PACKAGING ASPECTS OF THE UNIVAC 9200 DESIGN

W. F. Simon
Univac Division of Sperry Rand Corp.

ABSTRACT

The design of a small computer system for business data processing is as sensitive to the choice of packaging techniques as to the selection of memory cycle time or logic gate speed. The UNIVAC(R) 9200 was designed for volume production with complete freedom for innovation in all basic packaging parameters. The choices made and the technical justification of such choices will be explained.

There are two kinds of goals in the world of Digital Computers. One is to design a system so powerful that it can accomplish tasks which have thus far been impossible - or at least impractical - to complete. The other is to design a system so simple to build, use, and maintain that it becomes possible to extend the application of data processing to areas in which costs have previously been prohibitive.

This is not to imply that there isn't a large market between these extremes. In fact, very substantial income is derived from the existing market of experienced users, but the ability of a system to open new markets is the hallmark of success.

UNIVAC(R) has done well on both sides of this street, being fortunate to have the substantial resources needed to supply the full range of equipment needs and substantial software support as well. The challenge of building greater performance than any existing design can offer is the easiest to understand, and is limited only by the daring of the designer. The goal of the simple machine design is far less dramatic and yet requires a level of skill, insight, and innovation quite uncommon in the engineering profession.

The UNIVAC 9200 design represents a successful attack on the economic frontier of the computer world of the mid-1960's. It may seem clumsy in the LSI world of the 1970's, but it provides an excellent example of the importance of many small but related decisions.

In planning this system the basic dilemma was to provide a minimal system which could stand alone and yet retain the potential for growth to a performance level an order of magnitude greater. On one hand, nothing can be wasted in the minimal system; but on the other hand, each step upward to greater performance should be easy to make. Some considerable experience with this problem had taught us that both the minimal and the expanded design must be considered as a unified problem, and that success in designing a profitable minimal system would go a long way toward providing a highly competitive design throughout the desired range of performance.

What are the absolute minimum requirements for a viable business data processing system?

It must process data files. These can be perforated paper tape, punched cards, or magnetic tape. Since business applications are commonly associated with punched cards, the minimal system was oriented around a card reader and a card punch.

It must store and retrieve data and instructions rapidly in any arbitrary sequence. For acceptable performance, some minimum amount of random access storage capacity must be established.

It must prepare printed reports. The minimum system must therefore include a printer of some kind.

There must be some kind of control panel to initiate, monitor, and terminate operations.

And, of course, the minimal system must include the arithmetic and control circuits of the central processor.

Although the design and packaging of the mechanisms used for input and output of digital data represent a substantial portion of the 9200 System Design, we will pass over these problems with only a few general comments.

1. In most cases, mechanisms are housed in cabinets which are quite wasteful of useful volume.

2. Signals required to control mechanisms are commonly large in voltage and current levels; the responses returned from them are usually small, weak, and noisy.

3. Mechanical maintenance accounts for the lion's share of the field engineer's work on a small computer system. This favors the man who is a mechanical genius but a little fuzzy on electronic and logical concepts. When a failure occurs in the electronic segment of the system, it will take a long time to isolate and repair unless considerable planning has gone into the preparation of maintenance aids and fault isolation procedures.

These considerations had considerable impact on the basic decisions concerning 9200 electronic components. In the first place, we wanted a small system that looked like a small system. The considerations of operator convenience and mechanism accessibility made it clear that a minimum system would occupy a volume which could easily swallow a substantial segment of the electronics. We could avoid the over-shadowing bulk of a five foot high "electronics cabinet" and simplify intercabinet cabling if we made efficient use of the volume already available under the 40 inch high "work surfaces" of the various mechanisms.

This meant, however, that basic logic elements would have to co-exist with high powered solenoids. It also meant that volumetric efficiency of electronic packaging would be important, and that the problem of cooling the logic gates would suffer from the restricted space.

An even more serious consideration was the realization that efficient use of space would eliminate the use of duplicated printed circuit assemblies in the logic module. In "second generation" computers such as IBM 1401 or UNIVAC 1004 a few basic cards were used repetitively to supply most logic functions. When one card was suspected of malfunction, another could be found in an operating logic section for a comparison test. If interchanging the cards moved the malfunction, the fault was clearly isolated. This kind of trouble shooting requires a minimum of skill, which is an advantage more significant than other more obvious factors, such as reduction of spare parts inventory and increase of manufacturing volume on favored assemblies.

The difficulty is that it requires nearly all logic connections to be made via the card connector, so that a single component assembly may perform a wide variety of functions. Thus, any increase in volumetric efficiency of logic packaging would be limited to the practical improvement possible in connector pin density. At most, the number of connector pins per unit area of backpanel had only increased by a factor of four while electronics shrunk from discrete component technology to integrated circuit technology.

UNIVAC(R) decided in favor of higher density packaging, with increased support in the development of maintenance techniques to compensate for the loss of package duplication. The 9200 logic averages only one duplicated card in every seven card types used for basic logic (excluding such things as solenoid drivers, where power density forces lower component density). But the system is also supplied with a fault dictionary and a detailed sequence of fault isolation instructions. By following this carefully chosen sequence of simple operations, and recording the resulting digital codes (usually displayed on the operator's panel), a fault symptom is generated. These symptoms are sequentially listed in the dictionary, and the malfunction is identified with a printed circuit card location.

Obviously such a technique is aided by functionally grouped logic elements on each card. It is easier to identify an adder error and replace the entire block of adder logic than to identify the specific gate at fault. This is consistent with the goal of increasing logic density, which also profits from functional grouping of components.

The card connector chosen for UNIVAC 9200 printed circuit assemblies is a design originated for use in 1108. It has 55 terminals in an area of 1½ square inches. To keep the insertion and withdrawal forces low without loss of contact pressure we literally put "a square pin in a round hole". The "hole" is a split sleeve formed in a progressive die, which gives two independent spring elements shaped to make four lines of contact with the corners of the square pin. The sharp corners of the square pins create a large normal force on a small area, assuring a gas-tight seal without excessive friction forces.

The square pins simply continue through the connector block and provide appropriate terminals for wire wrap on the opposite side. A simpler male connector pin is difficult to imagine. The female half of the connector "floats" in a molded cavity for easy connector mating.

It is provided with a tang which bends down and solders in the printed circuit board at a distance great enough not to restrict the contact "float". This part is somewhat more expensive, but provides extremely high reliability.

The choice of a basic logic gate design for the 9200 was not an easy decision. Up to 1965, monolythic integrated circuits had not been cost competitive with hybrid chip technology or even with the use of integrated resistor networks to couple discrete semiconductor elements. An extensive survey of the semiconductor vendors surprised us by showing a significant cost advantage in most monolythic gate structures.

The only structure for which a monolythic version was distinctly poor in both cost and performance was a simple diode gate, without gain. The 14 lead integrated circuit packages had an obvious analogy to the 16 or 18 terminal printed circuit cards of the computer families popular in the early 1960's. We knew, therefore, that a simpler set of basic logic functions could be designed with two levels of gating - the first level having only diode gates, and the second level providing a complementary gate function plus power gain and signal inversion. Using this philosophy, we defined a set of integrated circuit NOR gates which could be coupled with simple diode AND gates.

In order to be sure that this unorthodox approach was justified, an existing processor design was studied to establish a basis for comparison of the different logic families. The original design had been expressed in two level logic, suitable for the diode-coupled NOR circuits which we proposed. To obtain a fair comparison we re-designed a segment of approximately 3000 to 4000 gates to take advantage of Emitter Coupled Logic which was available "off-the-shelf". When this proved unattractive, another attempt was made, using Complementary Transistor Logic features - another "off-the-shelf" item. Equivalent counts of components in each of the logic families were determined and comparisons of cost and performance were made.

Both ECL and CTL suffered from incomplete use of the package terminals. Although each circuit was available in a 14 lead package, as few as six leads were actually used for logic interconnections on some configurations. This was somewhat a function of excess dissipation, but also reflected the use of extra voltage and ground pins. The net result was less function per package, and hence more packages required per system. This more than offset the lower cost of a standard package as compared to a custom design.

The use of diode coupling between NOR circuits would not have so easy a time today, since the value of an integrated circuit terminal is generally appreciated. Nonetheless, the 9200 logic elements retain significant advantages. On the average, each integrated circuit has 17 diodes associated with its input terminals. If the typical chip has three outputs, we have the equivalent of an integrated circuit with 20 terminals instead of 12. In addition, a great deal of flexibility is obtained since the diodes are connected only where needed. Many different functions can be accomplished with relatively few kinds of parts.

Note also that the diode gates dissipate very little power. Thus the complex 20 terminal logic function has virtually the same power dissipation as the simpler 12 terminal configurations commonly used as an integrated circuit building block. Diode coupling has not only reduced the average cost of the logic functions; it has also made it easier to cool.

Provision of adequate cooling also has a bearing on the general layout of the 9200 System. A small system should not require too many forced air paths. Blowers are part of that "irreducible minimum" of cost which every system must contain before the first useful function is installed. As the system grows smaller, this base becomes more important. To keep the number of air paths at a minimum, all printed circuit assemblies were placed in one of two card modules - the memory module, or the central processor module. Thus, only two blowers are used in the minimum system. (I am excluding the small axial flow fans which are an integral part of the power supplies, and a small centrifugal fan which cools the printer actuators, because these requirements are independent of circuit packaging).

Why not go all the way, and merge the memory module with the processor module? This was considered, but was rejected as too dangerous. The 9200 was the first commercial use of a plated wire memory, and fully enough problems were anticipated without adding cross-talk from the high powered circuits associated with the printer and the punched card equipment. In fact, we felt quite brave in merging logic circuits with the I/O control circuits. To make this practical, we specified our custom NOR circuits with an extra measure of noise margin.

Where the usual DTL or TTL circuit must overcome two silicon diode thresholds to turn on, the 9200 logic element must overcome four diode thresholds. One of the extra diode thresholds is wiped out by the forward drop of the coupling diode, but the remaining noise margin permits a noise spike approximately 3/4 of a volt greater than the worst case noise limit of the more conventional circuit configurations.

This would not be practical if we did not move the power supply side of the logic waveform as well. Use of a six volt power supply instead of the usual five volt source increases dissipation somewhat, but not to a level which will offset the low power dissipation of the diode gate structures. In fact, the internal circuitry of the chip could be designed with current sources whose variation between OFF and ON conditions is less, since a larger source voltage is available.

Experience and measured tests have shown that the 9200 logic is even less susceptible to interference from external noise sources than we had expected. A standard signal generator with a 41 inch whip antenna three feet from the cabinet could not deliver enough power to disturb it. Severe line voltage transients and conducted voltage pulses applied to the cabinet structure caused no errors. This is scarcely surprising when one discovers that printer actuator drive circuits supply 170 volt pulses in the same wire-wrapped backpanel used by logic signal connections. As might be expected, some considerable care of wire routing in the vicinity of these high powered circuits, and careful attention to ground return paths were necessary to accomodate this circumstance.

Having chosen to package all electronics except the memory in a single area, there still remained a choice of selecting which area. The printer thrust itself upon us. It requires a heavy bundle of wires to couple actuators with their driving circuits. Up to a reasonable distance, these cables can be twisted pairs; beyond this coaxial construction would be needed. The cost increment would be dramatic. Therefore, the central processor found its home inside the printer cabinet.

Merged with the central processor card library we now find the printer control logic and associated drivers, sense circuits for paper motion and character sequence, and also its cable connections. Rather than wire a rack-and-panel type connector to the card library, we simply solder all cable connections on a printed circuit board and plug it directly into the wire-wrapped back panel.

Similarly, the card reader logic circuits and the card punch controls, with their associated photocell signal amplifiers and power drive circuits are part of the main card library. Cables branch out from the card insertion side of this module to connect the mechanisms. Although separate cabinets house the reader and punch mechanically, very little of the associated electronics can be found in these cabinets. Hence, no cooling is required, and no separate frame to mount replaceable cards is needed.

Substantial cable lengths are used, but these are made tolerable by providing that all cabinets bolt together in an L-shaped floor plan. Cables simply pass between such cabinets through holes in the adjacent sides. A continuous cabinet frame, with all sections electrically connected and grounded, provides a crude electrical shield. Where signals are small, separate cable shields and ground return paths (tied to frame ground at only one point in the system) maintain a low noise level.

At the corner of the L, there is a cabinet whose accessibility is rather poor. From the operator's station, only a six inch segment is exposed. At the rear and on the right side there are removable panels, but if one is chosen as a means of access, the other becomes somewhat redundant. A good part of it faces the side of a module whose front is the maintainable part.

In this cabinet we chose to mount the plated wire memory. It is accessed from the rear of the installation and this is the prime access to the cabinet. On the side, however, one has limited access to the volume behind the memory module. At this point we mounted the main system power supply. It requires little attention, and it is acceptable to slide it completely out of the cabinet for service. Above this supply the controls for power sequencing and protection are mounted. Access is provided by lifting the table top, except that the circuit breakers may be operated from the side.

Grouping the memory and the power supply in a common frame was a fortunate choice. Power requirements for the memory are more complex and more critical than for any other need. When memory expansion requires a duplicate module to be added, it is simple to roll up a second cabinet almost identical to the first, with the same simple power supply and a "slave" power sequence unit. Some power supply capacity is wasted, but not enough to justify a second design.

The UNIVAC 9200 processor was designed to take full advantage of the high speed and low cost of the plated wire memory. This was done primarily by letting a small part of the memory serve as the working registers. In the usual machine organization, these registers are flip-flop aggregates and comprise a significant portion of the processor cost. Not all working registers can be included in the memory; for example, the memory address register must be external in order to access memory. However, the 9200/9300 can carry this concept further than other systems while retaining a significant speed advantage because of the memory speed. Only 29 bits of flip-flop registers were required for the arithmetic and control portions of the processor. A total of 144 bits of control register storage are contained in the memory. In addition, a dual set of arithmetic and indexing registers are contained in the memory, providing a zero interrupt acceptance time. This feature is important in achieving maximal throughput rates since interrupt acceptance time effectively decreases I/O device speeds. This feature is believed to be unique for machines of this low price class.

To meet the cost goals of the 9200 system design a substantial reduction in manufacturing cost of the control panel was necessary. This was achieved by a unique combination of permanent magnets and reed switches which replaced the usual mechanical control panel switches.

Commercial computer control panels have traditionally been built with mechanical switches having either directly exposed contacts or, at best, covered for dust protection. Airborne contaminants, such as solder flux vapor, ultra fine particulate matter, and the normal atmospheric content of sulphur and oxygen, have all served to lower contact life and necessitate costly replacement or mainenance procedures.

Where neatness was important, the wiring connecting the switches had to be laboriously measured and formed into tied cables. Subsequent stripping and tinning of the wire ends preceded the often delicate job of hand soldering the wires to the many switch contact lugs. Errors in making the cables and marking its wires, mistakes made in wiring to the wrong contact lugs, poor soldering which led to cold joints, shorted switch contacts, and burnt insulation were common difficulties. These recognized manufacturing faults were controlled by setting up, at significant cost, quality control monitoring, testing, and rework procedures.

At the start of the 9200/9300 development program, it was proposed that glass reed capsules could produce a control panel that would eliminate many of the faults of the designs then in use. The immediate and most obvious advantages were:

1. The glass reed capsules themselves would be immune to atmospheric-borne contact contaminants. This in itself would completely eliminate any necessity for field procedures pertaining to contact cleaning.

2. Mounting the capsules as components on printed wiring boards would allow conventional dip soldering to be used to eliminate the delicate hand soldering as well as the cable making, stripping, tinning, and securing operations.

Briefly, the switch is formed as follows: A rocker-arm containing a permanent magnet is inserted into a plastic holding cavity which in turn is inserted into a control panel opening where it snaps in place. The rocker-arm is held and allowed to pivot in the cavity by a coil spring and either one of two spring retainers, the choice of retainer being based upon whether "momentary" or "alternate action" operation is desired. Depressing one end of the rocker-arm moves the magnet over the overlapped ends of the reeds, causing them to make contact. The reed contact is broken by depressing the opposite end of the rocker-arm, which moves the magnet away from the reeds.

Several major design efforts were necessary before a reliable magnetic switch could be produced. One of these efforts was to determine the aging characteristics of permanent magnets to ensure long term reliability. The resulting study was so thorough and unique that the magnet vendor requested permission to include the test results in a textbook being written on permanent magnet theory. Other design efforts were necessary to determine the proper mounting angle for the magnets in the rocker-arm that maximizes the manufacturing tolerances, and to incorporate a mechanical and magnetic "snap" action feel essential for operated switches.

This work has provided all the hoped for advantages. Manufacturing cost of a 9200 Control Panel is about 1/5 the cost of similar panels in earlier UNIVAC systems. Maintenance is simple and infrequent.

The plastic rocker arm seems more likely to fail than the reed capsule. In either case, repair is the replacement of a component; no adjustments are needed.

It may be that these details are rather far afield from conventional concepts of packaging design. Perhaps it does make a point, however, that bigger and faster machines do not have all the fun. Bringing the computer to the small business man can be challenging, and the audience is both large and enthusiastic.