were going out into the halls to talk. Gordon French was trying to hold a lecture on computer science up front, and people were screaming that half the audience was outside — which I was too, because I wanted to meet these people. I already knew that somehow we had to impose some order on this process.

Now [Steve] Dompier in the meantime had gotten his kit and built it, and he set it up at the third meeting with a low frequency weather radio sitting on top of it. He plugged in the Altair and hand-entered the program [with the switches] and of course somebody kicked his plug out, and he had to toggle it in over again. We all wanted to know what was going to happen and he wouldn't tell us, he said, "Wait, wait, listen, wait."

Finally he pushed the Run button, and we heard music, and we could identify "Fool on the Hill" quite quickly. It was the strangest feeling we'd ever had, like "My God, where's the music coming from? What is this?" The radio was picking up the RF noise generated by the computer, and the tones could be modulated by programming the computer to do things at different rates. It was both absolutely brilliant and completely serendipitous, and more than anything it spoke to the quality of Dompier's intelligence, because it takes a particular kind of mind to catch this stuff. I'm not at all sure I ever would have figured it out.

When "Fool on the Hill" was over, the Altair began to play "Daisy." Now "Daisy" was the first song that had been played by a computer, I think in 1957 at Bell Labs. They accomplished basically the same thing, but with a much bigger computer and a speaker legitimately wired up to one of the bits. That tune was historically recognizable as the beginning of all computer music, which was why, in the movie *2001*, the computer HAL is being gradually lobotomized and regresses to singing

"Daisy." Dompier, I'm sure, had picked it up from there — as had most of us — and he had done precisely the right thing by invoking this to say that we claimed this history as our own, that we were as much pioneers as Bell had been, but that the computers singing for us were computers we could build and own ourselves. There was a tremendous message in that and we all responded. He got a standing ovation....

HP's EARLY COMPUTERS

Part Three:

THE STRONGEST CASTLE:

The Rise, Fall and Rise of the HP 3000

by Christopher Edler
cedler@boi.hp.com

"The HP 3000 was God's gift to computing"
- Hank Cureton, HP engineer¹

"Listen, lad: I built this kingdom up from nuthin'. When I started here, all of this was swamp! Other kings said it was daft to build a castle in a swamp, but I built it all the same, just to show 'em! It sank into the swamp. SO, I built a second one! That sank into the swamp. So I built a third one. That burned down, fell over, then sank into the swamp. But the fourth one stayed up. And that's what you're gonna get, lad: the strongest castle in these islands."

- Monty Python and the Holy Grail²

THE ALPHA PROJECT

The HP3000 minicomputer was the first computer developed from the ground up by Hewlett-Packard.³ It was conceived of in 1969, designed in 1970 and 1971, and first sold in November, 1972 – only to be withdrawn from market several months later, and not re-released until 1973.^{4 5} It has since become one of Hewlett-Packard's most successful products, and to this day contributes significant revenue to the company.

The history of the project is fascinating. This article will describe the early days of the HP3000, what the machine was to be originally, what it actually became, and what changed (and why) in the meantime. It is a narrative of a machine that faced daunting failures in the beginning yet, after re-release, found relatively quick acceptance in the business data processing market -- a field then unfamiliar to Hewlett-Packard.⁶

The HP3000 was, and still is, HP's premier business data processing machine. It runs a proprietary operating system, called MPE for *MultiProgramming Environment*⁷, and supports both batch and termi-

¹ The task of deciding a name for the operating system was rather arduous, and was the cause of many a bellicose staff meeting. Finally, the name

nal-based users. It can support these different users running different computer languages and applications. The 3000 has been very successful in industrial applications, competing against such industry giants as IBM in manufacturing shop floor applications, order entry, warehousing, and production control.⁷

The 3000 was a very important factor in Hewlett-Packard's rise to eminence among the world's computer firms. In terms of dollar volume, MPE and MPE-based systems have only very recently been surpassed in sales by the HP-UX line of Unix-compatible systems that conform to HP's current "open systems" concept.⁸

But to the historian, the HP3000 is significant not only for its long and lucrative success, but for the advanced computing principles it embodied at the outset. The HP3000 was:

- the first 16 bit machine to have a stack and virtual memory.
- the first minicomputer to implement COBOL.
- the first minicomputer to run a Database Management System.
- one of the first computers to be completely programmed in a high-level language.
- one of the first machines designed by hardware and software engineers.

...all for under \$100,000.^{9 10 11}

The HP3000 was, and is, undeniably a significant machine, but its achievements were hard-won. It was actually released twice, with an abortive introduction, quick recall, massive redesign, and reintroduction.

ALPHA IS BORN

After the HP21XX series of real-time computers, released in the latter half of the Sixties, the next-generation Hewlett-Packard machine project comprised the 32-bit machine called Omega and a smaller, 16-bit machine named Alpha. Omega was a truly ambitious machine, essentially realizing the

was chosen by an engineering manager as one that offended everyone equally. From an interview with Frank Hublou, *op. cit.*

concepts of DEC's VAX while pre-dating it by almost 10 years. Unfortunately, it was a bit too ambitious for such a conservative company as Hewlett-Packard, and was cancelled in 1970, after nearly 2 years of work.

The Alpha project was supposed to be simultaneous with Omega, but there were actually a few engineers at most working on Alpha at any given time during the Omega development.¹² In fact, the Alpha 1's External Reference Specifications, included as Figure 2, show a release date of July 20, 1970, less than four months before the Omega was canceled.¹

Tom Perkins, General Manager of HP's Data Products Division and star of the HP21XX line, was the man who cancelled the Omega project. He was promoted in the Fall of 1970 to head Corporate Development, after which George Newman was made "GM," as HP-ers called their divisional manager.¹³

Figure 3 shows the original Project Data Sheet for the Alpha hardware system. Highlights from the "Objectives" section are:

1. a 16-bit computer system for multiprogramming and real-time application,
2. an optimum architecture for efficient execution of high level languages,
3. an I/O structure that provides CPU independent multiplexed data transfers of at least 1 megaword/sec,
4. a fast multiple level interrupt structure,
5. a modular hardware structure communicating over a time-shared buss (sic)¹⁴

Note the estimated costs for the Alpha hardware of \$896,000 in mid-1970. Figure 4 is a preliminary timeline that accompanied the Project Data Sheet, showing an MR (manufacturing release) date of February 1, 1972.¹⁵

As mentioned, there was great dismay over the cancellation of the Omega machine, and the decision to go ahead with the Alpha was not received enthusiastically. Going back to "just another 16-bit machine"¹⁶ had its discouraging aspects even beyond preoccupation with the word-size; the engineers at HP felt that they had been designing and working with a machine similar to Alpha (HP21XX) for a number of years. The prospect of going "back" to

the Alpha might also have seemed pedestrian in contrast to eagerly anticipated work on the 32-bit Omega. For whatever reasons, Omega was sorely and widely missed.

But the Alpha, and subsequently the HP3000, were to become Omega's memorial — intentionally or not — thanks to one significant decision. The developers, consciously or unconsciously, decided that the Alpha would be as much like the Omega as a 16-bit machine possibly could. As one engineer of that time put it, "We had an Omega mentality in an Alpha box;"¹⁷ in practical terms this meant that almost everything about the Omega was "cut in half" but survived into the Alpha plan. Both were to be virtual memory, stack-based machines with three modes of operation — timeshare, batch and real-time. The major differences were in the "virtualness" of Alpha memory (only application code was "virtual", user data was limited to 64K), and in the input/output scheme — a particular disappointment for the I/O engineers, who had designed elaborate input/output hardware systems¹⁸ for Omega's 32 bit data path, which could never be incorporated into the Alpha.¹⁹

Upper management saw the initial Alpha specifications and had some concerns over the functionality and feasibility of the machine. Arndt (Arnie) Bergh, one of the prime hardware designers for the project, recalls that, "...he [the general manager and one of the people who ordered the cancellation of the Omega] said 'You did it to me again. All I wanted was a simple, clean, inexpensive instrumentation computer that would do the job.'"²⁰

In spite of reservations, the project was approved and specifications for the Alpha were completed. By the time the External Reference Specifications were released in July 1970, the vision for the Alpha was complete in broad outline. It would have a 16-bit data word with a single storage register. It was to be stack-based with a maximum memory of 128 Kbytes, or 64K words, of ferrite-core memory.²¹ (A common confusion of the period devolved on the exact meaning of "XX K of memory". Most mini

¹ The original Figure 1 has been reproduced as the cover.




Figure 2 Alpha External Reference Specs²²

PROJECT DATA SHEET

1. **PROJECT:** (MODEL NUMBER, NAME, ETC.)
 HP ALPHA Processor (hardware)

DIVISION:
LAB NO.: Cupertino
DATE: 6/22/70

2. **PROJECT PERSONNEL:**

Bergh, Forbes, Katzman, McAninch, Olkowski and staff, Reid and staff.

3. **OBJECTIVES:** (SPECIFICATIONS, FEATURES, RELATION TO OTHER INSTRUMENTS, PACKAGING, AND NOTES ON DEVELOPMENT PROBLEMS, NEEDS, ETC.)

1. A 16 bit computer system for multiprogramming and real time application.
2. An optimum architecture for efficient execution of high level languages.
3. An I/O structure that provides CPU independent multiplexed data transfers of at least 1 Mega word/ sec.
4. A fast multiple level interrupt structure.
5. A modular hardware structure communicating over a time shared buss.
6. A modular asynchronous submicrosecond memory system expandable from 4K to 65K.
7. A class B environmental spec.
8. No options - Memory protect, arithmetic capabilities are standard.

Specification - Reference the ERS from the investigation project for the HP ALPHA.

4. **PROFIT/COST DATA:** (EACH 6 MONTHS OR WITH CHANGE IN OBJECTIVE, PRICE, ETC.)

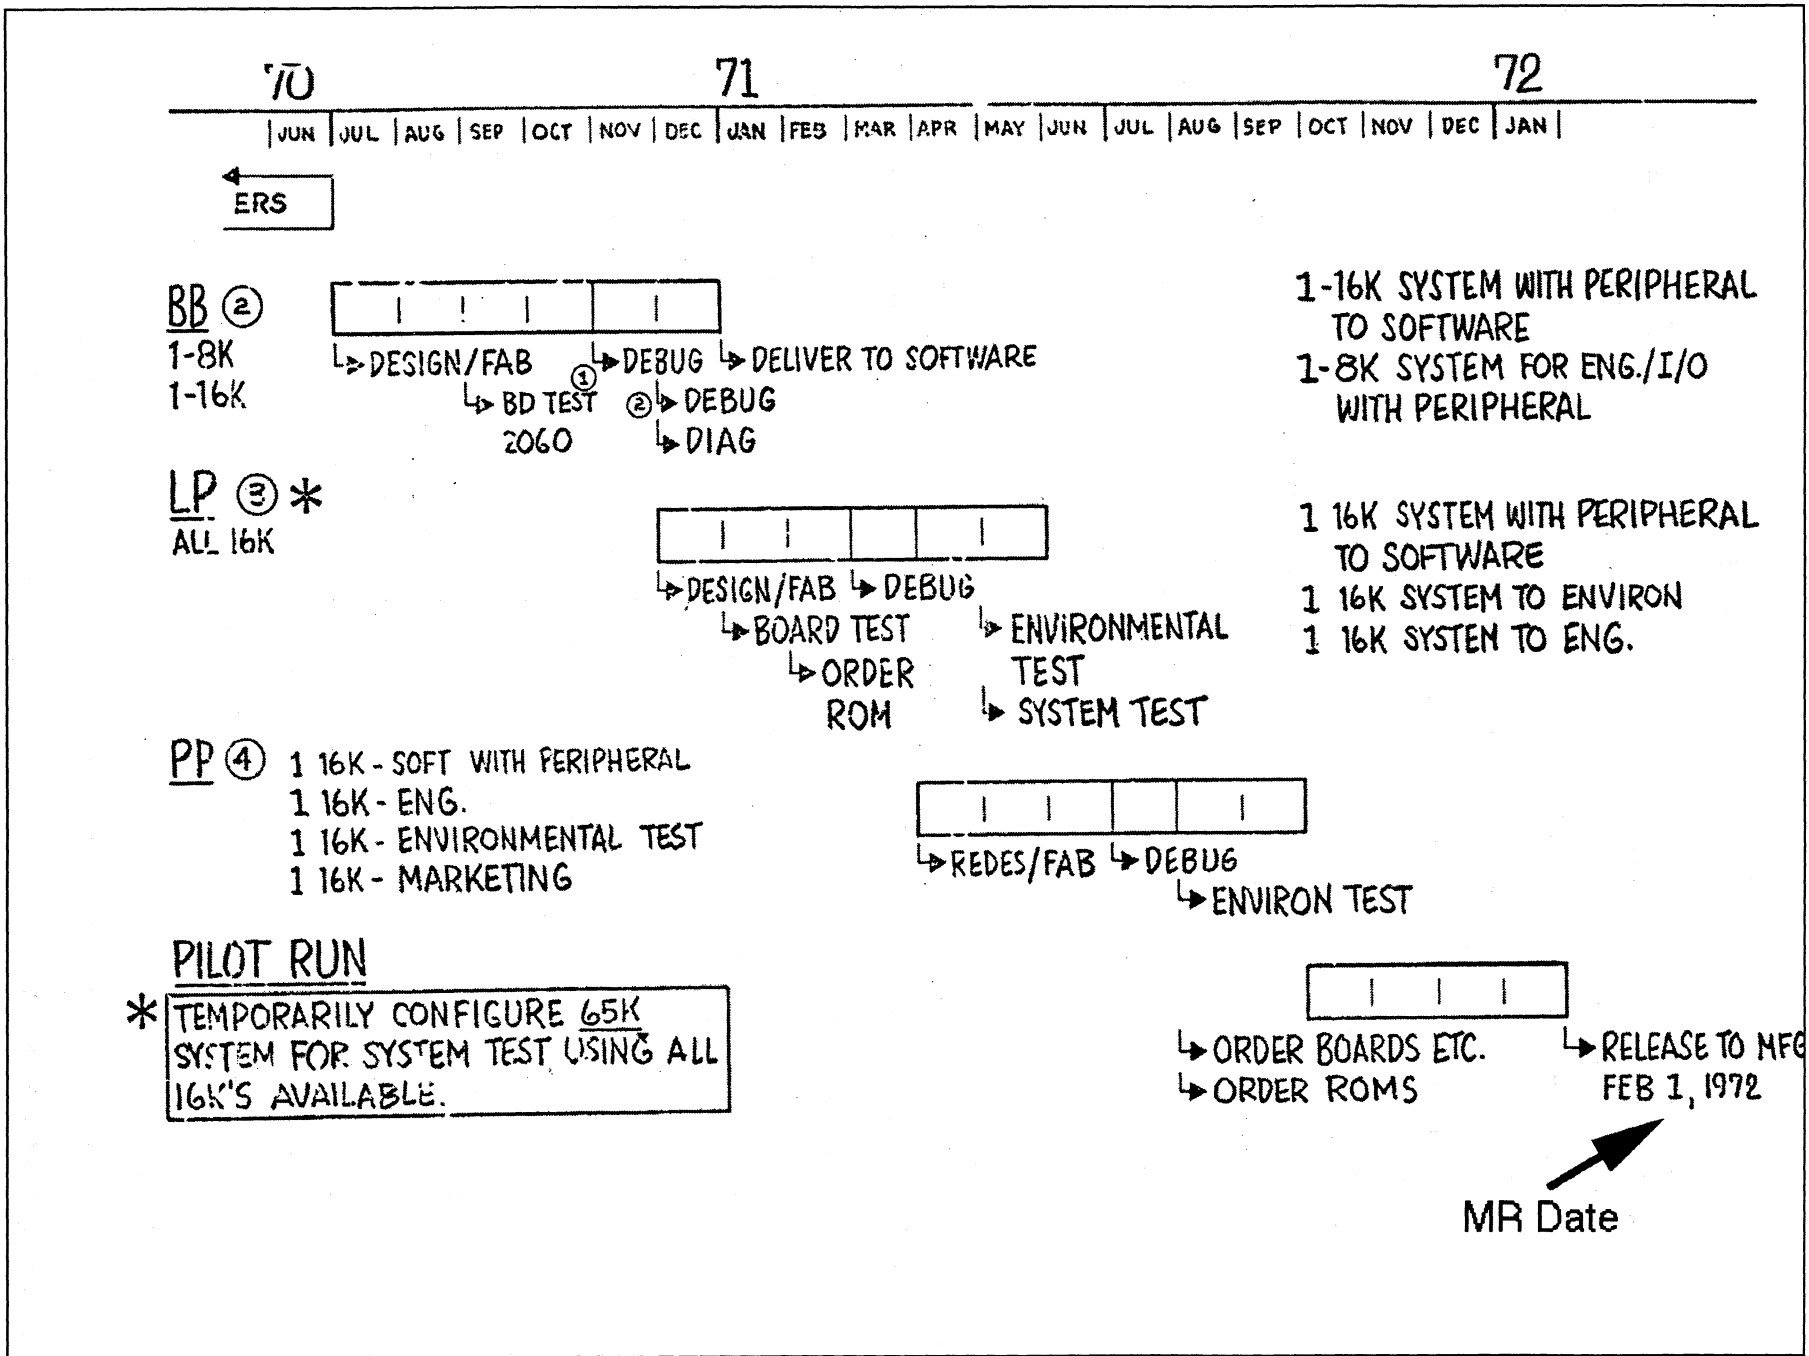
	ORIGINAL	FINAL				
DATE						
ENGINEERING \$ INVESTED	15K					
ESTIMATED ENGINEERING \$ TO GO	595					
DIRECT PROJECT CHARGES						
MODEL SHOP	8K					
TOOLING	73K					
NET PILOT RUN	10K					
MATERIAL	95K					
OTHER (3)	100K					
TOTAL PROJECT COST ESTIMATE (1)	896K					
TARGET PRICE	(2)					
CURRENT ESTIMATE OF QUANTITY/MONTH						
DOLLARS PROFIT AT % FOR YEARS						
PROFIT \$ GAIN OR LOSS ON MODEL NO.						
NET PROFIT EXPECTED						
RETURN FACTOR (PROFIT/COST)						

COMPONENT DEVELOPMENT, SPECIAL TOOLING, CONTRACTED ENGINEERING, ETC.

5. **NOTATIONS:** (ENTER DATE AND NOTATION AS NECESSARY TO EXPLAIN MAJOR CHANGE IN PROGRESS GRAPH, OBJECTIVE, PRICE, SPEC., ETC. CONTINUE ON ADDITIONAL PAGE.)

750041

Figure 3 Alpha Project Data Sheet



computer memories were measured in 16-bit data words and generally described as nK (thousand)-data-word memories. IBM, however, introduced the concept of 8-bit words called "bytes", which became an alternate general usage that clouded the issue.) The architecture as initially defined allowed for multiple CPU's, although this feature was never realized in the finished product.²³

The system was designed for modularity, with communication between modules over an asynchronous priority demand data bus. It utilized an "omnibus" concept in which the devices all shared the same bus. The maximum number of modules was seven.

THE OPERATING SYSTEM

A key distinction between the Alpha and other contemporary minis was to be its operating system — a feature then generally found only in mainframes. Minis were ordinarily bought with and for a single application, or were intended to be programmed in machine code by the user. HP envisioned Alpha as the first "general-purpose" or flexible-purpose minicomputer, and an operating system was essential to the concept. Alpha's was to be called MPE, for MultiProgramming Executive.

MPE was actually three operating systems in one. First and foremost, it was to be a timesharing operating system. This was to allow "multiprogramming multiaccess — in particular, time sharing with good term response".²⁴ The Alpha was intended primarily as a timesharing machine, and its architecture was chosen accordingly.²⁵

Secondly, it was to be a real-time system. The designers described the operating environment as commercial real-time and process control systems with fast interrupt response time.²⁶ Alpha in this context was intended to replace 21XX-series computers in real-time operation.²⁷ This was a key issue for HP, which considered real-time operating systems to be a "core competency" for the company, and a significant source of revenue.²⁸

The third capability of the operating system was a batch mode, in which jobs could be scheduled and run just as if from a user terminal. Proportioning among these three resources could be fine-tuned

through MPE, in essence turning off modes not currently needed.

How was this to be accomplished? The External Reference Specifications sketched some of the ways:

- "1. Program sharing of re-entrant code by two or more users simultaneously (e.g. compilers, intrinsics, etc.).
2. Automatic resource allocation and overlay of infrequently used memory.
3. Protection between users, and between users and the system.
4. Modular I/O organization to maximize effective device usage.
5. Users are segmented such that total residence is not required to execute a program.
6. Fast interrupt response time and environment switching, plus nesting of high priority interrupts."²⁹

LANGUAGE

The primary programming language for the Alpha was to be called ASPL (later SPL), for Alpha Systems Programming Language. It was an ALGOL derivative very similar to that of the Burroughs B5500, and — as ALGOL was the first language to handle recursion³⁰ — was an excellent language for stack-based machines. Alpha was designed from the ground up to accept SPL; in fact, MPE itself was written in SPL, while most other manufacturers wrote their operating systems in assembler.

FILESYSTEM AND I/O

The Alpha file system featured named files and directories with controlled access. Similar to modern-day UNIX machines, the Alpha relied on device-independent file access, and made I/O devices accessible as labeled files.

Input and output processing could occur independently of the CPU through an I/O processor (IOP) scheme. A master device, the MCU, controlled assignment of time slices of the bus to modules. This concept was unique enough to earn a United States patent for the machine; the first page is shown in Figure 5.³¹

United States Patent [19]

[11] **3,820,079**

Bergh et al.

[45] **June 25, 1974**

[54] **BUS ORIENTED, MODULAR, MULTIPROCESSING COMPUTER**

[75] Inventors: **Arndt B. Bergh**, Los Altos Hills; **Bert E. Forbes**, Palo Alto; **James O. Hamilton, III**, Sunnyvale; **Joseph C. Mixsell, Jr.**, Palo Alto, all of Calif.

3,470,542 9/1969 Trantabella..... 340/172.5 X
 3,480,914 11/1969 Schlaeppli..... 340/172.5
 3,560,934 2/1971 Ernst et al. 340/172.5
 3,623,011 11/1971 Baynard, Jr. et al. 340/172.5
 3,633,169 1/1972 Bickford 340/172.5
 3,731,283 5/1973 Carlson et al. 340/172.5

[73] Assignee: **Hewlett-Packard Company**, Palo Alto, Calif.

Primary Examiner—Raulfe B. Zache
Attorney, Agent, or Firm—A. C. Smith

[22] Filed: **Nov. 20, 1972**

[57] **ABSTRACT**

[21] Appl. No.: **316,429**

Related U.S. Application Data

[63] Continuation of Ser. No. 194,764, Nov. 1, 1971, abandoned.

A multiprocessing computer is structured in modular form around a common control and data bus. Control functions for the various modules are distributed among the modules to facilitate system flexibility. Modules separate from the central processor handle input/output operations to free the central processor for data manipulation. The central processor includes circuitry for instruction and data pipelining, single, double and triple shifts, preadding and memory mapping and interleaving. The central processor also includes a read only memory look-up table for micro-programming instructions.

[52] U.S. Cl. **340/172.5**
 [51] Int. Cl. **G06F 15/16**
 [58] Field of Search **340/172.5**

[56] **References Cited**

UNITED STATES PATENTS

3,295,102 12/1966 Neilson 340/172.5 X
 3,445,822 5/1969 Driscoll..... 340/172.5

19 Claims, 111 Drawing Figures

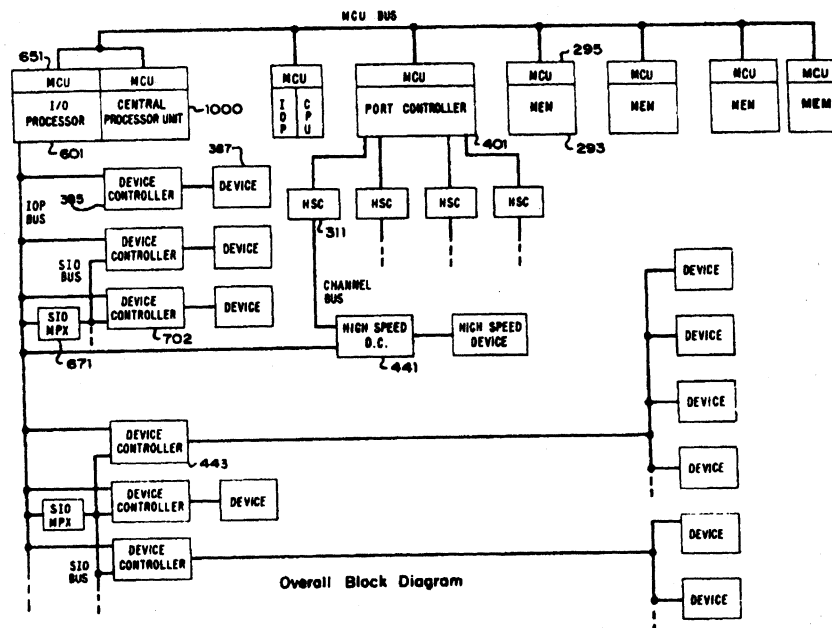


Figure 5 Original HP3000 Patent

The Alpha had 170 instructions (opcodes). For comparison, the IBM System/360 had 142 opcodes. Each code was a 16-bit word, but stack instructions were packed two per word. Instructions were stored on a ROM (read only memory) chip 32 bits wide, so that two instructions were stored per ROM address.^{32 33 34}

The instructions were divided into functional areas as follows:

- Memory reference
- Stack & control
- Shifts, bit tests, conditional branches
- Immediate
- Linkage and control instructions
- Special opcodes
- Move opcodes
- Mini opcodes³⁵

The machine was initially specified to give satisfactory response time for 16 users when equipped with 16k of memory, or 32 users with 32k of RAM. If the system was only running BASIC, the number of users could be increased by 50%.³⁶

The schedule for completion, as quoted in the preliminary document of July 1970, was:

- Approval 9/8/70
- Preliminary ERS 11/1/70
- Final ERS 4/1/71
- Begin system integration 5/1/71
- System integration 11/1/71
- Release 12/31/71
- IMS complete 3/1/72³⁷

The projected price of the machine was \$100,000. It was to be called the "HP System/3000".³⁸

THE SYSTEM IS ANNOUNCED

The premier showcases of the computer industry, where companies would announce and demonstrate their new technology, were the Joint Computer Conferences held in the spring and fall of each year. The industry and the press followed these conferences closely.

Hewlett-Packard announced its "System/3000" at the 1971 Fall Joint Computer Conference in Anaheim, CA. Although most Hewlett-Packard users and engineers remembered it vividly, the announcement only merited a small blurb on page 30 of the November 11, 1971 issue of *Computer-World*, with the headline "HP adds time-sharing system/3000".³⁹ The article stated that the 3000 had available RAM memory of 32 to 131K, a maximum cycle time of 950 nanoseconds (pretty good in those days), an instruction set of 170 commands, and could "be used as a front end to an IBM 360".⁴⁰ It was to have BASIC, FORTRAN, SPL, as well as COBOL "before the end of the year" (which didn't become available until 1973). The release was to be in August, 1972.⁴¹ Elsewhere, it was advertised to run up to 16 terminal users on a fully loaded system.⁴²

A CLOUD ON THE HORIZON

In early 1972 things started to go wrong at the HP3000 lab in Cupertino. The hardware was up and running in the form of three prototypes, PP1 (Production Prototype 1), PP2, and PR1 (Pilot Run 1).⁴³ The software effort was slowing down.

The first evidence of this can be found in office documentation in early 1972, when the extent of the 3000's functionality was being reconsidered in an effort to protect the late 1972 deadline. The real-time aspect of the operating system, specifically, would later cease to be included (or even mentioned) in internal documentation.

Figure 6 shows a memo from Ron Matsumoto, a manager in the O/S section of the HP3000 lab, dated February 1, 1972. Notice that on the 5/15/72 and 6/5/72 lines, real-time has been "stealthed out". This will be one of the last times that real-time is mentioned in a project scheduling memo. From now on, emphasis in Alpha development would be entirely on timeshare and batch functions.

MEANWHILE, OFF IN A CORNER...

Before we continue, Gentle Reader, allow me a brief aside concerning a project worked on in 1971 by four programmers in a secluded corner of the "Applications Lab", a section of the HP3000 lab devoted to non-operating-system applications such as the system editor, and to computer-aided instruction software.⁴⁴

File management software was being worked on in tandem with the rest of MPE, and the HP3000 designers sought to include a file inquiry tool (called QUERY) to enable searching. Two of the programmers involved were Dick MacIntire and Lee Bollinger.⁴⁵ Eventually the designers felt that they could expand their utility into a data management application that could be used by not only the operating system, but by all software languages on the HP3000. The key was a very new concept called database management systems, or DBMS — which may seem commonplace today, but were still in their infancy in the early 1970s.

DATABASES AND DBMS'S IN THE 60'S AND EARLY 70'S

Early DBMS concepts can be traced back to the late 1950's, in academic papers that discussed "generalized routines", which could sort any file regardless of its data content.⁴⁶

Early efforts at a DBMS included programs by GE, by SHARE, and by IBM itself — notably the first release of RPG for the Model 1401 in 1961 — and RCA's Retrieval Command-Oriented Language, developed in 1962⁴⁷; but DBMS theory languished through the mid-60s for lack of a perceived need. One of the first commercially available database systems was written by IBM, initially for the Apollo program, and publicly released for the System/360 in 1969, under the name IMS.⁴⁸

CODASYL, a group mandated to create standards for COBOL, meanwhile formed a List Processing Task Force (later called the Data Base Task Group) to construct COBOL database extensions. They released reports in 1969 and 1971 with recommendations for strategic "building blocks" of a database management system, such as a data description language (DDL)⁴⁹ and use of what was called the networked model of databases.⁵⁰

In 1970, a third-party software firm called Cullinane released their database product, IDMS, for IBM 360 and 370, which was extremely successful.⁵¹ Other key products of the period included DMS 1100 for the UNIVAC 1108, and DEC's DBMS-10, which ran on the PDP-10.⁵²

IMAGE

At HP, the file management group and the inquiry group merged in September 1971 and decided to work on an integrated data management package. This package, named IMAGE most probably by Dick MacIntire, would allow access to data files by any process on the 3000 by the use of intrinsic commands.

The developers looked at a number of database systems then available. The main data center at HP corporate headquarters had a mainframe running the TOTAL database management system. The team also read and investigated the work of Leo J. Cohen, an early expert in database management concepts⁵³, as well as the CODASYL report released in 1971. Image followed CODASYL recommendations in being a network-based database, but the product itself was not in full compliance with CODASYL standards.

Scheduled for a second release of the system⁵⁴, Image was not a high priority item at the time of the first release. Developers on the project had very little clout when competing for valuable and rare computing time on the HP3000 prototypes. Most development was done on either an HP2100 emulating a 3000 or (primarily during weekends) on the third HP3000 prototype, the one frequently cannibalized to keep the other two running.

Development took place throughout 1972 and early 1973, with quality assurance testing beginning in mid-1973. The product could have been released in late 1973, but a late hire of a technical writer postponed the introduction until the 1974 rollout of the HP3000 "CX" series.

...now, back to the HP3000 lab in early 1972...

COMPUTER SUDDEN INFANT DEATH

"...we fired it up, and it was just quite slow..." Engineer Arnie Bergh on the first HP3000⁵⁵

BACKGROUND : THE HP3000 LAB ENVIRONMENT

A tension always prevails between the people who develop a product and those who sell that product. This tension can be good, in that — in the words of Ed McCracken, then Marketing manager of the HP3000 and now CEO of Silicon Graphics Corporation — "you generally have a marketing person and a manufacturing person who provide some balance against the dreams of an engineering manager."⁵⁶ But in less fortunate circumstances, the tension can be so strong as to prove catastrophic.

There were between 60 and 85 people in the HP3000 lab during the machine's development. This staff were divided into 6 "sections", each devoted to a part of the project, such as the operating system, user applications, or quality assurance. General management was the responsibility of Steve Vallender. Figures 7 and 8 are organization charts of each section in the HP3000 lab in 1972.

The HP3000 lab was part of a much larger R & D organization devoted to the many products of the Data Systems Division, which included the HP21XX series of real-time computers.

Marketing was arranged somewhat differently. It was grouped by sector (Educational/Medical/Government, Industry), and comprised smaller product marketing and marketing support groups. Figure 9 shows an organizational chart of the Data Systems Division marketing organization in 1972.

There was little love lost between the HP3000 lab and marketing; or, to put it a better way, there didn't seem to be a lot of mutual respect between the groups — especially with regard to the "owners" of the project, those who would determine the functionality of the product. One of the original engineers put it this way:

"They [marketing] wanted to drive the project from marketing...if they said it could do something, they expected the lab to make it happen..."⁵⁷

This tension was so evident and strong that (in the words of Hank Cureton, an HP3000 engineer at that time,) "People from marketing were banned

from the lab".⁵⁸ The turf battle went even beyond that; engineers were not supposed to discuss the product or its progress with anyone in the marketing division. Lab manager Vallender "discouraged us from talking to marketing people", according to Jim Holl, Quality Assurance manager for the HP3000 software.⁵⁹

Denied access to development information, marketing was unaware of the progress of the product, and of its true specifications. They naively passed on whatever information they could glean to their sales representatives, who presented it to customers. Data sheets were created before the development was complete. Published performance data were, in general, more optimistic than what the lab could measure on running prototypes⁶⁰ — and this over-optimism was anything *but* naive. As Ed McCracken put it, "there was a lot of lying going on."⁶¹ Someone in the lab was even overheard knowingly giving out "rosy" specifications to customers on the phone, hanging up, and saying, "Ah, I've just got to tell them something..."⁶²

The tight lid on information about the 3000's progress provoked some speculation by members of the team. Some engineers suspected that the project was in danger of getting dropped. In their view, setbacks could have pushed the HP3000 off the drawing board — and there had been significant setbacks already.⁶³ In light of the antagonism between lab and marketing, failures on the development end could have persuaded top management to give overall project direction to marketing. Vallender used a lot of energy defending against this outcome; he "knew how poorly we were doing and he didn't want pressure from marketing."⁶⁴

THE FOLLOWING ARE WORKING ON THE 3000			
			TOTAL: 66
22-8100		22-8200	
Harlan Andrews	Ken Rothemueller	Bill Ansley	Carol Fuquay
Larry Birenbaum		Fred Athearn	Con Koreneff
Tom Blease		Jon Bale	Phil Taylor
Terry Branthwaite		Lee Bollinger	Doug Green
Steve Brown		Yann Cordelle	Bob Brown
Jean-Michel		Sam Edwards	
Alan Hewer		Hugh Field	
Jack MacDonald		Gary Gapp	
Dennis McEvoy		Bill Holden	
Bob Miyakusu		Ken Klingman	
Ron Matsumoto		Bob Mayer	
		Dick McIntire	
		Micki Miller	
		Scott Wallace	
		Fred White	
		Lee Johnson	
22-8300		22-8400	
Jerry Bausek	John Couch	Hank Davenport	
Waldy Haccou	Joel Bartlett	Tom Ellestad	
Terry Hamm	Chris Larson	Dan Gibbons	
Doug Jeung		Pete Graziano	
Dave Johnson		Tony Hunt	
Steve Ng		Walt Wolff	
Terry Opdendyk		Bob Bellizzi	
Paul Rosenfeld			
John Shipman			
Jerry Smith			
John Welsch			
John Yu			
Bill Foster		Mike Green	
22-8530			
Norm Alexander	Dan Low		
Don Coleman			
Ken Mintz			
Claude Robinson			
Jim Russell			
Earl Stutes			
Gary Wermuth			
Jim Holl			

Figure 7 R & D Staff List 1

<u>VALLENDER, STEVE</u>	22-8000	<i>was manager</i>	
			Werner, Gwen
Green, Mike	22-8000		
Matsumoto, Ron	22-8100		
Johnson, Lee	22-8200		
Foster, Bill	22-8300		
Bellizzi, Bob	22-8400		
Kintz, Joe	22-8500		
Taraldson, Jim	22-8010		
Quenneville, Paulette			
Rampone, Jody			
Witt, Linda			

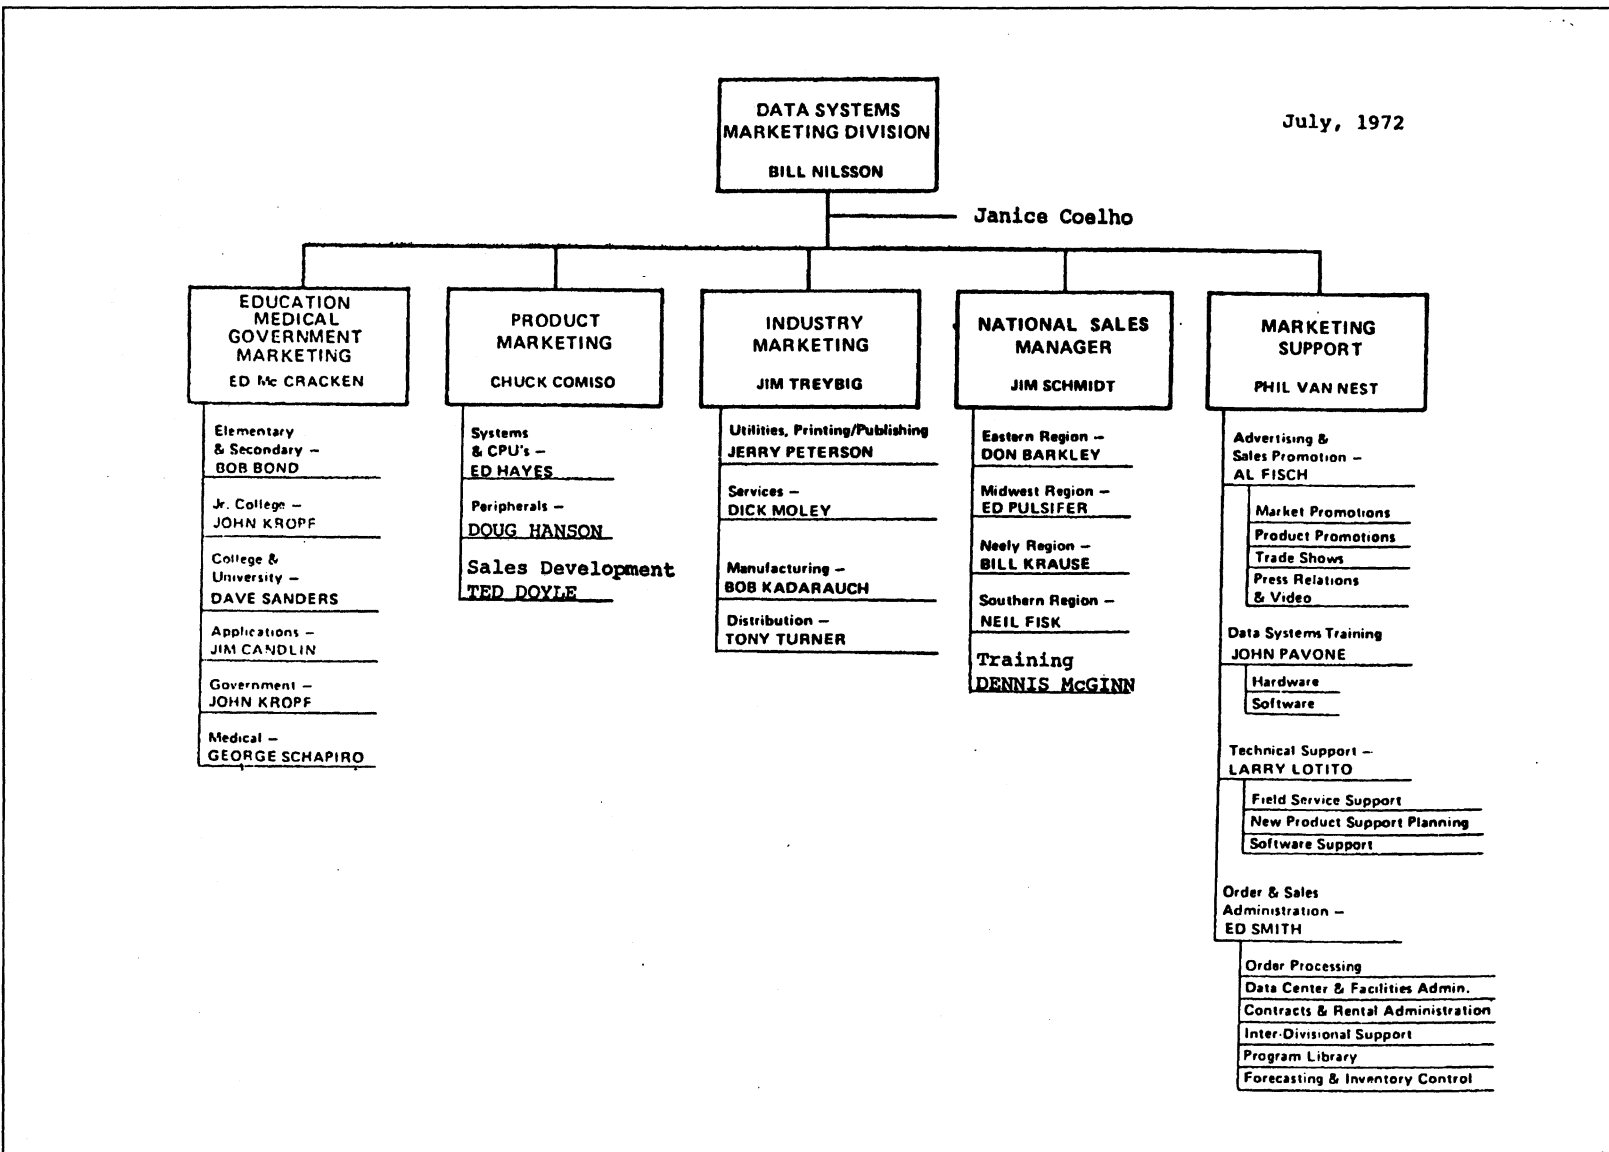
TOTAL: 12 (including Steve)

Total: 85 for locations - 22-8000, 22-8010, 22-8100, 22-8200, 22-8300,
22-8400, 22-8500, 22-8510, 22-8520, 22-8530

Revised: January 24, 1972

Figure 8 R & D Staff List 2

Figure 9 Marketing Org Chart



SPRING 1972: SYSTEMS MANAGEMENT GROUP FORMED

Top management was indeed concerned. In May of 1972, Dick Hackborn (who would later become Vice President of HP) hired Dave Crockett to head up a marketing team called the Systems Management Group, which would be part of the lab, but at the same time empowered to "product market-ize" the HP3000.⁶⁵

At this point the lab began a serious effort to determine what their new product could do. It is interesting to note that not until 1972, four years after the Omega/Alpha project began, did the lab hire a person — Jim Peachy of the Systems Management group — to analyze the performance of the new machine.

Peachy had done performance testing and tuning for one of the world's first timesharing systems at Dartmouth University in the mid-60s and had worked at Memorex and at General Electric. Three days after HP hired him, he announced "This isn't even a timesharing system."⁶⁶ After testing out the instructions, and timing out the loops, he stated categorically that there was "absolutely no way" for the HP3000 as it existed to meet its timesharing targets. Marketing meanwhile was assuring HP's customers that, at minimum, 16 users could run on a small HP3000, and 32 to 64 users could run on a 64K memory version of the machine.⁶⁷ The Systems Management Group stuck by its calculations and insisted that "this thing [the HP3000] won't time-share more than three users..."⁶⁸

This was not an occasion for great joy. Some of the engineers had scant regard for these "polished marketing types from Memorex"⁶⁹. Unfortunately, it was by then the late summer of 1972, and a full "fix" of the product by the release date was out of the question. Probably at this point, the August 1972 release date that had been announced at the 1971 Fall Joint Computer Conference was pushed to November of 1972.

SEPTEMBER-NOVEMBER 1972: DELAYS

The release date, although pushed, was looming by September. Management attempted a full definition of the code set to be distributed as the first release of

the MPE operating system. The new functionality contrasts materially with what was originally presented in the earlier memo. The new memo, dated September 25, 1972, is included as Figure 10 and 11.

This new memo proposes four major versions of MPE, and details what will be included in each. It mentions that no real-time capability and no spooling (for batch) functionality will be included in the first release. Console operator capability was proposed here for MPE 1, but would later be deferred to MPE 2.

The schedule for the other MPE releases was as follows:

February 1, 1973 MPE 2: spooling

April 1, 1973 MPE 3: real-time

June 1, 1973 MPE 4: system recovery, performance monitoring

This information was to be released to the sales force, and lab members were to go to the regional sales offices and tell them personally.

To say that the salesmen were dismayed would be an understatement. Many of them had already collected commissions from HP3000 purchases.⁷⁰ HP was forecasting that 120 units would be sold the first year, and a \$100,000-plus HP3000 carried a nice sized commission.⁷¹

NOVEMBER 1972: "NOVEMBER IS A HAPPENING"

No one knows who first spoke this phrase, but in the weeks prior to the release of the first HP3000, dozens of the posters shown in Figure 12 were placed all over the factory floor.

The picture shows an HP3000 going out the shipping dock door, ostensibly to its first customer, the Lawrence Hall of Science at UC Berkeley. Unfortunately, as Frank Hublou put it, "they should have put it on the truck, drove it around the block, and brought the machine back. That's what should have happened."⁷²

M.P.E. Release 1 - November 1, 1972**Capabilities:**

- . Full timesharing and concurrent batch modes of operation with accounting and spooling for later release.
- . Full file system capabilities with private volumes, tape label processing and file security scheduled for later release.
- . System generation capability from terminal. Command and update modes scheduled for later release.
- . Full terminal type support via hardwired and 103 type modems.
- . Full organizational management capabilities with accounting scheduled for later release.
- . Console operator and system supervisor capability.
- . Special capabilities including:
 - Extra data segment, process handling, multiple RIN and privileged mode. Core residency and real time capability scheduled for later release.
- . Libraries and library handling capability.
- . Languages include SPL, BASIC and FORTRAN.
- . Other capabilities are STAR and Text Editor.
- . Off-line and on-line diagnostics.

(continued)

P 3000 SOFTWARE RELEASEM.P.E. Release 2 - February 1, 1973

Capabilities:

- . Complete file security.
- . Spooling for more efficient batch operation.
- . 202 Modem support.
- . Complete accounting and logging facilities.
- . Increased reconfiguration capability with card reader configuration facility and update mode.
- . Operator and system manager capabilities.

M.P.E. Release 3 - April 1, 1973

Capabilities:

- . Real time capability.
- . Private volumes by group.
- . Tape label processing.
- . Core residency capability for user code.
- . Text Formatter.
- . Capability to switch job input streams via command.

M.P.E. Release 4 - June 1, 1973

Capabilities:

- . On-line performance monitoring via monitor command.
- . System recovery from error.
- . Power failure recovery.



Figure 12 "November is a Happening" Poster

It didn't. HP installed the machine, turned it on, and discovered — along with the customer, of course — that the 3000 could only accommodate one or two users before falling to its knees. It was immediately returned.⁷³

After the Lawrence Hall of Science debacle, many engineers quipped that they weren't sure if the "November is a Happening" poster showed a man sending the machine out or bringing it back in.⁷⁴ Feelings were prevalent throughout the lab that the machine was released much too early, although the engineers knew it could not meet specifications. The November release had been perceived as a firm deadline, which they had to honor regardless.⁷⁵

Perhaps there was an unofficial "1 year after announcement law" in effect which had something to do with the November release. For example, the IBM System/360 was announced in 1963, and the first version was released (arguably in poor condition) in 1964; IBM's System/370 was announced in 1970, and released a year later.⁷⁶

DECEMBER 1972 -FEBRUARY 1973: FURTHER DELAYS

Things didn't look that much better a month later, although a prototype used for training in December was able to run reasonably well with 4 users, rather than the earlier 2. During the two-week training session with customers, the prototype crashed about once every two hours, and the operating system "was patched more than 50 times."⁷⁷

December 11 brought another schedule change, detailed in the memo included as Figure 13. Note that spooling, originally promised by February 1, was moved to MPE release 3 (April 1, 1973).

January 31 saw another memo, and another schedule push-back. It is included as Figure 14. In it release 2, scheduled for the following day (February 1), is moved to June 30, 1973. Release 3 is pushed to "late '73", and release 4, mentioned in the September 25, 1972 memo, is not even mentioned. Real-time is gone.

Figure 15 shows a memo dated February 1, one day later. It is essentially the same as the January 31 memo, except it leaves out "complete segmentor" from the capabilities of release 2.

In the meantime, the January, 1973 issue of the *HP Journal* (a magazine devoted to the research and products of Hewlett-Packard), devoted a cover story and two other articles on the HP3000. They were entitled

"An economical full-scale multipurpose computer system", by Bert Forbes and Mike Green,

"Central bus links modular HP3000 hardware", by Jamshid Basiji and Arnie Bergh,

"Software for a multilingual computer", by Bill Foster,

and "Single operating system serves all HP3000 users", by Tom Blease and Alan Hewer.⁷⁸

This was Hewlett-Packard's most formal announcement yet of the HP3000. The first article stated that the primary purpose of the HP3000 was "to provide, at low cost, a general purpose computer system capable of concurrent batch processing, on-line terminal processing, and real-time processing, *all in the same software.*" (emphasis added)⁷⁹

Also in early 1973, George Newman, the divisional manager since after the cancellation of the Omega project, moved to HP's calculator division. He was replaced by Bill Terry.

SPRING 1973: THERE IS A PROBLEM...

By Spring 1973 the HP3000 project had been worked on — either as Alpha or as Omega — for almost 5 years, and had cost over \$20 million to develop. It was "by far the costliest [program] in the company's history".⁸⁰

Some of the engineers felt that considering the HP3000 as having been in development since 1968 was a little unfair. They felt that the "clock should have been reset" when the Omega was canceled and Alpha began. But management didn't see it that way.⁸¹

EWLETT ¹²/₇ PACKARD

19110 • Head Office: 100 Lakeshore Blvd., Toronto, Ontario M5S 1A5, Telephone 416 291 4100

Jim Cockrum

December 11, 1972

Bob Grimm 70
Bill Nilsson

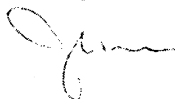
MPE Release II (February 1973)

cc: Dave Crockett
Bill Foster
Mike Green
Dick Hackborn
Larry Turner - Geneva
Steve Vallender

Will you please make the necessary people aware of the new features to be in MPE.

1. The complete file security mechanism will be implemented. This includes user specification and system verification against the account, group and user structure.
2. Accounting and logging will be provided.
3. System bump and reload capability for user files as well as for the system. The configuration update mode will also be available which allows the merging of new modules in the system library with existing modules.
4. A dynamic debugging capability which permits inserting program breakpoints from a terminal.

These new functions will be reflected in the MPE reference manual. Spooling and 202 modem support are scheduled for release III (April 1973).



JC/mr

Figure 13 Memo 12/11/72

HLETT ^{HP} PACKARD

DATA SYSTEMS • 11020 Wolfe Road, Cupertino, California 95014, Telephone 408-257-7000

FROM	Bill Foster	DATE	January 31, 1973
TO	Distribution	SUBJECT	MPE Schedule

Release 2

1. To QA on 4/13/73
To Distribution on 7/15/73
ERS: 3/9/73
IMS: 6/30/73

2. Capabilities

Complete accounting and logging
Complete file security
Operator and system manager capability
202 Modem (2 wire-half duplex)
Backup/Reload
Disc Error recovery
Scheduling & I/O performance improvements
Complete segmenter

Further investigation is required for the implementation schedule of drivers.
The priority of driver development is as follows:

1. Paper tape punch
 2. Paper tape reader
 3. Calcomp plotter
 4. 250 CPU card punch
- } (probably April-June window)

Figure 14 Memo 1/31/73

HEWLETT  PACKARD

DATA SYSTEMS • 15000 Zolfe Road, Cupertino, California 95014, Telephone 408-257-7000

FROM Jim Cockrum

DATE February 1, 1973

TO Hank Cureton
Jim Treybig
Dave Sanders
Bill Foster
Tony Turner
Joe Bailey
Roger Ueltzen
Bill Nilsson

SUBJECT MPE Schedule

On Wednesday, Jan 31, 1973, a meeting was held to review and discuss the proposed second release of MPE. The results of this meeting follow:

Release A consists of all capabilities described in the presently published manuals (excluding the accounting capabilities scheduled for release B). Any described capability that does not work should be considered an error and so documented via Al Mitchell.

- . Release B will be made available to QA on 4/13/73.
- . Release B will be given to distribution on 7/15/73.
- . Release B ERS will be completed on 3/9/73
- . Release B IMS will be completed on 6/30/73
- . The additional capabilities in release B will consist of
 - .. complete accounting and logging
 - .. complete file security
 - .. operator and system manager capability
 - .. 202 Modem (2 wire-half duplex)
 - .. Backup/Reload
 - .. Disk error recovery
 - .. Scheduling and I/O performance improvements

The major capability difference between the above mentioned release B and the previously mentioned release 2 of MPE (memo HP 3000A Status from Jim Schmidt, dated 25 September 1972) is the lack of spooling.

Figure 15 Memo 2/1/73

In late 1972, the sales forecasts for the HP3000 had been downgraded to 80 for the year, instead of 120ⁱ. But then, in April of 1973, Bill Terry appeared at an IEEE conference and said that it was now downgraded to 30. Further, at a meeting with security analysts, Mr. Terry said that the software wasn't running as well as hoped, and that it wouldn't support the original plan of 16 users.⁸² (Note the "spin control" of stating that the original plan was 16 terminal users. The original plan had envisioned more than 16 users.)

Even though the software was not yet defect-free, the hardware was remarkably robust. One of the other early customers, Yale New Haven Hospital, stated that "the hardware has given very little trouble".⁸³ But other early customers were not having as easy a time with the HP3000. An early system was loaned to RAIR Inc., a timesharing company in Mountain View, to test timesharing, but it was returned. National Bank of Detroit had trouble with putting more than 4 terminals on it.

It was becoming clear that the HP3000 was having serious problems. These problems were noticed all the way up the company chain, to Bill Hewlett himself. Hewlett asked Barney Oliver (who would soon become head of all of HP's corporate R & D programs) to come in and save the product. He had stepped in earlier and pulled the HP21XX line out of a mess. Oliver refused.⁸⁴

Word was all over the company about the HP3000. That summer, at the HP company picnic, there were conversations at how much money Hewlett-Packard was pouring into the 3000 "rathole" and when the project would be shut down.⁸⁵

It was time to take action. Paul Ely, General Manager of the Microwave Division, was asked to take over Data Products Division and fix the HP3000 problems. Ely had turned around the Microwave Division earlier and had a reputation as a "get the job done" kind of manager, a driver who could intimidate anyone easily.⁸⁶ His management style was not necessarily a typical HP one.

ⁱ HP eventually shipped 120 units for the year by the end of 1975. From "Hewlett-Packard takes on the computer giants", *op. cit.*

About this time there were a few changes in the Marketing area. Jim Treybig, one of the sector marketing managers, left the company and eventually helped found Tandem Corporation, and Ed McCracken became Marketing manager for the entire division.

SUMMER 1973: "WOW OUCH"

The new management team, once in place, acted quickly. Production of the HP3000 was halted, and more importantly, the decision was made to recall the entire existing base of HP3000s.

At this time, Dave Packard sent a memo to the HP3000 team. It was only two lines long and said, essentially, that they would never again announce a product that did not then currently meet specifications. Its succinctness revealed how displeased Packard was with the outcome of the program.

Steve Vallender, the lab manager, made a copy for every member of the department. Before sending it out, he made a little comment on Packard's memo, scribbling the words "Wow" and "Ouch" in the margin. The memo was henceforth (and still is over 25 years later) referred to as the "Wow Ouch" memo.⁸⁷

The next few weeks were, as Ed McCracken would later say, "one of the worst two weeks of my life".⁸⁸

McCracken went back to the customers and "de-committed" each and every HP3000 sold. In essence, he said that:

1. HP could not bring the software components of the system up to full specifications before Fall 1973.
2. HP was devoting "maximum resources" to correcting the problem.
3. The 3000 system would support no more than 4 to 6 simultaneous users.
4. Image, for those beta testing the product, would be delayed until January 1, 1974.⁸⁹

Some customers literally broke down and cried. Many of them were extremely loyal to HP, and believed in the machines as strongly as any developer at the Cupertino lab.

The customers were offered HP2000 timesharing systems instead. Some accepted them as a stopgap, but still pushed for their HP3000's as soon as possible.

One customer, Anderson College in Indiana, threatened the company with a breach-of-contract lawsuit. It was only averted by the decision of the plaintiffs to contact Bill Hewlett first before filing the papers. Hewlett issued an edict to do anything HP could to appease Anderson rather than going to court.⁹⁰

THE (NEW, IMPROVED) HP3000, PART II

"if you have a flop in a void, you may have time to resurrect it and be a success still in that same void" HP3000 Engineer Jim Holl⁹¹

Soon after the withdrawal of the HP3000, lab manager Steve Vallender left the company. Bill Foster, who is now CEO of Stratus Computers (a competitor to Treybig's Tandem), took the reins of the lab from his earlier position as a section manager in the department.

MPE-B

During the shutdown it was the operating system designers who were in the "hot seat". The hardware was rather well-evolved; it had very few bugs upon release. Hewlett-Packard was in essence a hardware company that did not yet consider software a "core competency."⁹²

After the 3000 withdrawal, the MPE lab staff diminished to about 8 programmers, all working on MPE-B. Mike Green, one of the brightest engineers at Cupertino, worked personally on the operating system. Bob Miyakusu, from the file system area, led the effort. The other applications were essentially finished.^{93 94}

The MPE section attained a frenzy of development. The lab was still constrained by availability of machine time, so the staff worked 24 hours a day on the software. The operating system section split into two 12-hour shifts, to make sure that the computers were being used whenever possible. When one of the recalled machines came back, the lab grabbed it and put it to use.⁹⁵

Functionality was also redefined at that time. The remaining real-time code in MPE, unofficially cancelled earlier, was deleted. On several occasions, naturally, a real-time module in MPE was removed, and other, supposedly non-related code promptly crashed. There is probably some real-time code buried deep in MPE to this day.⁹⁶

After about six months spent ironing out the bugs in MPE, the HP3000 was ready for re-release. It was now able to run 8 interactive users. This in itself was impressive; after the recall, outside experts had predicted that it would take up to "two man-years" to fix MPE.⁹⁷

OCTOBER 1973: RE-RELEASE

The new HP3000 was re-released in October, 1973. It now had MPE-B, and was thirty per cent faster thanks to tweaking by the hardware guys. It also cost twenty per cent less.^{98 99} The press praised the new machine for having its "...sights and promises better aligned with performance..."¹⁰⁰

Figure 16 shows an advertisement published after the re-release; it emphasizes the batch and timeshare capabilities of the machine. Notice that, near the bottom of the ad, the customer is urged to get "your new brochure" (emphasis added).

NOVEMBER 1973: EXIT TO TANDEM

Meanwhile, the HP3000 project lost several key contributors. Jim Treybig, a marketing section manager who had left in early 1973 to join Tom Perkins (ex-General Manager of Data Products Division) in a venture capital company, started Tandem Computing. Tandem was to produce (and successfully has to this day) a "fault tolerant" mini-computer. Their machine would have redundant circuitry and other hardware to insure maximum up-time even in a catastrophic situation.

Hewlett-Packard was going through a slow period, and the company announced that it would shut down all of its divisions during the Thanksgiving holiday. During that week, Treybig interviewed many of his old comrades at HP to recruit them for his new company. More than a few accepted the offer to join Tandem at its inception. It should be noted that these people left with no ill-feeling from Hewlett-Packard, and there was very little animosity between HP and the newly-formed Tandem.¹⁰¹ Nonetheless, the departure stripped HP of considerable talent: Mike Green (who wrote Timeshared BASIC), Tom Blease (who had worked on SPL and was an early proponent of the stacked architecture

**more than batch...
more than timeshare...
less than \$170,000***

**Hewlett-Packard's 3000.
The New
Multiprogramming
Multilingual Computer.**

Full system capability from multiple interactive terminals. For your new brochure call (408) 257-7000, Ext. 3000.

Or write



Sales and service from 172 offices in 65 countries.
1501 Page Mill Road, Palo Alto, California 94304

22408 *Domestic USA Price Only.

Figure 16 Advertisement c.1973

for the Omega/Alpha), and Bob Miyakusu (who led the MPE re-work during the recall period), were among the many who left.¹⁰²

DECEMBER 1974: ADVENT OF THE CX

In the twelve months after the re-release of the product, the lab was working on the next version. It was to be called the HP3000 CX, to distinguish it from the earlier models (which were to be called, eventually "Series I").¹⁰³

The CX embodied several key improvements in architecture, of which the most important was probably the main memory. The hardware designers returned to the original bus architecture and found a couple of free bits that had been set aside. Using these, they could increase the system's capacity to 128 Kbytes of main memory. The type of memory was also upgraded when Paul Ely issued an edict, "No more core memories", and they promptly went away¹⁰⁴; HP used semiconductor memory in the new series and finally left the '60s behind.

Also gone was the wire-wrapped backplane, where the circuit boards were attached. The Series I backplane looked like a board splashed with spaghetti, and generated some interesting parasitic electrical currents that plagued the early 3000s. Mysteriously enough, these effects went away if one took some cardboard and stroked the backplane wires.¹⁰⁵

A further improved operating system, MPE-C, was included, as well as RPG — to compete with IBM — and, probably, the first COBOL compiler on a minicomputer. The COBOL was included to comply with a 1960 Department of Defense specification stating that all machines purchased by them must be supplied with a COBOL compiler.^{106 107}

Bill Hewlett originally opposed installing COBOL on the 3000, for fear of suggesting that HP wanted to take on IBM; but the demands of the market eventually changed his view.¹⁰⁸

And finally, HP's database management system IMAGE was sprung on an unsuspecting market. It was a \$10,000 option at first, but a year later was offered free and bundled with every HP3000 system.

Many people mark this as the time when the 3000 left infancy, and began its successful ascent. This is also where my story ends.

Copyright [c] 1995 by Christopher Edler as an adaptation of a thesis submitted in partial fulfillment of the Master's of Science degree at California State University, Chico.

About the author:

Christopher Edler works as an Information Technology Specialist at Hewlett-Packard Company in Boise, Idaho. He is currently completing a Master's of Science degree at California State University, Chico. He also holds a Master's of Business Administration degree from the University of Arizona and a Bachelor's degree in Economics from the University of Michigan.

He has a wife, Kirsten, and three young children.

NOTES

- ¹ Hank Cureton, interview by author, tape recording, Cupertino, CA., 14 February 1995.
- ² Monty Python and the Holy Grail, "The Tale of Sir Lancelot", Scene 18.
- ³ Carson Kan, Interview with the author, tape recording, 14 February 1995, Cupertino, CA.
- ⁴ "A 10 Year History of the HP3000," *The HP News* (1982):4.
- ⁵ Bob Green, Correspondence with the author, December 5, 1994.
- ⁶ George Cullison, Interview with the author, 13 February 1995, Cupertino, CA.
- ⁷ James W. Cortada, *Historical Dictionary of Data Processing* (Westport, CT; Greenwood Press, 1987).
- ⁸ Rick Justice, interview by author, tape recording, Cupertino, CA., 7 April 1995.
- ⁹ Rich Edwards, "HP 3000 Systems Family History [September 1981]." Internal memorandum.
- ¹⁰ Michael D Green, Bert E Forbes, "An economical full-scale multipurpose computer system," *HP Journal* (January 1973):2-8.
- ¹¹ Bob Green, *op. cit.*
- ¹² Rocky Graziano, interview by author, 21 March 1995.
- ¹³ Phyllis Loise Egan, "Hewlett-Packard, From Instrumentation to Computers: A Study of Strategy" (Masters Thesis, San Francisco State University, 1984).
- ¹⁴ *HP ALPHA Project Data Sheet* (Palo Alto, Hewlett-Packard, 1970).
- ¹⁵ *Ibid.*
- ¹⁶ Graziano, *op. cit.*
- ¹⁷ Kan, *op. cit.*
- ¹⁸ Bill Foster, interview by author, 27 April 1995.
- ¹⁹ Jim Holl, interview by author, tape recording, Cupertino, CA., 14 February 1995.
- ²⁰ Arnie Bergh, interview by author, tape recording, Los Altos Hills, CA., 13 February 1995.
- ²¹ Bob Green, *op. cit.*
- ²² *Alpha 1 External Reference Specifications* (Palo Alto, Hewlett-Packard, 1970).
- ²³ Mark Matoza, interview by author, tape recording, Palo Alto, CA., 7 April 1995.
- ²⁴ Steve Vallender, "Alpha Multiprogramming System, [28 August 1970]." Internal memorandum.
- ²⁵ Bergh, *op. cit.*
- ²⁶ *Alpha 1 External Reference Specifications, op. cit.*
- ²⁷ Bob Miyakusu, interview by author, tape recording, Cupertino, CA., 14 February 1995.
- ²⁸ Kan, *op. cit.*
- ²⁹ Vallender, *op. cit.* pp. 1-2.
- ³⁰ Rene Moreau, *The Computer Comes of Age*, (Cambridge, MA; MIT Press, 1984).
- ³¹ U S Patent Number 3820079 "Bus oriented, modular, multiprocessing computer." issued to A. Bergh, B. Forbes, J.O. Hamilton, J. O. Mixsell.
- ³² Daniel Siewiorek, C Gordon Bell, and Allen Newell, *Computer Structures: Principles and Examples*, (New York; McGraw-Hill, 1982).
- ³³ *System/3000 Software General Information* (Palo Alto, Hewlett-Packard, 1971).
- ³⁴ Bill Foster (lecture presented to HP3000 R & D staff, Cupertino, CA., c1971).
- ³⁵ *Alpha 1 External Reference Specifications, op. cit.*
- ³⁶ ALPHA Multiprogramming System, *op. cit.*
- ³⁷ ALPHA Multiprogramming System, *op. cit.*
- ³⁸ *HP ALPHA Project Data Sheet, op. cit.*
- ³⁹ "HP adds time-sharing system/3000," *ComputerWorld* (17 November 1971): 30.
- ⁴⁰ *Ibid.*
- ⁴¹ *Ibid.*
- ⁴² "System/3000 runs into time-share problems," *Electronics News* (30 April 1973): 1
- ⁴³ Bob Green, et. al., *Image/3000 Handbook*, (Seattle; Wordware, 1984).
- ⁴⁴ John Bale, interview by author, tape recording, Cupertino, CA., 13 February 1995.
- ⁴⁵ Bob Green, et. al., *op. cit.*
- ⁴⁶ W C McGee, "Generalization: Key to successful electronic data processing" *Journal of the ACM* 6 (January, 1959):1-23.
- ⁴⁷ James W. Cortada, *Historical Dictionary of Data Processing: Technology*, pp. 126-27 (Westport, CT; Greenwood Press, 1987)
- ⁴⁸ E H Sibley, "The development of data-base technology," *ACM Computing Surveys* 8 (1976): 2-7.
- ⁴⁹ Cortada, *op. cit.*
- ⁵⁰ J P Fry, E H Sibley, "The evolution of data-base management systems," *ACM Computing Surveys* 8 (1976): 7-42.
- ⁵¹ R W Taylor, R L Frank, "CODASYL database management system," *ACM Computing Surveys* 8 (1976): 67-101.
- ⁵² Cortada, *op. cit.*
- ⁵³ John Bale, *op. cit.*
- ⁵⁴ John Bale, *op. cit.*

- ⁵⁵ Bergh, *op. cit.*
- ⁵⁶ Ed McCracken, interview with the author, tape recording, 24 April 1995, Mountain View, CA.
- ⁵⁷ Graziano, *op. cit.*
- ⁵⁸ Cureton, *op. cit.*
- ⁵⁹ Holl, *op. cit.*
- ⁶⁰ "HP 3000: A new way of thinking," *Measure* (Magazine for Hewlett-Packard Employees) (November, 1973).
- ⁶¹ McCracken, *op. cit.*
- ⁶² Frank Hublou, interview by author, tape recording, Cupertino, CA., 14 February 1995.
- ⁶³ McCracken, *op. cit.*
- ⁶⁴ Holl, *op. cit.*
- ⁶⁵ Jim Peachy, interview with the author, tape recording, 6 April 1995, Mountain View, CA.
- ⁶⁶ *Ibid.*
- ⁶⁷ *Ibid.*
- ⁶⁸ *Ibid.*
- ⁶⁹ Graziano, *op. cit.*
- ⁷⁰ Cureton, *op. cit.*
- ⁷¹ "System/3000 runs into time-share problems," *op. cit.*
- ⁷² Hublou, *op. cit.*
- ⁷³ "System/3000 runs into time-share problems," *op. cit.*
- ⁷⁴ Bob Heideman, interview by author, tape recording, Cupertino, CA., 13 February 1995.
- ⁷⁵ Graziano, *op. cit.*
- ⁷⁶ "As time goes by," *Datamation* 28 (September, 1982):65.
- ⁷⁷ Thomas Harbron Ph.D., correspondence with the author, Anderson University, Indiana, 5 February 1995.
- ⁷⁸ *HP Journal* (January 1973).
- ⁷⁹ *Ibid.*
- ⁸⁰ "System/3000 runs into time-share problems", *op. cit.*
- ⁸¹ Graziano, *op. cit.*
- ⁸² "System/3000 runs into time-share problems", *op. cit.*
- ⁸³ "System/3000 runs into time-share problems", *op. cit.*
- ⁸⁴ Foster, *op. cit.*
- ⁸⁵ Cullison, *op. cit.*
- ⁸⁶ Bergh, *op. cit.*
- ⁸⁷ Hublou, *op. cit.*
- ⁸⁸ McCracken, *op. cit.*
- ⁸⁹ Harbron, *op. cit.*

- ⁹⁰ Harbron, *op. cit.*
- ⁹¹ Holl, *op. cit.*
- ⁹² Bergh, *op. cit.*
- ⁹³ Miyakusu, *op. cit.*
- ⁹⁴ Graziano, *op. cit.*
- ⁹⁵ Holl, *op. cit.*
- ⁹⁶ Miyakusu, *op. cit.*
- ⁹⁷ "System/3000 runs into time-share problems," *op. cit.*
- ⁹⁸ Bergh, *op. cit.*
- ⁹⁹ Laton McCartney, Harvey Wilkinson, "The year mini users shed the security blanket," *Infosystems* (January 1974):24.
- ¹⁰⁰ *Ibid.*
- ¹⁰¹ Miyakusu, *op. cit.*
- ¹⁰² Kan, *op. cit.*
- ¹⁰³ Edwards, *op. cit.*
- ¹⁰⁴ Bergh, *op. cit.*
- ¹⁰⁵ Cullison, *op. cit.*
- ¹⁰⁶ Edwards, *op. cit.*
- ¹⁰⁷ Andrew L Freidman, *Computer Systems Development: History, Organization, and Implementation* (Chichester; Wiley, 1989).
- ¹⁰⁸ Foster, *op. cit.*

IN MEMORIAM: JOHN V. ATANASOFF

Dr. John Vincent Atanasoff, a pioneering researcher in binary electronic calculation, died on June 15, 1995 in Frederick, MD, USA. He was ninety-one.

Dr. Atanasoff is widely credited with exploring or developing several concepts crucial to the construction of the digital electronic computer. He first became interested in the problem of computation at the University of Wisconsin during 1929 and 1930, while he was finishing a dissertation for a doctorate in physics, using only a Monroe calculator and scratch pad to solve equations. Like Konrad Zuse a few years later, he was convinced that complex mathematical problems could be solved by machines that had not yet been built.

At Iowa State University in 1938 Atanasoff began a collaboration with a graduate student in mathematics, Clifford E. Berry, with the goal of building an electronic digital computer. After a year of conceptual development, they began construction of the Atanasoff-Berry Computer (ABC) which could store numbers using capacitors and Bakelite drums, and solve differential equations using concatenations of simple arithmetic steps. Input was by punched cards or keypunch, and the ABC converted decimal input to the binary numbers used for computation. Though it was a very limited device, it contained elements — input-output, arithmetic units, and storage memory — which would later become qualitative to standard computer architecture. If Atanasoff and Berry's work had not been interrupted by the Second World War, it might have resulted in a complete and operable digital computer; but in 1942 the collaborators were obliged to close down the project.

In 1940 and 1941, Atanasoff struck up an acquaintance with John Mauchly, later a primary developer of ENIAC, EDVAC, BINAC and the early UNIVAC computers. The meeting was fateful and in part unfortunate, since any knowledge that the two shared was to become a point of argument in the controversial *Honeywell v. Sperry Rand* case, which pivoted on an attempt to determine who had actually invented the electronic digital computer. That case was settled in favor of Honeywell — and implicitly of Atanasoff — in 1973, but the

controversy has barely cooled in the intervening years.

During and after the War, Atanasoff worked primarily at the U. S. Naval Ordnance Laboratories, in specialties including fire control, acoustics, phonetics, projectile tracking and package handling. He was widely recognized as a pioneer in computing only after the verdict in *Honeywell v. Sperry*, during the middle 1970's; in 1984 he summed up his contribution in a long memoir published in the *Annals of the History of Computing*. For his efforts he received the Computer Pioneer Medal from the IEEE in 1981, and the National Medal of Technology in 1990.

The ANALYTICAL ENGINE extends condolence to his wife, Alice Crosby Atanasoff, to his daughters Elsie Whistler and Joanne Gathers and his son John V. Atanasoff 3rd, and to colleagues and friends.

IN MEMORIAM: J. PRESPEER ECKERT

John Presper Eckert, co-developer with Dr. John Mauchly of the first programmable electronic digital computer in the United States, died in Bryn Mawr, PA, USA, on June 3, 1995. He was seventy-six.

Mr. Eckert was a lifelong computer engineer, continuing in consulting work until shortly before his death; but his name is most often associated with his trailblazing work on problem solving in ballistics, between 1943 and 1946, which culminated in the creation of the ENIAC computer at the Moore School of Engineering, University of Pennsylvania. ENIAC, a truly vast undertaking that contained over 18,000 vacuum tubes and was said to consume 150 kW of power, could compute almost 2,500 times as fast as a human operator with a desk calculator; it was destined to become a legend of electronic engineering. The device ran from 1946 to 1955, performing military computation including gunnery tables, and design calculations for the Manhattan Project.

Taking the lessons of this work to private industry, Eckert and Mauchly co-founded the Electronic Control Company, soon renamed the Eckert-Mauchly Computer Company. This firm was purchased in February 1950 by Remington Rand,

which bankrolled the production of UNIVAC I, the first commercially successful American computer. The prototype UNIVAC was delivered to the U. S. Bureau of the Census in March 1951; ultimately forty UNIVAC I's were sold. An example, possibly unique, is currently displayed in the Computer Museum in Boston, MA.

Eckert was also a primary designer and developer of EDVAC (the "von Neumann computer") and of a pioneering airborne guidance device, BINAC, constructed for Northrop Aircraft. Thereafter he continued to make substantial contributions to computer engineering, earning a reputation for brilliance. In 1963 he reached his highest position at the UNIVAC Division of Sperry Rand, as Vice-President and Technical Advisor to the President. In 1989 he retired from Sperry's successor, Unisys. He received a doctorate *honoris causa* from the University of Pennsylvania in 1964, and the National Medal of Technology in 1968, among many other honors.

The ANALYTICAL ENGINE extends condolence to his wife Judith, to his daughter Laura Phinney and his sons John P. Eckert 3rd, Gregory Eckert and Chris Eckert, and to colleagues and friends.

IN MEMORIAM: GERARD SALTON

Gerard Salton, a leading authority on the design of database retrieval systems and the primary developer of the SMART text retrieval engine, died in Ithaca, NY, USA on August 28, 1995. He was 68 and a professor in the Cornell University computer science department, which he helped found in 1965.

SMART, an acronym for System for the Manipulation and Retrieval of Texts (but sometimes parsed as Salton's Magical Retriever of Text,) was one of the earliest search engines that relied on the frequency of a term's occurrence as one criterion for searches. This innovation increased the efficiency of retrieval as well as convenience for the user; it is the basis for many retrieval systems in use today.

Professor Salton was born in Nuremberg, left Germany during World War II, and reached the United States in 1947. He received a bachelor's degree in mathematics in 1950, and a master's in 1952, from Brooklyn College; he then studied at Harvard, earning a doctorate in 1958 and serving on the Harvard faculty in various capacities. He joined the faculty of Cornell in 1965.

The ANALYTICAL ENGINE extends condolence to relatives, colleagues and friends.

A WEB PAGE FOR THE CHAC

Weeks of work and a lot of deeply appreciated advice have put a CHAC page up on the Web. Its topics include:

- What the CHAC does, and why
- Every issue of the ANALYTICAL ENGINE available in plaintext via ftp
- The story of our (formerly!) imperiled SDS 930
- Pictures of notable micros from our collection
- A comprehensive list of other computer history pages indexed by topic
- A list of other computer history organizations and periodicals
- Plenty of room for suggestions!

Our friends the Wombats e-mail us a hit log every night, and we're gaining a new appreciation of the phrase 'World Wide Web' as the accesses pour in, not only from the US, the UK and Western Europe, but — in smaller but still tantalizing numbers — from Mexico, South America, Japan, and the former Soviet republics. Once more we feel the exhilaration that was so welcome two years ago, when the CHAC announced itself to USENET and to the mailing lists. Honestly, computer history is sparking interest all over the world — and attracting a population of enthusiasts dedicated enough to keep that interest vibrantly alive.

Give us a hit! Cruise over to <http://www.chac.org/chac/index.html> and sample the CHAC's latest online goodies. We'll be looking for you.

SVERDLOFF SUCCEEDS WALLACE AT COMPUTER MUSEUM

Brian Wallace, former curator of historical collections at the Computer Museum, Boston MA, has left that post to continue graduate studies. Brian has given crucial support to the CHAC since our earliest days, and while we certainly wish him luck, we hope he doesn't drift too far from computer history.

His successor is Brent Sverdloff, who comes to the Computer Museum from an administrative position at the Department of Germanic Languages and Literatures, Harvard University. Before that, he served as Archivist of Special Collections of the Getty Center for the History of Art and the Humanities (Santa Monica, CA,) for five years. He has also taught, and translated to and from, Romance languages extensively.

Mr. Sverdloff received a bachelor's degree in Romance linguistics and literature in 1986, and a master's in Spanish and linguistics in 1989, both from UCLA; he was principal flutist with the University's symphony orchestra while still an undergraduate.

The CHAC welcomes this variously talented curator to the rough-and-tumble field of professional computer history. His genius for languages, his musician's timing and his appetite for detective work will all be exercised to the fullest.

Tech Corner: MAKING NEW BATTERY PACKS FOR THE HP-35 CALCULATOR


Douglas W. Jones
University of Iowa

Today, an old HP-35 Calculator is likely to be as useful as it was when it was new, but finding new battery packs for one is difficult. Thus, most people who use one today must put up with the tether to its plug-in power supply.

Fortunately, HP-35 battery packs are fairly simple to make. Three 1.25 volt NiCd cells, wired in series, produce exactly the 3.75 volts the calculator is rated at, and three AA sized penlight cells fit comfortably inside the calculator's battery compartment. Furthermore, the battery charger side of the HP-35 power supply is well matched to the rated 45 milliamp trickle charge current of the AA NiCd cells I bought from Radio Shack.

Before going into detail on the battery pack, it is worth documenting the details of the HP-35 "wall wart" power supply and battery charger. This unit contains a transformer, a voltage regulator, and a current limiter, with the connection to the HP-35 through a 3-pin plug, as follows:

(Figure 1)



The plug is shown as viewed from the bottom of the calculator. The switch S2 is the on-off switch on the front of the calculator. Switch S1 is a spring shim that shorts pins a and c when the calculator is unplugged from the power supply, allowing the battery to power the calculator.


When the supply is plugged in, battery charging is provided by a current limited 50 milliamp 30 volt supply between pins a and b, while operating power is provided by a regulated supply between b

and c; a 3.75 volts, 0.15 amp supply should suffice, but the HP supply I have puts out 5 volts.

I gleaned the above by disassembling both my power supply and calculator; I don't recommend disassembling HP-35 calculators if they work, but mine was broken when I got it. After disassembly, I found that the only problem was dirt under the slider of the on-off switch; and after cleaning and reassembly, it works quite well.


To make a new battery pack, start with three AA NiCd cells, arrange them side by side with the positive ends pointing in alternate directions, and bundle them together with vinyl electrical tape, as shown:

(Figure 2)



The next step is to jumper the cells together and add contact pads at each end. I gently clamped my stack of cells in a vise to solder the necessary connections to each end. To prepare the batteries for jumpering, I melted a button of solder on each end of each cell. Be careful soldering to batteries — the heat boils a bit of the electrolyte inside the battery, and you'll do the least damage by using a very hot iron very briefly; I used a soldering gun. Also be careful not to fill any vent holes in the battery with solder.


(Figure 3)



It is critical to get the polarity right! Make your battery pack exactly as shown above, with the positive pad at the bottom left side of the pack. If you accidentally build a mirror image of this, you will either end up reverse charging your NiCd cells (which can damage them) or you will end up putting a reversed voltage into the calculator (which might damage it).

Finally, wrap tape over each end of the battery pack, as shown:

(Figure 4)



The contact pads should be securely held by both their top and bottom edges so that they are between 1/8 and 1/4 inch apart. The tape covering the top half of the battery pack should cover the pads to a point about 1/4 inch beyond the center line of the battery pack, and the tape over the bottom half of the battery pack should extend about 1/4 inch from the bottom shoulders of the battery shells. I used a narrow strip of tape around the pack for the bottom end, while a full width piece of tape worked fine around the top. At both the top and bottom, I taped over the solder and wires at the ends of the battery pack before securing everything with tape around the pack.

The battery pack I made this way came out a bit on the small side; better small than too large, but as a result, it occasionally slips out of registration with the spring finger contacts inside the HP-35 battery compartment. The problem is not severe enough to make me consider rebuilding my new battery pack; the 3/8 inch squares of exposed contact pad are big enough that it can slop around quite a bit inside the calculator and still make good contact, and the geometry is such that if the battery pack is ever accidentally inserted backwards, it won't hurt anything because it only makes electrical contact when it's inserted the right way.

I don't guarantee that you won't damage your valuable antique calculator by following my instructions, and I certainly don't want to be held responsible for any errors you might make in trying to duplicate my work.

Nonetheless, if you feel confident in your ability to handle a soldering gun and electrical tape, and if you know enough electronics to verify my reverse engineering of the HP specifications for the HP battery charger, you should find these notes useful.

[And one more note: If you aren't comfortable with a soldering gun, throw money at the problem instead by calling

Edu-Calc
27953 Cabot Road
Laguna Niguel CA 92677 USA
+1 800 677-7001
+1 714 582-2637

to see if they still have replacement HP-35 battery packs for \$14.95. They're also a source for other HP accessories. Thanks to Guy Ball for the tip. — Ed.]

SUPPORT FOUND FOR TANDY PORTABLES

Remember that great Tandy laptop you had in college, in the desert, or in some minor war? Is it still in your attic or garage, gathering dust because the evolution of microcomputing seems to have passed it by? Well, you might want to buy some batteries today, because Rick Hanson can help you bring your Tandy or NEC portable back to life.

Hanson, the proprietor of Club 100 in Pleasant Hill, CA, USA, sells a package that will connect the Tandy's disk drive directly to an IBM-compatible PC serial port. His Lapdos II software will then display a file directory of the Tandy drive on the PC screen and permit a variety of file and disk manipulation. Lapdos II is available for \$39.95 and the appropriate CompLink cable sells for \$17.50.

Club 100 stocks an assortment of other useful goodies for the Tandy 100, 102, 200, WP-2, and the NEC PC8201a, and you can request a free catalog by calling +1 510-932-8856 or faxing +1 510-937-5039. Hanson also operates a Tandy support BBS at +1 510-939-1246.

Quick Take: SILVER ANNIVERSARY OF XEROX PARC

To the rigorous aficionado of comp. hist., it can seem that almost every day is an anniversary. (Or is that just us?) But even among the current welter of commemorations, June of 1995 stands out — as the silver anniversary of the founding of the Xerox Palo Alto Research Center. An exhaustive list of PARC's inventions and developments would take several pages, but let's just mention

- WYSIWYG text processing
- Cut and paste
- Bitmapped screens
- Icons and windows
- Mouse-based drawing
- Laser printing
- Ethernet
- Smalltalk and InterLisp

Which means that when you're sitting and working at a Mac, or an Intel box running Windows or GeoWorks, or a newer Sun, an Atari ST, an Amiga, or....well, lots of things....you have cause to remember Xerox PARC fondly. Certainly we do, and the ANALYTICAL ENGINE hopes to bring you a feature article on PARC's history in the very near future.

Quick Take: APPLE II RESOURCE LIST

A new friend of the CHAC, Roger Aitken of Foster City, CA, sent us an Apple II resource list that we planned to include in this issue; but we circulated it to a few friends and they keep adding to it! If you have any interest in the newest Apple II apps, don't miss this resource list — in February's ENGINE whether it's finished or not.

Quick Take: AMIGA PREVAILS!

David McGlone reports in *Z-Letter* #38 that Escom AG, having succeeded in purchasing Commodore's assets, plans production of Amiga 4000 video production systems, entry-level Amiga 1200's, and the Amiga CD32 game machine. The Escom Amiga 4000's will feature a SCSI interface and include SCALA, a multimedia prototyping and production tool from Scala of Norway; they should be available by the time you read this, probably from video hardware retail outlets. The ANALYTICAL ENGINE, once again, declares itself unequivocally in favor of *more Amigas* and congratulates Escom for backing a courageous decision with a big hunk of money.

SPOTTER ALERT

Copies of the ENGINE, the FAQ, and project information have been pouring out to print and broadcast media, especially in Silicon Valley. We do have tearsheets of most of the ink we know about. But is there ink we haven't heard of? Once more, with feeling: If you spot any mention of CHAC or the ENGINE in any periodical, *please*,

- If your copy of the piece is clippable, clip and mail to the Palo Alto address.

- If you can't spare the physical copy, send the text as net.mail to engine@chac.org, or photocopy and fax to the Palo Alto address.
- If you're too busy for that, just send the publication name, date and page number and we'll do the hunting.

Thanks! (And thanks to the spotters who have given us invaluable help with keeping up so far.)

SPOTTER FLASH

"Save vintage mainframe from extinction," trumpets the September 1 issue of *EDN*, leading a tidy quarter-page about the 930 Rescue and our urgent need for funds. The column winds up with "If you have any ideas or want to help, please call or e-mail CHAC...." and the power of the press delivered another solid hit as, through that, we located almost a dozen former members of SDS' technical staff or management. Enduring thanks to Max "Bebop" Maxfield for suggesting the story, and to *EDN* editor Fran Granville for writing and running it.

MONEY, THE WORLD, AND THE ORDERLY PROGRESS OF DAYS

Just a reminder that, even though we raised enough in donations to save the 930 from getting tipped, the CHAC is still — in the old phrase — broke at a higher level. As soon as the last bills from the Rescue are paid, and this issue of the ENGINE is printed, no more than pocket change will be left in our bank account.

Take a careful look at this robust copy of the ANALYTICAL ENGINE and ask yourself — isn't this the time to subscribe? Thirty-five dollars a year will still bring you four of our new, thicker, illustrated issues. Subscribe today, if you don't yet, and join the CHAC's perennial celebration of the world's most fascinating technical history.

Book Review: BEBOP TO THE BOOLEAN BOOGIE

Clive "Max" Maxfield

Solana Beach, CA:
High Text Publishers, Inc., 1995
471 pages, \$35.00 (paper)
ISBN 1-878707-22-1

Reviewed by Tom E. Ellis

The first impressive thing — and not the last — about this book is its sense of humor. Start, naturally enough, with the cover: three-dimensional gate logic symbols flying over a printed circuit board, trailing a glowing blue staff of musical notes. Follow the trail down to a couple doing the bebop, to the music of a saxophone player who looks remarkably like the author, Clive ("Call me Max") Maxfield. This is a survey text in micro-electronics? Yes, and a very good one.

The humorous and inimitable style of the *Foreword*, *Acknowledgments*, and *About the Author* set the tone for the entire book. Maxfield's keen interest in his subject matter becomes obvious as he takes you on a journey that you won't soon forget. As Pete Waddell mentions in the *Foreword*, technical books are usually "as dry as West Texas real estate" — which I take to heart, having been born and raised there — but this book is anything but dry; while it is jam-packed with highly technical information, the format and style make it incredibly readable. Illustrations almost outnumber paragraphs as complex topics are rendered in the most understandable terms.

As you make your way through the book, Maxfield's talent for the absurd will drive you forward, while you wonder what's around the next corner. One minute you'll explore the traces inside a memory IC; the next, you'll learn how to perform duo-decimal calculations on your fingers, using your thumb to point to finger joints "while maintaining a free hand to throw a spear at someone...." Yes, this is a textbook with nearly 500 large-format pages, but it becomes as compelling as a Christie mystery. And although the foreword suggests that you "plunge in and out [of the book] as you wish," the material presented has been elegantly arranged to build fundamental understanding before drilling

down to the nitty-gritty. Once hooked by Maxfield's light-hearted wackiness, you're likely to read *Bebop* straight through and absorb a huge amount of valuable knowledge, somewhat in spite of yourself.

Bebop assumes no prior understanding of electronics and begins, engagingly, with "a fun-loving fool sliding down a ramp" to illustrate the difference between analog and digital viewpoints. Even when you move on to *Atoms, Molecules and Crystals*, you discover reassuringly that electricity is only "vast herds of electrons migrating from place to place," and that the science of electronics is all about "deciding where they can roam, and determining what they are going to do when they get there." Like a good *Jeopardy* board, *Bebop* reveals each new truth with a just-out-of-reach familiarity that rings the "I knew that!" bell in your head.

But in almost no time you're saying "I knew that!" about "correlated electron movements in conducting planes," because the sure-footed author shepherds you through inner workings at the particle level and then connects, through the most lucid illustrations, the theoretical world of circuit design to the physical world of substrates and semiconductors. Packaging technologies are compared, as well as common techniques of forming conducting paths, vias and lead-holes. Breadth as well as depth of knowledge is served with full explorations of *Alternative Numbering Systems, Binary Arithmetic, Boolean Algebra, ICs, Programmable ICs, ASICs, Circuit Design, Hybrids, and Multichip Modules*. To read *Bebop* is to understand absolutely the progression from particles, to paths, to circuits, to modules, to components, to processes.

Maxfield then declares that his "potpourri of technologies...[has] been offered for your delectation and delight." But as the final chapter rolls to a close, he quotes Churchill: "Now this is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning." This is the literal truth since the next hundred and fifty or so pages are filled with "oodles of yummy appendices" several of which discuss various logic concepts, and one of which includes several C programs for "the computer programmer that lurks within us all." *Bebop* is rounded out with a robust *Abbreviations and Acronyms* listing, a wonderful *Glossary*, and — rarity of rarities — an index comprehensive

enough and well enough organized to be truly useful. Even the plentiful sidebars and some of the wacky quotes have been indexed, lest you miss “the blind obedience of fools” and “the purple tint in San Francisco Bay.” And at the very end of this extraordinary wealth of information, the diligent reader is rewarded with a great “No-Holds-Barred Seafood Gumbo” recipe. I’m not kidding.

You could not find a more delightful way to absorb complex electronics theory, meanwhile speculating on such things as how life developed on Earth and why window shades spontaneously roll themselves up. Again and again you’ll be fascinated by the previously obscure, and you’ll come away thinking that microelectronics isn’t difficult after all. But don’t kid yourself; Maxfield’s unrelenting fun only facilitates his expert, lucid and very approachable descriptions of how and why things work. *Bebop* is a serious asset for scientist and layperson alike.

ACQUISITIONS

[Even though we currently have no more storage for micros, we can’t seem to stop collecting them, so this summer we concentrated on little ones.]

CONVERGENT WORKSLATE

Douglas Mandell

Sometimes an idea before its time has a certain charm that’s hard to articulate. So it was with the Convergent Workslate, which first appeared in 1983 and was determined to be a notebook computer, no matter how many hurdles it had to jump to get there.

At 8.5”x11” and about an inch thick, it’s a true notebook and weighs roughly three pounds without peripherals. The screen is a 16-line by 46-char LCD, integral with the computer and the size of a large index card. Overall the machine is very handsome, in “slate” gray with lighter gray and green keys, and a diamond-shaped cursor wobble-key that looks like it escaped from a video game.

While the Workslate could be fairly versatile, it was fundamentally designed as two things: a portable spreadsheet and a communications terminal. We’re not sure yet whether ours has an internal modem or whether communications required the optional, external, and (so far) missing box called

the CommPort; but we’ll know, and say, more about this intriguing little slab after we get it booted. Thanks, Douglas!

HP 110 and PORTABLE PLUS

Roger Sinasohn *et al.*

In these days of fan-cooled processors and expanding keyboards, when a five-pound notebook can pack fully as much punch as a “real” (desktop) computer, it’s salutary to remember the time — 1984 or so — when an easily portable computer was inseparable from real adventure.

The HP 110 and its gutsier successor, the Portable Plus, were some of the world’s first laptops. As usual with HP’s computers, they feature robust construction — the off-white plastic clamshell cases would probably survive a small explosion — and use of innovative techniques to assure expandability; the “bottom drawer” expansion chassis of the early HP desktop computers translated well to these portables. As for the sort-of-green-on-sort-of-yellow LCD screens....that’s part of the adventure, and the intrepid user quickly gains the knack of getting the desk lamp at the exact proper angle to the screen.

With 80C86’s as processors, these are slow by modern standards but certainly usable. A suite of applications and applets in ROM, including Lotus 1-2-3, eases impact on conventional memory.

Warnings to the unwary: <Return> and <Enter> are *not* the same keystroke, and the re-boot sequence is *not* Control-Alt-Del. Now let’s see — was that Shift-Extend-Contrast, or Shift-Contrast-Enter, or....

Thanks to Roger and another, anonymous donor, we now have one 110 and several Port Pluses, as well as two external stiffy drives and two Think-Jets. Your favorite Computer History Association is ready for the road, in style.

PMC MICROMATE

Kevin Rudd

Nowadays a desktop computer tends to be manufactured in one of two formats — beige flat box or beige tall box. But not so long ago, computer designers were more experimental in their case layouts....resulting in fascinating devices like the Micromate.

This Z80-based micro is the size and shape of a small toaster. One half of the case is occupied by a 5.25" floppy drive mounted on its side; the other half contains a quite tidy motherboard and some ports. If it were sitting next to a terminal with a swivel base, the whole shootin' match would hardly take up more room than the terminal.

We haven't booted this one yet either but — given our fond memories of Z80 micros — we can hardly wait. Thanks, Kevin.

LETTERS

DATA DEGRADATION ON 7-TRACK TAPE

The University of Iowa Physics Department has a longstanding problem with the preservation of archival data. They built the Explorer I instrument package back in the 1950's, and they have been building satellites and instrument packages for satellites ever since. Over the years, they have accumulated a huge archive of data recorded off of satellite downlinks, and most of this is on magnetic tape in various recording formats.

The old data is quite useful today, particularly when long-term trends are being studied, but maintaining access to old data is difficult. One of the major headaches the University of Iowa physics department faces, for example, centers around the maintenance of their 7-track tape drives.

These drives are connected to a VAX-11/780, itself something of an antique these days, but it is hard to find newer machines with reel-to-reel tape interfaces that are compatible with any kind of 7-track drive.

Having working 7-track tape drives, the University of Iowa physics department has found that data conversion to modern media can be a significant source of income, but recently, they have had diffi-

culty finding new heads for their 7-track drives, so they are using them only very sparingly for their own internal work.

It is clear, at this point, that the availability of new (or as good as new) tape heads is currently the factor that limits the ability of the few sites that still maintain 7-track tape drives to read and convert the data from old tapes into modern data formats. I suspect that, unless someone has a cache of spare 7-track heads hidden away somewhere, or unless someone somewhere starts manufacturing small lots of new 7-track heads on a custom basis, we will altogether lose our ability to recover data from the vast storehouses of 7-track tapes that are out there!

The data we will lose includes not only the scientific data that motivated the University of Iowa physics department to preserve and maintain a set of 7-track drives, but also the 1960 census and large libraries of corporate records.

Doug Jones

jones@cs.uiowa.edu

HP'S EARLY COMPUTERS: OLIVER INTERVIEW

In reading your interview with Barney Oliver, I was tickled that you cited my Spirograph program for the HP 9100, but you were wrong in saying that the program was an old one! While I used a 9100 back in 1972, I didn't write the Spirograph program until about a month before I sent it to you. I had just been given a surplus 9100B, with plotter, and the obvious exercise was to turn it into a Spirograph. I'd written a similar Spirograph program for a CDC 6600 back in the 1970's, and with both a plotter and a the polar-to-Cartesian function on the HP 9100, the algorithm was an obvious application to try out.

Later in your interview, you commented on the relation between HP and DEC. Barney Oliver mentioned that, for a while, David Packard was thinking about buying DEC, but that he hadn't heard of HP buying the PDP-8 on an OEM basis. My understanding was that an HP subsidiary was, briefly, DEC's largest OEM customer for the PDP-8, in the early part of their production run. Combining what with what Oliver said, it sounds like the subsidiary was probably Dymec, and I wonder

if this relationship might explain HP's brief flirtation with the idea of buying DEC.

Finally, I want to note that, although I have a somewhat battered HP 21MX sitting next to me as I write this, ready to be boxed up and sent to CHAC, my only real experience with the 2100 family is with the HP2115B; I spent the summer of 1971 building custom I/O interfaces for one at the University of Michigan Physics Department. That machine, with the interfaces I helped build, ran for many years as part of a high energy physics experiment at Fermilab outside Chicago. It's worth noting that in addition to some TTL, our interfaces used an awful lot of DTL chips!

Doug Jones
jones@cs.uiowa.edu

HP'S EARLY COMPUTERS: 21xx MEMORY

Two points about the Schoendorf interview:

The 2100 was *not* based on semiconductor memory; at least initially, it used core. I suspect that it may later have used semiconductor memory (and the guy who wrote the emulator seems to think it is within reason to make it do so), but the schematic set indicates core planes. *sigh*. Now if I could just find some for my machine.

Early in the article he makes reference to the 2100 "making provision for" a megabyte of memory, at Bob Frankenberg's initiative. That's very interesting, because the 2100's memory cage is laid out so you can't put more than 32KW into it; there are two sections, each with space for two core assemblies, each of which can be 4KW or 8KW depending on the driver card installed for that section.

The successor 21MX CPU had a Dynamic Mapping System (DMS) feature that might be described as a bank-switching scheme -- you could map different physical regions into certain sections of the address space under program control, and thus address more than 32KW of real memory, but no more than 32KW was visible at any given time. I think this was an option comprising additional hardware and microcode ROMs to extend the instruction set.

I really need to pull the 21MX manual and check this, because it's been a while. I should also pull the 2100 schematics and find out whether it really brought out additional address lines -- for all I

know, maybe the DMS stuff was originally done on the 2100 after its introduction.

This in conjunction with the mention of semiconductor memory makes me wonder if what's identified as the 2100 in the interview is an inadvertent confusion of the 2100 and the 21MX.

Frank McConnell
fmc@aphasia.us.com

QUERIES

ASI SPEEDER: NEW OWNER PUZZLED

Does anyone know of, remember, or use a beastly such as the "ASI Speeder" Addressing Systems International Model 131 A. What is it? I've "obtained" one, and it seems to be functioning - a keyboard/printer/disk drive unit but that's about all I can tell. It wants a disk on startup - if anyone has a copy of this disk (maybe an OS disk?) or manuals, or *anything* about it, it would be useful!

Michael Brown
mjb@dcs.warwick.ac.uk

ATARI ATW800: DOCS AND VENDORS WANTED

Hi. I'm trying to get some information about an Atari Transputer Workstation, model ATW800. It seems to date from about 1990. It's in a large tower-style case. On a cursory examination, it seems to contain most of the innards of an ST, along with a large board with a T800 transputer on it and another board that seems to be a video card of some sort. All I've got is the box and the keyboard; no monitor, no mouse, no manuals. I'm told that it runs Helios.

Basically, I'm looking for any information about this machine; in particular, I'd be very grateful for pointers to any companies that might still be able to supply documentation for the machine and its software.

J. Lothian
jlothian@castle.ed.ac.uk

BAUDOT CODE LISTING WANTED

Back in the ancient days...Teletype made a model that used a five level (read *five bit*) code called Baudot. I have a couple of interesting (to me) questions (or requests):

- 1) Can someone provide a listing of the *entire* Baudot five level code?
- 2) From info I have, the code did *not* represent letters in alphabetical order. What are the reasons behind the arrangement of the letters in the code?

Any other information or anecdotes about the Baudot code or Teletype would be appreciated. Please mail me and I will post a summary. Thanks in advance!

Charles Richmond
richmond@plano.net

BLIT TERMINAL: PRETTY MUCH ANYTHING

I have acquired an authentic Bell Labs Blit for which I have no documentation. It works, and the 13-year-old termcap entry from Rob Pike also seems to work, but I'd like to know how to get it running faster than the current 1200 baud. The 'SETUP' key does nothing obvious. Any tips or pointers to Blit resources on the net? Or Blit manuals you don't need anymore? Thanks,

Jon Leech
leech@cs.unc.edu

CANON LBP-CX PRINTER: POSTSCRIPT CONVERSION FEASIBLE?

I've a couple of Canon LBP-CX printers that I'd like to continue using "as is", and also be able to add some circuitry to allow them to be used to print regular ASCII text, Postscript and DVI files. These are printers with the "raw" CX interface. The type of computer system that I now use them with requires this raw interface, so it's not just a simple matter of getting new printers to use.

Besides, the printouts from these printers look very nice. What I'd like to do is be able to flip a switch and use them, via a serial port, with my Sun, or use the raw CX interface with the other systems.

A while back, I recall seeing a book that is supposed to tell one how to convert a Canon LBP-CX to a Postscript printer, but I don't remember the title, or who the author is. Can anyone comment on what all's involved in such a conversion, or on any other things about this book?

Thanks very much in advance for any information that anyone can provide about this!

R. D. Davis
rdavis4@umbc.edu

CROMEMCO XXU CARD: MORE THAN A PAPERWEIGHT?

I was prowling the KC area old and funky computer shops yesterday. I bought a Cromemco card, XXU XDOS 2.14. It has two Motorola chips, a 68020 and a 6888. There is a sticker that says "19k-CM" on the heatsink (if that's what it is). The sticker is hand printed and looks like it was trimmed out with a pair of scissors so I don't have confidence in its significance.

The solder side of the board has two "factory" stickers (i.e. round cornered and dot-matrix printed) with

	L/T #3
and	520-0176-G
	REV. AC

The Motorola chips say

	MC68020RC16E
and	MC6881RC16B.

The aforementioned sticker is on a bent metal plate that covers some socketed IC's and what I take to be an op amp. But I can only see the sides. Silk-screened on the plate is

XXU (tm) Cromemco U.S.PATENT PENDING

A poster in comp.os.cpm said that they might have made a Unix box and this might be for it. Soooo, now what? How do I find a box to put it in, software to make it work and a little expert advice?

C. W. Fertig
cwfertig@axp2.umkc.edu

CURTA CALCULATOR NEEDS REPAIR

Any idea where I might get my Curta repaired? I bought the smaller model over the net but when I received it the mechanism appears to hang up now and then. The crank doesn't have any sort of tactile feedback on being pulled out and I guess that someone pulled it part way out and then forced the mechanism. It will work at times but I would like to see if there is any U.S. source of repairs. I am told that attempting this myself is not a good idea.

Thanks,

Don Taylor

dont@agora.rdrop.com

[We'd tend to agree. A Curta, for those who've never met one, is a hand-held mechanical drum calculator made in Liechtenstein and popular in the late 1950's; it looks like what you'd expect of a black crank-type pepper grinder with eight digits' precision, if it were made in Liechtenstein. They're complex, obsessively precise, and cost lots more now than they did new. — Ed.]

DEC VT101 SETUP/B CODES WANTED

I have a VT100 with the VT101 upgrade board, and I was wondering if anyone had a list of the Setup/B codes (1-4)

All I can figure out so far is speed (obvious), Inverse, and Block-Cursor. TIA,

Jason McMullan

jmcc@m5.vi.ri.cmu.edu

DIGICOMP PLASTIC COMPUTER WANTED

I'm building a small museum of the computers that I have used, and I would love to add the first computers I played with. They were called the Digi-comp I and Digi-comp II Computers. The Digi-comp I had used little flip-flops and was programmed with sections of a drinking straw. The Digi-comp II ran on marbles and gravity.

Bob Roswell

broswell@syssrc.com

EAI ANALOG COMPUTER WANTED

I'd love to find one of these marvelous machines as a nostalgia item. A TR-20 or TR-48 would do fine. Functioning or not, doesn't matter. Must be affordable.

Dick Mills

dmills@albany.net

GENIAC: NOTED COLLECTOR SEEKS COROLLARY MATERIAL

Still busy with non-computer-history activities, but I have managed to capture a 1955 Geniac "Electric Brain Machine" (original model) unconstructed kit in original box. Inventor and manufacturer: Edmund C. Berkeley, who was also author of the classic 1949 book, *Giant Brains*.

I am looking for a advertisement or brochure that portrays the computer in finished form.

If anyone can help, please contact me. Thanks.

Hal Layer

P. O. Box 27676

San Francisco, CA 94127

email: hlayer@sfsu.edu

ph: +1 415-338-2637

HONEYWELL H-316: MONITOR DOCS WANTED

Anybody got any info on OP-16, a real-time monitor program that used to run on the Honeywell H-316 mini? I spent several years working with it about 20 years ago and would like to find a copy of the docs for both OP-16 and the 316 hardware.

Bruce Grant

bgrant@montego.umcc.umich.edu

IBM PCJR SUPPORT DESPERATELY NEEDED

Does anyone have any old manuals or info on support groups for IBM PCjr's? I gave a friend my old PCjr and they are still refusing to upgrade.

They are driving me CRAZY with technical questions....

Any info or help would be greatly appreciated. Thanks.

Michael Ribeiro
LdgImage@ns.net

INFORMATION STORAGE 525WC OPTICAL DRIVE: INFO WANTED

I just acquired an Information Storage Inc. type 525WC optical disk drive. This company is (or was?) in Colorado Springs. The drive is a 5.25" full height device which accepts standard 5.25" disk cartridges.

What I would like to know (lacking *any* docs):

- what is the interface used? It has a 20pin and a 34pin edge connector (like a ST412 or ESDI interface)
- what is the capacity of this beast?
- is it WORM or MO read/write?
- anything else worth knowing.

Help is much appreciated.

Wilko Bulte
wilko@yedi.iaf.nl

MICROPOLIS 1568 AND WD7000ASC: SETTINGS WANTED

I have a Micropolis 1568 760MB ESDI drive jumpered for 1024 bytes/sector. I need to jumper it for 512 bytes/sector for a PC interface.

Anyone got a refcard for this baby? I also need a copy of the jumper settings for a WD7000ASC.

Peter da Silva
peter@nmti.com

OHIO SCIENTIFIC C3 OS NEEDED

I have an OSI C3 OEM. I believe that the hardware is in good working order. The machine will no longer boot. I think that this is because the magnetic image on the disk has not been refreshed in over 10 years.

I got the machine with no software except what was on the hard disk. There was no floppy formatting program, so I was unable to back anything up or make a bootable floppy. When the machine booted, it came up in OS-65U Timesharing. I have three 48K user partitions and would love to see this machine come up in something other than the ROM monitor.

Thanks,

Bill Sudbrink
bill@umsa7.umd.edu

OLYMPIA OL-H004 PROGRAMMING DOCS URGENTLY WANTED

I have recently acquired a "new" Olympia OL-H004 (Panasonic RL-H1400) HHC (Hand Held Computer)....the HHC is a "palmtop" computer circa 1982, with 6502 MPU, 4k RAM (Wow! <g>), its own built-in operating system/programming language called SNAP, 26 char x 1 line LCD display; it accepts custom options & ROM modules called Capsules. I have found it very handy for note taking....[and] would like to tinker with it to make it even more handy.... I am very keen to obtain programming information/software/etc. for it and for the SNAP built-in interpreter. If anyone could supply me with these materials, or direct me to a source for them, I'd certainly appreciate it! (I would, of course, gladly recompense you for your shipping/media/etc. expenses.) Alternatively, if you could suggest a better place to ask around, I'm open to suggestions.

Cordially,

Richard Kanarek
72371.111@compuserve.com

RCA COSMAC VIP: NOT MUCH WITHOUT DOCS

I have two RCA COSMAC VIP computers for which I'd like docs. Anyone out there have a copy they'd either be willing to be part with or copy? I'd be willing to pay a reasonable fee for either.

Please e-mail if you can help me out! Thanks.

Barry Kline

kline@juncol.juniata.edu

SYLK DATA FORMAT SOUGHT

I need the specification for SYLK, which was used in the MS-Multiplan spreadsheet. I think it was described in the Multiplan user manuals. Does anyone have manuals they could check? TIA,

Russell Schulz

Russell_Schulz@alpha3.ersys.edmonton.ab.ca

TELEVIDEO: THAT DARN CAT

Ann Radnich, annr@scs.unr.edu, would like any available information on a Televideo desktop computer she inherited. We think it's a TeleCat, which was Televideo's 286 AT clone, but we can't find any docs for it. Anyone who has docs, e-mail please.

ARTICLES NOTED

"Let's Boot Up the Trash-80 and Play Some Oldies," by George Johnson, *New York Times* Week in Review section p. 2, Sunday, August 20, 1995. A summary of issues related to hardware collecting and software archiving, written with tongue firmly in cheek, but a welcome sight in the mainstream press even so. The author seems to think that software emulators, virtual antique stores, and problems of media deterioration belong to the dim, distant future.

PUBLICATIONS RECEIVED

Australian Computer Museum Society Newsletter.

#5, 8 June 1995. Committee news; Historically Significant Computers and Australian Top 10; Golden Oldies Profile; General Meeting notice. Unpaginated.

#6, 1 August 1995. Committee news; Memberships and New members; WWW sites; Book review; PDP-8 Story part 1 by Doug Jones. Unpaginated.

Charles Babbage Institute NEWSLETTER, Volume 17 Number 3, Spring 1995. Year of the Computer in MI; Fundraising campaign; Business History Conference; Bruemmer elected to SAA Council; Herbert Simon lecture; SHARE 40th anniversary. 6 pp. From Judy O'Neill.

Hewlett-Packard *Journal*, recognizing and publicizing technical contributions made by HP personnel.

Volume 46 Number 3, June 1995. Capillary electrophoresis instrumentation; Fault-tolerant mass storage; CASE tools for microprocessor design. 98 pp.

Volume 46 Number 4, August 1995. 100VG-AnyLAN; AccuPage 2.0; Large LCD flat-panel display; Economic modeling for programming; Benchmarking ASIC evaluation. 74 pp. From the editors.

HISTORY OF COMPUTING: An Encyclopedia of Computer History by Lexikon Services. Deluxe Edition, July 1995. Approximately 1,000 pages and 70 digitized photos; requires 5 Mb disk space, VGA graphics and MS-DOS 3.3 or better. US\$15.00. From Mark Greenia.

International Calculator Collector, Issue #9, Summer 1995. Dallas Show, Summit: A Man and an Idea, Pocket Calculators: The Early Years, HP Handheld Calculator and Computer Guide, classifieds, resources, more. 16 pp. US\$12 per year with membership (\$16 foreign). From Guy Ball.

Random Output, monthly newsletter of East Bay FOG.

Volume 11 Number 6, June 1995. Educational software; Using CAD; Annual elections. 4 pp.

Volume 11 Number 7, July 1995. CAD for Construction; Web browsing; Q&A. 4 pp.

Volume 11 Number 8, August 1995. BASIC through Windows; Passage to Vietnam; Taglines. 4 pp. From Pete Masterson.

The Z-Letter, newsletter of the CP/M and Z-System community.

Number 37, May/June 1995. TS-802H system tracks; Z-System on HP 125; reversing BACKSPACE and RUBOUT in DR CP/M; resources, publications, letters and classified. 20 pp.

Number 38, July/August 1995. Spellbinder soft keys for Televideo; introduction to Z-System shells; resources, publications, letters and classified. 20 pp.

US\$18 for 12 issues (2 years); Canada/Mexico, US\$22; International, US\$36. From David A. J. McGlone.

The ANALYTICAL ENGINE

Volume 3, Number 1, November 1995
ISSN 1071-6351

newsletter of the Computer History Association of California, is published four times a year — in February, May, August and November — at Palo Alto, California.

This magazine is available both on-line and on paper. Basic, domestic subscriptions are \$25 electronic and \$35 paper, with \$25 deductible as a charitable donation. For information on institutional, international, and low-income subscriptions, contact the Association at:

3375 Alma Street

Suite 263

Palo Alto, CA 94306-3518 USA

Internet: engine@chac.org

WWW: <http://www.chac.org/chac/>

ADDRESSES OF CORRESPONDING ORGANIZATIONS

Australian Computer Museum Society, PO Box 103, KILLARA 2071, NSW, Australia. Michael Chevallier, secretary.

Charles Babbage Institute, 103 Walter Library, 117 Pleasant Street SE, Minneapolis, MN 55455. Judy E. O'Neill, associate director.

The Computer Museum, 300 Congress Street, Boston MA 02210. Brent Sverdlhoff, curator of historical computing. Note change of contact.

East Bay FOG, c/o Pat Watters, 5497 Taft Avenue, Oakland CA 94618. Tom Lewis, president.

Hewlett-Packard *Journal*, Hewlett-Packard Company, Box 51827, Palo Alto CA 94303-0724. Richard P. Dolan, editor.

Historical Computer Society, 2962 Park Street, #1, Jacksonville FL 32205. historical@aol.com. David A. Greelish, director and editor.

International Association of Calculator Collectors, 14561 Livingston Street, Tustin CA 92680-2618. Guy Ball, Bruce L. Flamm, directors.

Lambda Software Publishing, 149 West Hilliard Lane, Eugene OR 97404. David A. J. McGlone, editor and publisher.

Lexikon Services, 3241 Boulder Creek Way, Antelope CA 95843. lexikon2@aol.com. Mark Greenia, director.

Perham Foundation, 101 First Street #394, Los Altos CA 94022. Don Koiijane, president.

Silicon Valley Engineering Council, 145 W. San Carlos Street, San Jose CA 95113-2006. Edwin V. El-Kareh, director.

Unusual Systems, 220 Samuel Street, Kitchener, Ontario N2H 1R6, Canada. Kevin Stumpf, president.

GUIDELINES FOR DISTRIBUTION

The ANALYTICAL ENGINE is intellectual shareware. Distribution of *complete, verbatim* copies through online posting, Internet mail or news, fax, postal service or photocopying is encouraged by the Computer History Association of California.

Excerpting or brief quotation for fair use, including review or example, is also permitted, with one exception: Any material copyright to or by a third party and reprinted in the ANALYTICAL ENGINE by permission shall not be used in another periodical or context, unless the permission of the copyright holder is separately secured for the new use.

Alterations, abridgments or hacks of the ANALYTICAL ENGINE which change the intent or meaning of original content; or which contrive to bring income to any person or organization other than the Computer History Association of California; or which contrive to injure the Computer History Association of California, its officers, contributors, volunteers or members; are PROHIBITED. Reproduction of the ANALYTICAL ENGINE without its subscription coupon is abridgment in this sense.

GUIDELINES FOR SUBMISSION

The ANALYTICAL ENGINE solicits manuscripts of 750 to 2500 words on the general topic of the history of computing in, or with significant reference to, the State of California. Articles should focus on one interesting or illuminating episode and should be written for a technically literate general audience. Submissions are welcome from both members and non-members of the CHAC. Article deadlines are: July 15 for the November issue, October 15 for the February issue, January 15 for the May issue, and April 15 for the August issue.

Each author may publish a maximum of one signed article per year. This restriction does not apply to letters, queries, book reviews or interviews. Thank you for cooperating to protect diversity of voices and topics. Previously published material will be republished only in clearly attributed quotations or citations; or when its publication in the ANALYTICAL ENGINE will

bring it to the attention of a significantly broader audience; or when the original publication is materially obsolete or inaccessible.

Decision of the editors is final but copyright of all published material will remain with the author.

The preferred document file format is Microsoft Word for DOS or Windows, but almost any DOS or Macintosh word processor file will be acceptable. Submit manuscripts on DOS 5.25" or 3.5", or Mac HD (1.4) diskettes. Alternatively, please send your article as ASCII or ISO Internet mail. Please avoid submitting on paper unless absolutely necessary.

NEXT ISSUE / COVER ART

As much SDS as we can cram in, more Felsenstein, and a review of David Packard's autobiography. Definitely an "ENGINE Plus." Don't miss it!

The cover: An early marketing shot of an HP 3000, *circa* 1973. From the HP Archives.

NINES-CARD

SOME BUGS TAKE A LONG TIME TO SURFACE...

by Tom Van Vleck
<http://www.best.com/~thvv/index.html>

In 1972, I wrote a program for Multics that created wall calendars. I wanted to find the date of Easter, so I went to the library, and discovered a simple algorithm for calculating when Easter came in about 7 steps, attributed to Gauss. I copied it, used it in my program, and tested a few years to make sure it worked. The program worked fine for years, and then I got a trouble report in late 1979, that it was off by a week for Easter 1980. Sure enough, it was wrong.

Dennis Capps kindly researched the fix, and included a page-long comment in the revised program, quoting an article from the *American Mathematical Monthly* that said that Gauss's algorithm failed occasionally: the first time since he wrote it was 1980.

What I learned was that others may get by for 400 years with trying a few test cases, but If *you* try it, you'll get caught.

CONTENTS

Editorial: X-VICTORY!!	1
LATE BUT GREAT (we hope)	2
CONVIVIAL CYBERNETIC DEVICES: An Interview with Lee Felsenstein (Part 1)	2
HP's Early Computers, Part Three:	
THE STRONGEST CASTLE: The Rise, Fall and Rise of the HP 3000, by Christopher S. Edler	16
IN MEMORIAM: JOHN V. ATANASOFF	44
IN MEMORIAM: J. PRESER ECKERT	44
IN MEMORIAM: GERARD SALTON	45
A WEB PAGE FOR THE CHAC.....	45
SVERDLOFF SUCCEEDS WALLACE AT COMPUTER MUSEUM.....	46
Tech Corner:	
MAKING NEW BATTERY PACKS FOR THE HP-35 CALCULATOR by Douglas W. Jones.....	46
SUPPORT FOUND FOR TANDY PORTABLES	48
Quick Take: SILVER ANNIVERSARY OF XEROX PARC.....	48
Quick Take: APPLE II RESOURCE LIST	49
Quick Take: AMIGA PREVAILS!	49
SPOTTER ALERT	49
SPOTTER FLASH.....	49
MONEY, THE WORLD, AND THE ORDERLY PROGRESS OF DAYS	49
Book Review: Max Maxfield's BEBOP TO THE BOOLEAN BOOGIE	50
ACQUISITIONS	51
LETTERS	52
QUERIES.....	53
ARTICLES NOTED	57
PUBLICATIONS RECEIVED	57
ADDRESSES OF CORRESPONDING ORGANIZATIONS	58
GUIDELINES FOR DISTRIBUTION	59
GUIDELINES FOR SUBMISSION	59
NEXT ISSUE / COVER ART	59
NINES-CARD by Tom Van Vleck	59

SUBSCRIBE!

US\$7.00

UK£5.00

12.50DM

¥1000



*We
did
it!!*

CHAC • 3375 Alma #263 • Palo Alto CA 94306 • FIRST CLASS!

The 930 lives!!

Thanks to tireless work, generous donations, team spirit, and brilliant advice — as well as total determination — the magnificent SDS 930 mainframe from Space Environment Center in Colorado now has a safe home in the Bay Area. Every admirer of legacy hardware can take pride in the rescue of this pristine, five-ton computer from looming destruction. See inside for details! And now....

Onward to the Museum!

The Computer History Association of California