

Sol: the inside story

By Lee Felsenstein
ROM magazine, July 1977

"I designed the Sol!"

These words are made to be spoken from a pinnacle of technical authority, preferably by a gimlet-eyed Herr Doktor who pursues exact solutions to the nineteenth decimal place and who reigns over a limitless sea of subordinates slaving away over rows of drafting boards. Or they could come from a furry little gopherlike creature with a piece of string for a belt who sleeps all day and occasionally surfaces to deposit a few dog-eared pages of scrawled diagrams with his custodians.

Since I fit neither of these descriptions, I hesitate to make that claim (except as part of a put-down), for it is only partly true. Besides, as I look over the reasons for making certain design decisions along the way, I am struck by the fact that most of those reasons had little to do with the ultimate advantages of the decisions.

The Sol, therefore, got designed partly by me, partly by Bob Marsh, and partly by chance and circumstance. My description of that process is intended to instill confidence in those who feel that there are great secrets involved in the design of products and that mastery of most of the universe is a prerequisite to successful design. It is also intended as a warning to those who think that the design process is deterministic and uncomplicated.

WHAT IS THIS THING CALLED SOL?

For the benefit of future historians, I shall state that the Sol is a single-board computer built around the 8080 microprocessor and the S-100 bus structure. It incorporates an integral video alphanumeric display circuit, serial and parallel interfaces, and random-access and read-only memory on the board along with an audio cassette tape interface. A keyboard plugs into a connector on the board, and a video signal comes off through a coaxial cable. Regulated DC power is supplied to the board through another push-on connector, and that's all that is needed to make it compute.

Sol's main feature is a 100-pin edge connector that provides all the signals of the S-100 bus to any number of memory, I/O, or other peripheral cards available from different sources. An important secondary feature is the "personality module," a tiny (3-by-1 1/2 -inch) printed circuit card on which sits the ROM. The personality module plugs into a small edge connector on the Sol board. By this means the personality of the Sol can easily be changed without technical skill. Also, all the serial, parallel, and audio connectors are mounted along one edge of the card, requiring no external harness or connector assembly.

The Sol-10 and the Sol-20 both put this card, along with a power supply, into a metal chassis with walnut sides and a typewriter-style keyboard. The Sol-20 has a five-slot "daughter board" that plugs into the 100-pin edge connector and itself provides five more similar connectors, so that cards of the S-100 type can be plugged in within the cabinet. It has additional power supply capacity to feed these extra cards.

BEFORE THE BEGINNING

In 1974 I was helping to run a public-access "computerized bulletin board" system (called Community Memory), which was essentially a labor of love for me and several other people. It had two terminals in public where people could come in off the street and enter information items as well as search for them. Since the terminals were unattended, hardware reliability was an obvious problem, especially when we postulated much larger systems.

My way out of this future problem was to design an all-purpose "convivial cybernetic device" as a terminal/concentrator/processor - in such a way that amateurs would be encouraged to get their hands on it. In theory, each place where one of these "Tom Swift Terminals" was installed would develop a computer club. Then, when a terminal broke down, relief would be a local matter, and people would not have to place their trust in a remote maintenance structure.

It was by placing a notice about this conceptual design on the system and inviting respondents to form a discussion group that I met Bob Marsh for the second time. (The first had been during our college days when we lived in the same co-op residence hall at Berkeley. We did not share many interests then.)

Now, however, Bob had raised himself to the state of an unemployed electronics engineer (self-taught) who had nearly won an encounter with a glorified version of Don Lancaster's TV Typewriter that he had built, improved, and fixed from scratch. After a while Bob suggested that I go in with him on the rental of a workshop. I agreed, and we signed a three-year lease on a garage in industrial Berkeley, commencing from January of 1975. I moved my workshop out of my living room, and Bob took up residence in the upstairs office, trying to find a product to manufacture. A plan to produce a limited-edition digital clock with a fancy wood case never materialized, which was a disappointment both for Bob and for Steve, a friend of his who did woodworking (we shall hear more of this later). Bob spent much time investigating the possibilities for a logic analyzer similar to others then available, doing much design before he gave it up as impractical.

In March 1975 I took Bob to the second meeting of the Homebrew Computer Club, where about twenty-five people stood around the first Altair 8800 to reach the area and watched it blink its front panel lights. That was all it could do, since it had no I/O circuitry and only 256 bytes of memory. There was a lot of empty space inside that cabinet.

Processor Technology was founded as a partnership in April of that year (I was not one of the partners) and began designing ROM, RAM, and I/O cards for the Altair. Incorporation

followed in July, and I began to, get bits and pieces of work - redrawing schematics, writing preliminary manuals, and other minor chores.

In July Bob finally got finally got to me. He had a proposition: he would pay me to design the video display section of the Tom Swift Terminal. In October the VDM-1 video display module first saw Revision A. The Tom Swift Terminal never did get designed exactly as I dreamed it. But, as the hordes of amateur computer enthusiasts swarmed out of the woodwork, its purpose was being fulfilled.

THE SOL SOLUTION

I like to say that, in the process leading to the birth of the Sol, Leslie Solomon, technical editor of *Popular Electronics* magazine, performed the act equivalent to that of the male. I still don't know who solicited whom, but Les agreed to carry a construction article on "an intelligent terminal" on the cover if a working model could be supplied in thirty days. This proposal was made in the middle of November 1975. I was summoned to Bob's office.

Bob tactfully asked me if I thought such a project was impossible. I reluctantly admitted that I did not but strongly suggested that an unintelligent terminal be designed, mainly because I wanted badly to exercise the features designed into the VDM-1 for such an eventuality. Besides, I had already turned down an offer from Bob to design an 8080 CPU board for Processor Technology.

Alternate designs were roughed out and prices compared. Both were for terminals using the basic circuitry of the VDM-1. His had an 8080 thrown in, mine had decoders and counters. Mine was cheaper, but only by about ten dollars. Gradually I came to realize that if the inevitable were to happen, it would be better to be on the inside than on the outside. After a day or two of discussions, I agreed to contract to do the design. I pulled out of a volunteer project on the grounds that "the roof was about to fall in again." And somewhere in the midst of all this I looked up at Bob and said: "Let's advertise it as having 'the wisdom of Solomon.'" From the comment came the name Sol, which is meant to be written in biblical-movie-poster letters chiseled out of stone. Les will never live it down.

Bob had already developed an architecture for the device. It involved taking the on-card memory and I/O devices, disconnecting the S-100 "DI" data input bus, and creating a separate low-drive "internal bus" that would get data to the CPU through a four-way data selector or multiplexer. One input of this would be the regular S-100 DI bus, another the internal bus, and the other two would be used for keyboard and parallel data inputs. The reasoning behind this was to minimize the need for high-power tri-state drivers on the card and to allow the low-power tri-state output features of the RAMS, ROMs, and UARTs to serve their intended functions.

I should explain here that designing for Bob Marsh can be somewhat of a trial. At least at that time, when he had little else with which to concern himself, he was continually turning up with new features and economies that he suddenly wanted incorporated in the design. He would explain the problem or opportunity and then preface his technical solution with an inevitable "All's ya got to do is" This would be forgivable if he were not so often right and

possessed of a truly useful and valuable idea. Were the designer a prima donna, the relationship would terminate after the second such incident, with the designer fuming about "professionalism" and "interference." Of course, since my workshop was in the same room as his, I could not have gotten very far if I had wanted to stamp out in a rage. The situation did, however, call heavily on my sense of futility, absurdity, and ultimate irrelevance.

A few days after I started on Sol, Bob had another idea. He wanted all clocks on the card derived from a single crystal. The VDM-1 had been using a 13.4784-Mhz crystal, and Bob had been doing some division problems with his calculator. It seemed to him that not one but three relevant clock speeds could be extracted from this frequency, allowing operation with 2.5-Mhz and 3-Mhz 8080 chips if desired. I protested, but without grounds, and very soon gave in and designed a simple clock generator which allowed (through a variable-modulo flip-tail ring counter, if you must know) the kind of clocks Bob wanted.

After a decent interval Bob informed me that he would like to increase the crystal frequency to 14.31818, which would be necessary if color video peripherals were to operate with it. Please to redesign all relevant circuits accordingly. I fumed and grumbled but found a way out eventually. After all, the clock frequency to the 8080 would be of a 488-nanosecond period, a whole 1.8 percent shy of the spec sheet minimum! Why worry?

The biggest upset actually had the flimsiest reason. Bob wanted to be able to extend the S-100 bus of the Sol to an external cabinet through a flat ribbon-style cable, which came in 50-conductor maximum widths. With my avid agreement, Bob wanted to ensure that there were plenty of ground return lines sprinkled through this cable - a practice that had been neglected on the Altair and that probably caused many of the noise problems associated with it. To do this would require more than 100 wires.

Bob and I did some figuring while driving down to a Homebrew Club meeting, and he reached the conclusion that it would be permissible to take the DI and DO buses and connect them in parallel, making one single DIO bus. We assumed that every manufacturer who was anybody was doing the same as Processor Technology, using the DBIN signal from the 8080 as an enabling signal for data to the DI bus. DBIN could therefore be used at the Sol as a "direction signal" for the DIO bus and data could be sent both ways on one set of wires. Think of all the cable you'd save!

As it turned out, those eight extra traces saved on the Sol board nearly made the difference between a buildable board and an impossible one. The fact that some manufacturers' (notably MITS's) boards did not use the DBIN signal and were therefore useless with Sol, unless modified, is still a source of defensiveness in discussions of the design. I take the position that Bob made me do it, and he takes the position that history will absolve him.

As the design progressed, we realized that we were building a general-purpose computer rather than just "an intelligent terminal," but the decision was made to soft-pedal the fact until the last possible moment. Once published, all the fuss possible was to be made about its general-purpose nature; but until it actually saw print, it was to be treated first as a terminal.

To jump ahead a bit, when I finally delivered the working prototype to Les Solomon's desk and pointed out its salient features, his eyebrows began twitching. Why couldn't he, he wanted to know, plug in a ROM board with BASIC burned in (as he could do with a Bytesaver) and run stand-alone? I smiled my blandest smile and muttered, "Beats me."

THE BIG PUSH

Originally I had been given to believe that I would be required by the terms of the contract only to provide a schematic diagram and to help interpret it for the benefit of the layout artist. When it came time to start laying tape, I discovered that I was expected to pitch in with a will and help stick down the many miles of crepe tape that would be required to produce the printed circuit artwork. Time was of the essence, as I agreed, and the layout artist who had been engaged was somewhat out of his depth with a 110-package layout to be done at four-to-one enlargement. Besides, he apparently didn't believe in making preliminary pencil sketches to test possible routings.

Steve the woodworker had built a large light table to Bob's specifications, and this was set up in the only available space in the now crowded garage-aloft above some offices that Processor had installed. I personally made sure that an electrical conduit running at forehead level was padded and tried to get someone to put up a fence to prevent one of us from stepping backwards off the edge of the loft. That, the knocking of heads against rafters, and the sticking or scratching of various parts of the head with needle-point X-acto knives, which we constantly had to hand, were the main hazards of working there.

In addition, there was the chill at night and the heat that accumulated during the day. People downstairs never had the same opinion about a comfortable temperature as we did. We soon settled into an insane schedule of fourteen- to seventeen-hour days, seven days a week. I kept going on orange juice, the younger layout man used Coke. He succumbed at the very end and I had to finish up alone. The task took almost three weeks of standing up looking down into a fluorescent-lit white background on which crawled worm tracks of black tape.

The crowning moment came when Bob changed his mind for the last time about which side he wanted the 100-pin connector to emerge from (this was determined by whether the 5-100 boards were to mount above or below the Sol board). Layout had already started, and fortunately that connector lay in the other guy's area of concern. He had to cross over about seventy-five traces in order to flip the connector, and he did not succeed. Some traces were brought out to dead-end pads with no hole through them, leaving some of the pins to be connected by tack-soldered jumper wires later during assembly.

I should point out that this printed circuit board was the prototype; no wire-wrap or other version had been made. Only two boards were ultimately produced (by a specialty shop run by a work addict), and one of them sits on my window ledge at home to remind me of how rough things can become. Besides, I like to point out my half of the layout and how obviously superior was my technique compared with the other half.

The other board I loaded with parts, tack-soldered with almost a hundred jumper wires out of sight on the rear side where traces were incomplete due to the hasty layout, checked

out, and presented to the software people. to see what they could do. I have no recollection whether that was before or after Christmas of that year; I think it was before.

IN THE BAG

Bob had laid out the basic outlines of the cabinet and had patched together a power supply and a sample keyboard. One of his design criteria from the absolute beginning was that the cabinet have walnut sides of a certain height or less. Steve had told him of a great bargain to be had from center-cut pieces of walnut, which were ordinarily almost thrown away. There is some cause to believe that the primary reason for the existence of the Sol was to provide an outlet for this inexpensive wood. At any rate, Steve is now in the wilds of Wisconsin running a mill to make walnut Sol sides.

As the project proceeded, Bob was undecided as to whether a cassette interface or drive would be included in the production version. As photographic deadlines approached, he decided to put a dummy cassette drive in the cabinet. This was to symbolize the fact that cassette tape would ultimately be available in one way or another - how he did not know. The text of the article, when it appeared, mentioned nothing about it, and the mysterious tape deck has proved to be the biggest single source of questions from the readership of that article.

The printed circuit board was ready forty-five days after the start of the project. About sixty days from the go-ahead, the cabinetry, power supply, and software were coming down the home stretch. Bob scheduled a night flight to New York, to save money, and informed me that I was coming; everybody went into a final homestretch panic lasting the better part of twenty-four hours.

We made the helicopter only because my watch was fast. Bob forgot to bring his tickets to the heliport and broke numerous traffic laws going back to get them. Carrying the Sol shrouded in two paper bags, we embarked for the big-time world of New York.

Of course we got no sleep on the plane and arrived an hour too early for our purposes. We went directly to Les Solomon's home and shared breakfast with him, discovered that the Sol wouldn't work (you could see it trying behind a veil of "snow"), then took the train to Les's offices on Park Avenue, where we met the boss and everyone else. Leaving behind a trail of excuses, we emplaned for Boston, where the folks from *Byte* magazine took us to dinner and drove us up to Peterborough, New Hampshire to see the snow. I conked out on this last trip, after attempting to stay awake by gorging myself at every available opportunity. I fear that I made a poor impression by so doing.

Back at the workbench a day or so later, I traced the trouble to a tiny speck of wire obviously loosed from the shield braid of the coaxial cable as it was stripped. This crumb had been captured underneath a socket, where it had plainly stayed during all the checkout without causing mischief. In transit it had shifted to the worst possible position - shorting two obscure but critical traces together. Murphy's Law confirmed!

After that it was bundle up again, this time myself alone - but carrying my oscilloscope and tool kit as well as the Sol. An hour or two of demonstration in the offices of *Popular*

Electronics convinced them that Sol would indeed work. Then it was back home directly, where I passed "Go" and collected the balance of my lump payment.

DEBUT

A great deal of work remained to be done after the prototype Sol was finished. Bob wasn't through with suggesting new things and better ways - I suspect he never will be. One major change in the design bears pointing out, though.

The personality module, which owes its name to Don Lancaster, came about as a result of unpredictable EPROM supplies. Some way had to be found to plug in three different EPROMs without taking up extra area. The first suggestion was for a "piggyback" card with upright pins that plugged onto a row of socket connectors on the board. I take credit for extending this idea to one of a miniature printed-circuit plug-in card with a row of edge-connector "fingers" that would mate with a right-angle PC connector and be guided by rails so it could be changed from outside the cabinet. This leads to fantasies of employees swapping in game personality modules during lunch time and returning to the business modules when the whistle blows.

Almost everything except the circuit was redesigned between prototype and production, and even the circuit was the object of much straightening out of wrinkles. The printed-circuit board was totally re-laid out, and the layout was digitized for computerized photo-plotting. More expediting, twenty-four-hour days, and dead runs from plane to plane.

In late August the Sol, in its production form, had its debut at the Atlantic City Personal Computing Show. I understand the cocktail party was splendid. Personally, I spent the time trying to trace down a bug that I was later told was in software.