07/26/95

HEWLETT PACKARD

Visit date: Sept. 13, 14, 1994

Host: Peter Zivkov

Background

Hewlett Packard (HP) is one of the largest makers of small computers and instruments in the US. It began about 60 years ago as an analog instrument company, became a digital computer company in the 1970s, and today is largely a maker of small to medium size computers, microprocessors, laser and ink jet printers, and software. Instruments are now a tiny fraction of its sales. (See Table 1.)

H-P'S BUSINESSES	1994 REVENUES
COMPUTER-RELATED	19632
ELECTRONIC TEST	2722
MEDICAL ELECTRONICS	1141
ANALYTICAL INSTRUMENTS	754
ELECTRONIC COMPONENTS	742

Table 1. Distribution of HP's Revenues in 1994.[1] (Revenues in 000,000's)

Since 1990, net revenues have risen from $13,233 million to $24,991 million while the total number of employees has fluctuated narrowly between 89,000 and 98,400. Net revenue per employee has therefore nearly doubled while, as we all know, the cost of computer-related equipment has fallen.

A number of factors may be cited for the increase in revenue per employee, several of which are discussed in more detail below:

- More efficient design processes, use of multi-function teams, and better computer design tools, including direct engineering aids and wide use of networking

- More reliance on suppliers to help design as well as manufacture components and subassemblies, including direct electronic linkages and use of the Internet for supplier interaction

---

[1] HP Annual Report for 1994.

An important purpose of this visit was to discuss both mechanical design and VLSI design with experts in a company that does both very well. Important issues include the state of computer tools for component and system design and analysis, and the ability to partition designs and share design and manufacturing with other companies. Another purpose was to discuss partnering with suppliers and strategic partners in fast-moving industries like microprocessors and disk drives. Important issues include sharing design requirements and rationale while preserving proprietary knowledge and retaining core competencies.

The Kittyhawk Disk Drive

The Kittyhawk drive was not a central focus of the visits reported here but was evolving at the same time. Information about it was obtained from other sources at HP.[2] It represents a number of the issues noted above and may serve as an introduction to the rest of the report. It was co-designed by HP and a number of key suppliers. A few years before, HP had decided that it could not afford to keep up technologically with all the elements of disk drive technology, so it developed supplier partners. HP set the main requirements for size, data storage capacity, and power consumption. These requirements in turn set the overall requirements on the main subassemblies, including a number of critical tolerance chains.

A disk drive contains several distinct and largely independent subassemblies, each of which operates in mostly well-understood ways with relatively mature technologies. Steady advances have occurred in each area (magnetic media, heads, head clearance over the platter, and spindle bearings) but radically new technologies have not been introduced. Because the main elements operate fairly independently and their interrelations and interfaces can be described fairly cleanly and precisely, it is relatively easy to divide up the detailed design and manufacture. The important system engineering skill is to determine the net tolerance required to meet the data storage requirements and then to allocate those tolerances to the separate elements and their suppliers. The remaining challenge is to find skilled suppliers for these high precision elements, especially the spindle motor with its micro-inch range runout requirements and the head suspension arm with its delicate balance requirements.

The main tolerance chain loops through the read / write head, suspension arm, main foundation, drive spindle bearings, disk shaft, and disk platter. This loop determines the drive's data capacity, which is based on track width and spacing within a given platter diameter. Spindle bearing non-repeatable runout is one of the key errors to keep under control, but

---

[2] A more detailed description of this product and its procurement may be found in [Whitney, Nevins, De Fazio and Gustavson].

others in this chain are important as well, including servo dynamics for controlling head position.

Once HP had determined the elements of these tolerance chains, it contracted out the design of key elements, namely the spindle motor, disk platters, read/write heads, head suspension arms, and electronic chips. HP designed the disk operating system software and the read/write control system. Citizen Watch in Japan was selected to manufacture the drive because of its ability to ramp up rapidly, its low margins, and its ability to manage the manufacture of complex precision equipment. The practical effect is that each element in the chain was made or assembled by a different company.

I have asked several people inside and outside of HP what would prevent Citizen from assembling the same set of suppliers and going into the disk drive business itself. The answer heard most frequently from insiders was that HP's skill lies in the ability to look ahead two years and see what price and performance the market will demand, and then to convert those requirements into engineering specifications. This is basically requires two skills: market knowledge and system engineering. In the case of the drive, the issues were size, data capacity, and price. Size and price are being driven down while capacity is being driven up. Within the given technology of magnetic moving disk storage, the only way to raise capacity and reduce size is to make the heads, tracks, and track spacing smaller and make the platters thinner so that more could be used in a given overall height. HP's strategy was thus to control creation of the specifications and rely on suppliers to deliver most of them.

HP people cite the laser printer as another example. HP depends on Canon for laser engines but has nearly driven Canon from the market for finished printers.

When I first heard of the Kittyhawk in early 1993 it was just coming onto the market. As of this writing, it has been withdrawn. It was known at the time to be too expensive for laptop computers and was targeted at higher priced products like business printers. Apparently HP has decided to exit this product line entirely.

This product therefore raises a number of fascinating questions:

- Who really has (had) the knowledge necessary to design this product, especially given the fact that no single company could make all of the critical components?

- Of all the kinds of knowledge required - market prediction, conversion of customer needs into engineering specifications, precision mechanical design and fabrication, electro-magnetic modeling, ultra miniature electronics, precise servo controls, operating and control software, mechanical assembly, and product integration and test - which are the ones needed by HP and which could safely be allowed to wither inside and be obtained from outside companies?

- What made this product susceptible to subdivision into conclusively defined elements that could be made by separate companies and still be integrated by yet another company in spite of a high degree of required precision?

- What kinds of product definition methods, system engineering, interface definition, CAD tools, electronic data interchange, supply chain management, tolerance chain definition (trading off the capabilities, cost, and development time of each supplier), tolerance chain oversight (seeing that each supplier delivers on its part of the tolerance chain) are needed in an age when products are increasingly being designed and made by teams of companies? Which of these companies needs which of these skills?

- Who is in the driver's seat - the company that can conceive the product or the company(ies) that can manufacture, assemble, test, and deliver at high quality and for low cost? Or are they in a state of inescapable mutual dependence?[3]

- What gives a company the capability to understand the market and develop complex electro-mechanical-optical-software products: those with their own ability to design all the elements; those that develop their own computer tools that suit their product development methods; those that retain the ability to manufacture the elements so that they can evaluate suppliers if they want ultimately to buy the elements; those that can make their own critical manufacturing equipment, tooling, and molds?[4]

---

[3] See [Liker et al] for a discussion of supplier relations which frankly recommends embracing mutual dependence as a fact of life.

[4] A purely mechanical example is provided by automobile half shafts which connect the transmission to the wheels in front wheel drive cars. Half shafts are highly stressed, safety-critical assemblies which contain highly engineered and precisely made constant velocity (CV) joints. A major US car firm makes half shafts for itself plus US and Japanese competitors and notes differences in the degree of oversight each such customer provides. The US customer provides basic requirements like torque and mechanical interfaces. The Japanese customer provides as much and more: a highly detailed set of test and evaluation specifications that the design must pass. A half shaft design engineer said of the Japanese customer, "They wouldn't dream of telling us how to design the CV joints but they will become very 'helpful' if our design fails any of the tests. They will lead us to find the answer they already know is right. They know because they make similar shafts themselves and are among the best in the world." In other words, both US and Japanese customers

A pattern emerging in the last decade is that Japanese companies are now the world's great manufacturers but American companies are the world's great product innovators. In fact, one can read of Japanese who worry that Japan will simply become America's factory while America retains the creativity and market savvy to set the specifications and design new products. This is quite the reverse of America's worry that it will lose the ability to manufacture. In addition, America has shown that it can learn many of Japan's manufacturing methods. Many Japanese companies are superb product innovators as well, so it remains to be seen what will happen.

Discussion with Peter Zivkov

Zivkov leads a group of mechanical engineer-programmers whose mission is to improve HP's capability in mechanical design. He has formulated a long term strategy that is built on use of HP's 3D solid modeler called ME 30 and its successor SolidDesigner plus systematic development or adoption of simulation tools. Of 2500 mechanical CAD users in the company, 75% use ME 30 while the rest use the predecessor drafting software called ME 10. (See below for a discussion of the differences between SolidDesigner and ProEngineer.) He estimates that every one of these designers has his/her own workstation, made by HP. These workstations and other computers are linked into a huge internal internet with about 90,000 nodes. The kinds of products designed by these people include workstations, medical instruments, test equipment, disk drives, and other similar electro-mechanical-electronic items.

New Design Process

In the last few years, HP has been moving toward cross-functional team design and concurrent engineering. The design process steps are

- concept design
- a concurrent process of detail design/simulation/validation
- documentation and release to manufacturing

Team membership is broad, including mechanical engineers, electrical engineers, manufacturing people, and experts from safety, procurement, and marketing. These people are not necessarily co-located, and both design and manufacture are distributed. The process depends critically on e-mail, teleconferencing and other networking facilities, design data management, a common database, and common CAD tools.

"depend" on their half shaft supplier, but the US customer depends in an entirely different way.

In this mix of skills and tasks, it appears that the MEs have emerged as the system integrators, while the EEs tend to divide into their own critical specialties, such as logic, circuit boards, VLSI, electromagnetic interference, etc. The MEs could just as easily have subdivided into chassies, packaging, thermal management, and so on, but at least some have not. However, the trend appears to have been spontaneous, inasmuch as there is no officially designated system integrator job classification. Nominally, the program manager of the design project is the system integrator, but MEs appear to have taken over technology integration during design.

This is interesting because it is rare to find MEs in this role, certainly in the US and Europe. EEs move into this role more naturally in the US, at least in part because their education has a higher system content. MEs tend to focus on components until late in their education and similarly keep a component focus in their careers, mainly because mechanical components often require redesign and integration with each new design. In VLSI, components are usually created at a separate time and place by different people, leaving product designers with a clear track to focus on systems. This issue is discussed more below in connection with the differences between mechanical and VLSI design, as well as in [Whitney, Nevins, De Fazio, and Gustavson].

Development and Use of Design Tools

HP's team-oriented design process has evolved from a serial process dominated by repeated use of prototypes into a more integrated process that makes more use of computational prototypes. It begins with a more thorough market analysis than in the past, complete with sales projections, technology surveys, and planning. The project manager utilizes a tool called "the return map" to estimate the total cost of product development and the completion date. A product is deemed successful if it earns back its development cost in sales relatively quickly.[5]

Once the basic ideas for a product are set, a number of concepts are evaluated quickly. Three dimensional models are the key to the new process, and CAE tools for many analyses are available. These include visualization of the exterior, internal arrangements, assembly and DFA, thermal hot spots, stress, wire routing, and rapid prototyping. Less amenable to simulation but still attempted are shock, vibration, and EMI.

A new tool called Sheet Advisor has become available to help design bent sheet metal structures such as cabinets and chassies. It contains a knowledge base and can predict cost based on manufacturing process as well

---

[5] See [House and Price] House, Charles H. and Price, Raymond L., "The Return Map: Tracking Product Teams," Harvard Business Review, Jan.-Feb., 1991, pp. 92-101

as what shop or vendor is being considered. It was developed in-house using internal manufacturing knowledge and is now sold commercially.

A key element in Zivkov's activity is to promote mechanical engineering design using such tools. They improve design and raise the professional level of the engineers. Besides, the tools are economical to use compared to physical prototypes and are developed in-house, which reduces their acquisition cost. The workstations and basic CAD software are economical, too, since they too are an internal product. Once everyone uses the same design tools, there will be a high degree of integration and uniformity in the process and its results.

The speed of the process depends on intensive use of electronic communication plus delay or even elimination of traditional documentation. Traditional detailed 2D drawings are not made. (Figure 1 is an example. Note the clutter.) Instead, simplified isometrics or perspectives are made of each part or assembly, and only the key tolerances ("critical to function" or CTF) are indicated explicitly. (Figure 2 is an example. Compare the clarity with Figure 1.) For the majority of dimensions, a blanket tolerance is defined, based pragmatically on what the supplier can deliver. Many of the parts are molded or machined based on NC data taken directly from the computer models.

See Figure 1. Design for A "Center Support" Portrayed in the Traditional Dimensioned/Toleranced Form (Used by permission.) This kind of drawing is used only for rare cases where drawings like Figure 2 cannot be used.

See Figure 2. The "Center Support" Portrayed via the CTF Method. (Used by permission.)

All tolerance judgments are made by engineers, who make these drawings themselves. There are no draftsmen. There is a guidebook two inches thick to help engineers choose tolerances. It is becoming available via MOSAIC to HP insiders and selected suppliers. HP has tried to use VSA, a commercial software package that uses Monte Carlo methods to predict tolerance stackups. Results have been mixed, since there is still a tendency to use tolerance analysis reactively when parts don't fit. The discipline to define tolerances early in the design, as well as to strive for robust designs, has not yet taken firm hold. But VSA is said to be hard to use, especially on 3D stacks that include angular errors. Yet these predominate over simple 2D stacks, so a good method and a robust design process that leverages tolerance analysis is not yet in place. When thorough tolerancing is done, a project is very likely to be successful, and vice versa.

Suppliers are increasingly being integrated electronically into HP's design and manufacturing process. Most have a PC and an NC machine tool. HP has an electronic design exchange (EDX) standard in place as well as ISO standard EDI (electronic data interchange) for business data. As of February, 1995, over 90 suppliers are tied in directly via EDX. From inside HP a design package can be assembled, encrypted, and sent out via FTP. Even though many nodes might see the package, only those encrypted to receive it will be able to. This process goes most smoothly when the supplier also has ME 30, but so far there is no plan to equip all suppliers with it. IGES is sometimes used as a substitute but it is not entirely satisfactory.

What Skills Does HP Need In-house?

HP is still exploring how far it can go with farming out design and manufacture. In some cases it may have already gone too far. Opinions are evolving regarding what are the key disciplines to have in-house as well as when to use suppliers.

Different divisions in the company make different choices regarding what to give suppliers. In some cases low volume specialized manufacturing is done inside while high volume manufacturing is contracted out. In other cases, fabrication is contracted out while assembly is done inside. No clear pattern has emerged, but suppliers are generally an increasing presence in both design and manufacture.

Six or seven years ago I was told by the then Director of Corporate Education, Marvin Patterson, that top management, being from the analog circuit design era, needed to understand that software was a profession. Since that time, computer products have become the main source of profit, so software-oriented management has emerged. Most engineers, regardless of base discipline, write or use software some of the time, and programmers are by far the majority of engineers.

Patterson also felt that the time would come when all engineers would be free agents, part of what the Agile Manufacturing community calls the Virtual Enterprise. They would coalesce around a project for the time necessary to carry it out, and then dissipate. Zivkov feels this would work satisfactorily for routine things, such as updating designs. "But HP should not do much routine stuff."

Long Term CAD/CAM Strategy

Today, HP's mechanical design is characterized by the following:

- All new products are designed with 3D solids. SolidDesigner is replacing ME30 for CAD and CAM.
- It is commonplace to use stereolithography.

- Other kinds of prototyping are done without recourse to 2D drawings. For example, Sheet Advisor is used for sheet metal part design and fabrication.
- Manufacturing release is in the form of 3D models rather than detailed drawings.
- Standard simulation tools are in wide use (Rasna, VSA, Flotherm...) as well as internally developed ones.
- EDX is used with both internal and outside shops.
- The Internet and other electronic communication means (WWW, 3D Model Library, SharedX, and News/Notes) are commonly used.
- A divisional Product Data Model is in use.

The intermediate goals (one year from achieving the above) are:

- Design intent will be captured in 3D models, and conventional drafting will be obsolete.
- Product data management systems will be in use company-wide in R&D and manufacturing.
- A library of useful 3D models of purchased parts (fans, connectors, fasteners, etc.) will be purchased from suppliers and made available for all designers to draw from.
- Use of Internet-based commerce will begin.
- Design intent will begin to be attached to 3D models, replacing the CTF drawings.
- Associativity will be implemented one way from 3D models to 2D drawings.[6]
- ME and EE design processes will be linked more tightly, and more automatic and direct links would be established between CAD and engineering analysis tools.
- Transparent links will be established between CAD and CAE tools.
- CAD/CAM will be linked directly to inspection tools, processes, and data.
- Design advisor tools will begin to be used.

HP's long term goals are

- Engineering infomation and services will be accessed through the Internet.
- Simulation tools will be fully integrated into the product life-cycle
- There will be a simple, smooth paperless information flow throughout design and manufacturing.
- Design, manufacturing, and procurement systems will be linked.
- Geographically distributed virtual work teams will be in place.

---

[6] Associativity means that when the 3D model is changed, the corresponding 2D drawing changes automatically. ProEngineer pioneered associativity in CAD.

- Virtual reality applications will be available for MEs.

Comparison of SolidDesigner and ProEngineer[7]

ProEngineer has become established as the fastest-growing 3D CAD system in the US, with good growth in Europe and Asia as well. Its main claims to fame are parameters, associativity, and constraints. Parameters are numerical values, usually dimensions, that the designer may attach to portions of the model, say the radius of a hole or its distance to an edge. Associativity means that different models, such as 2D drawings, derived from a parent design will update automatically when the parent is updated. Constraints permit the designer to relate two or more parameters, for example requiring that the distance of the hole to the edge must always be twice the radius of the hole.[8]

Parameters and constraints permit a designer to impose intent on the design. Associativity enforces that intent on later phases of the design that utilize derived models such as NC programs or 2D drawings. This sounds like a good idea but it has serious drawbacks. Parameters and constraints can build up into what used to be called spaghetti code in software. That is, things are linked to each other in hidden chains. These chains build up gradually as design proceeds. However, there is presently no convenient way to view them. When someone tries to modify a design, these hidden changes can be activated, causing other changes to ripple unpredictably through the design and into 2D models linked by associativity. This is more likely to occur when a new person attempts to modify another person's design. Thus a design often becomes the property of its originator.[9]

Companies have found that these properties of ProE tend to force them to change their design practices. Some new practices are healthy, such as encouraging expression and enforcement of design intent. But the unpredictability of ripple effects from changes has caused other companies to turn associativity off, sacrificing its configuration control for ease of design modification. Other companies have adopted design rules for constructing models so that there will be fewer surprises. Feature libraries, supported by ProE, make standardizing of modeling easier in some cases.

---

[7] Information in this section is based on two reports by D H Brown Associates plus discussions with their author.

[8] Many other CAD products offer similar capabilities, including Computervision CADDS5, Unigraphics, Intergraph, and SDRC's I-Deas Master Series.

[9] When associativity is active in a large CAD model, the performance of the software is often affected because it takes a long time to make all the associated changes.

It may be said that ProE has challenged companies to look concurrent engineering squarely in the eye and decide if they really want the design and the design process linked together in so many ways.

By contrast, SolidDesigner does not permit the designer to impose any persistent constraints by means of parameters or any other method. Any association, such as that two faces touch, can be used by the designer to facilitate making the design and may be broken later by anyone else. This makes modification of a design easy. In this sense, HP's approach puts less pressure on companies to seriously modify how they design in order to accommodate the CAD system.

D H Brown Associates' report on SolidDesigner puts the issue this way [D H Brown August 1994]: "This tradeoff between editing flexibility and capture of design intent may be a fundamental technology conflict. [CAD] Product architecture that offers SolidDesigner's level of editing flexibility may have to sacrifice an ability to retain design intent. Conversely, [CAD] products that allow users to permanently capture design intent through persistent constraints and object history orientation may never match SolidDesigner's editing flexibility. Today, no other product offers SolidDesigner's unique capability of surface grouping to define different representations of features on the same object. Feature definition depends on initial model construction in all [other CAD products] while retaining the advantages of permanence. Thus, SolidDesigner offers users a distinct choice in design methodology."

The following sections of the report deal with specific product or process design capabilities: mechanical assembly, computer servers, mass spectrometers, and microprocessors.

Mechanical Assembly (Lawson Fisher)

Fisher's group helps factories with special automation problems, typically by using commercially available robotic and hard automation systems and customizing them where necessary. They work exclusively for internal customers within HP. The group has seven members with a mix of backgrounds in ME and software. An interesting feature of their lab is the ability to work remotely with factories in such places as Oregon, Colorado, and Puerto Rico. This capability includes transferring software and remotely programming and operating a robot at such locations.

This capability is supported by HP's internal network and software called Shared X. This software permits local and remote people to share an X window and to use it as a whiteboard. A parallel telephone connection is used, and, by convention only, the one speaking "has the mouse."

When I visited, the group was working on a system to robotically plug a connector into a circuit board. The product in question was designed nearby and will be assembled in Colorado. The robot assembly equipment consists of a successor to "Robot World," a product originated by Automatix, Inc. under the direction of Victor Scheinman, a pioneer in robot system design. This system consists of cameras and robot hands that can move more or less freely in two dimensions while suspended magnetically from above. The cameras are used to find fiducial marks on the circuit board to aid in placing the connector accurately. Programming is done from a workstation using an array of programming windows that permit visual servoing, joystick control of the robot, motion planning, and straight C programming where necessary.

Over half a dozen such systems are in use doing precision jobs at low volume for prototype manufacturing. In addition to mounting odd-shaped components onto boards, some units are used to make prototype production runs of ink jet print heads.

Computer Server Design (Jim Smeenge)

This division's product is a network server, a high performance computer that is made in lower volumes (3000 - 5000 per month) than desktop computers. The main components are a 60486 or Pentium processor, a disk drive made by HP or another drive vendor, circuit boards and packaging (cabinet, board racks, ventilation, exterior appearance) designed by HP, and all other procurement and manufacturing done by suppliers in Singapore. Companies who started manufacturing in Singapore many years ago benefit from tax breaks, and in any case the costs of labor and many materials are very advantageous.

HP writes the software, which is complex. It involves network system management, storage and retrieval of data from many disk drives, and redundancy management among the drives. Lots of testing is required to be sure that software from many vendors (Novell, Microsoft, Banyan, etc.) functions properly. These systems are sold all over the world, so HP has adopted a practice called "postponement" to handle customizing the product to local situations. In this approach, the distribution center installs the software and provides the keyboard, which are both customized to the local language. The factory therefore makes one version of the product, which simplifies manufacturing greatly.

The design process begins with a specification received from Marketing which includes target cost, software functions, number of network connection board slots, and number of disk drives. Design begins with mechanical layout and electronic design and is usually cautiously evolutionary. Mechanical layout includes exterior appearance and user interface. Some rapid prototyping (stereolithography) is used. After prototypes of all types are available, a total cost estimate is made using the costs of similar parts on past products. Historically, this estimate has exceeded Marketing's target, so negotiation occurs over product features, and some are eliminated. When design is nearly done, quotes are obtained from Singapore. If the cost is still too high, more cuts in features or (rarely) a price reduction from the supplier are sought. But these are suppliers that HP has worked with for years so there are few surprises and few cuts available. Apparently no attempt is made to predict the cost prior to making prototypes and no attempt to predict what the cost drivers might be.

Since design is done in Palo Alto by HP while manufacturing and assembly are done in Singapore using mostly local suppliers, the design process is interesting. People travel back and forth a lot and there are many phone calls. Electronic design is aided by close ties with Intel, which permitted HP to field the first server using a Pentium.

Mechanical design comprises sheet metal and plastic parts, which are prototyped in Singapore. The longest lead time prototypes are plastic injection moldings, which take 12 weeks. ME 30 is used for all of these parts, and the designs are converted to IGES format for transmission to the vendors. It often takes two iterations to get the plastic design finalized, whereas sheet metal designs usually are right the first time. Plastic tooling design is simplified by frequent re-use of the previous product's molds, with minor changes for the sake of appearance or to permit a new arrangement of the disk drives.

Sheet metal design is aided by the fact that hard tooling is rarely cost-justified given the low volumes and short product life. The "non-hard" tooling used by HP is apparently robust, perhaps because people are more involved in shaping the parts. Hard tooling would take a long time to design and be very expensive. Design errors would take time to fix and the tooling would have to be brought to proper operation because hard tooling must operate unattended.

Circuit boards are designed and tested by HP, and early testing is done without the final housings because they are not ready, requiring substitute arrangements to simulate thermal loads. HP designs the circuits and a local board design supplier takes about 2 weeks to design a board. The design house's supplier takes 2 more to make it for testing.

Altogether, this is a well-oiled product development process, taking 6 to 9 months from start to finish.

Mass Spectrometers (Ed Cirimele)

Mass specs are ultra precision instruments which contain a vacuum chamber, high voltage electronics, and elements that generate large amounts of heat. Production rates are around 100 per month and the cost of a complete system of mass spec and gas chromatographer is about $65000. Design is a meticulous matter of thermal analysis, 3D tolerance analysis, and control software. Today it takes about three years to design such a product, a time HP would like to cut in half.

The product operates by taking in a sample of a gas whose chemical composition is to be determined. It is ionized by being heated and is then drawn through a quadrupole magnet. The field of this magnet is tuned to oscillate at frequencies at which particular chemicals will resonate. The response of a given chemical is sensed as heat. To test for a set of chemicals, the magnet is set to the appropriate frequencies one at a time. Other design implementations are made by both HP and several competitors.

The quadrupole magnets are made to extremely fine tolerances in the millionths of an inch and are implemented as quartz tubes plated inside and out. (See Figure 3) All the parts are designed by HP, fabricated by suppliers, and assembled by HP.

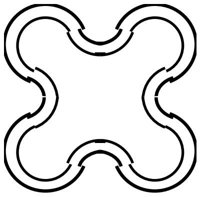

Figure 3. Approximate Cross-section Shape of Quadrupole

The main design tools available to help this process are ME 30, some commercial thermal analysis simulation software, and highly developed in-house software to simulate the ion beam. No suitable 3D tolerance analysis tools are available, although VSA was considered. The main design challenge is obtaining and maintaining the vacuum. A lot of lore and experience are involved because the strangest details can be sources of minute amounts of gas that spoil the vacuum. For example, no blind screw holes are used because the tiny amount of air trapped under the screw could leak out slowly. Great care is required in choice of materials. No plastics are allowed and this often means that wires have no insulation. Just the right chemicals must be chosen for cleaning and for cutting fluids during machining, since each of these could leave a tiny residue on the metal parts tat would later act as a contaminant.

Ed demonstrated ME 30 and commented on its strengths and weaknesses. The designer creates simple solids and then modifies their shape using a machining motif: cut, drill, etc. Reference surfaces can be called out and other features can be located with respect to the references by manual placement or typing numbers as dimensions. To make a part that fits into a slot, the easy way is to pick lines or vertices indicating the slot's width, declare them references, and then create new geometry that matches the references.

This type of design permits the designer to see the parts in the correct relative positions and orientations, so gross errors in nominal geometry do not occur. Tolerances, however, are not provided directly. He explained that there are three roles played by tolerances in product development: for design, for fabrication, and for quality control checking and measuring. Each presumably has its own rules and language, but the ANSI Y14.5 standard is really written from the point of view of measuring. He feels that, for design, tolerancing should be approached as one would think about assembly: a datum on this part establishes a datum on the next part, and so on. However, tolerancing is rarely done that way and the ANSI standard does not say how assemblies should be toleranced.

VLSI Design (Fred Harder)

Harder's group is responsible for providing VLSI design tools for use by HP's designers. Both commercial tools and homegrown ones are integrated into tool suites. His customers design microprocessors. This is important because two other kinds of VLSI should be kept in view for perspective: memory chips (DRAMs) and application-specific integrated circuits (ASICs). DRAMs represent the cutting edge, requiring the smallest features and redesigned individual device elements at every new generation. ASICs are at the other end, having relatively large devices and line widths and relatively fewer devices on a chip. Microprocessors are in between. To say it another way, a team of five designers, using high level design tools (more on this below) can design a control chip for a laser printer. But 300 people are on the design team for an advanced microprocessor with millions of devices on it.

The design of VLSI consists of a series of steps, approximately as follows:

1. Processor requirements are set, comprising data bus width, speed, operations in the instruction set, onboard facilities such as cache memory, floating point, etc., plus a target power use and dissipation.

2. A physical implementation is chosen, such as CMOS, which presents a family of candidate processes and materials.

3. A target chip size is chosen, and a preliminary division of the real estate is done, including connecting the main elements with "wiring."

4. As early as possible, some simulations are done to see what kinds of timing problems might occur; these problems arise due to the time it takes signals to propagate along the wires between the main elements as well as the time for an element to process a signal.

5. Logical implementations of the operations of each element are designed; the overall logic of the design can be tested at this point using simulators, and errors in the logic can be fixed.

6. Electronic implementations of the logic are designed, comprising circuit designs; these are done locally, that is, each logic element is individually converted into an electronic equivalent; the circuits repeat the topology of the logic. Detailed routing of conductor lines is done with the aid of computer algorithms.

7. Physical implementations of the circuit elements are drawn from a device library and linked together electrically using conducting lines of a given width and separation; width and separation are important because they affect resistance (power dissipation) and capacitance (which slows down the signal transmission speed or causes false signals to arise in adjacent lines). All devices and conductor lines consist of layers of different material. Each layer is possibly a different shape or made from a different material. The shapes of these layers are supposed to conform to "design rules" that indicate the capability of the manufacturing processes in terms of minimum line widths and spacings, radii of curvature, and so on. Associated with each device is a circuit behavior model that is used in the simulations mentioned next. These are "verified models" as they come from the library, but, as discussed below, it is not always possible to use the devices as is, so the models may not retain their validity.

8. The circuit is simulated using models of the devices and conducting lines, giving rise to possible timing or heat problems. Problems and errors are addressed.

9. Artwork similar to that used in multi-color printing is generated to represent the individual layers of the devices and conducting lines. This artwork will become the basis for masks that are used to control the deposition or etching of materials in the process of building up the circuit.

10. Masks are generated from the artwork. In order to improve the yield of the process, the process designers may alter details of the masks, generally thickening or thinning lines and features, adding layers to permit some of the features to be made, and so on.

11. The masks are taken to a wafer fab, a cleanroom with highly controlled filtered air and an array of plating and etching equipment. Here several test wafers are made. If the chip is a simple ASIC then the design usually works the first time. If it is a complex microprocessor then several iterations are often necessary, with the total time being a year to 18 months to the first test wafer and as much as another year debugging. The reason is that simulation is very hard to do because it takes so long to run even a few seconds of simulated time.

It is important to note that the process described above really works best for ASICs. For that kind of chip, it is an effective top-down approach in which many of the steps are done automatically. Few items in all of technology can be designed so automatically by proceeding from step to step, algorithmically converting requirements and symbolic representations of behavior into specific geometry without intervention by a person. Certainly it cannot be done in mechanical systems, at least where efficiency in space, weight, or power are required.

The power of this method lies in its ability to deploy verified elements and the ability to combine them in previously determined ways to carry out purely logical functions. Each element performs one function and one function only, in one-to-one correspondence with the symbolic or logical representation of the design. As long as each element and the conducting lines obey the design rules, a manufacturable design can be produced.

These elements have almost no effect on each other electrically when they are hooked together even in endlessly complex ways, so system design is effectively decoupled from element design. This fact is due to the high ratio of input to output impedance displayed by the devices. The designer can concentrate on designing the required system, getting its logic and main timing correct, without having to create the necessary elements or worry that they will not behave in the system the way they behave alone. An almost perfect divide and conquer design method can thus be used with high confidence that predicted element and system behavior will be obtained.

The analog of this in mechanical design is to pick a single element, such as a gear or shaft, and place it on a breadboard in approximately the same location as its functional equivalent appeared on the designer's first sketch or stick diagram. But several problems arise with this technique. First, mechanical elements transmit significant power and exert loads on each other. It is not practical to duplicate the high input:output impedance ratios characteristic of electronic elements. This means that system design cannot be separated from element design. Linked to this is the fact that such systems are not logic processors but instead energy movers and converters. Function includes but is not limited to logic. An example is a transmission in a car: logic is expressed in the gear ratios and the ability to go in reverse. But power transmission is the main task and most design effort goes into guaranteeing that the logic can be executed while the power is passing through.

A final mechanical design issue is that one cannot merely pick a separate element to perform each function, because the item would end up too heavy, too big, or too hungry for energy. Mechanical elements perform multiple functions anyway, so it is better to take advantage of this fact. The outer case of a transmission holds all the internal parts in their correct relative locations, keeps the oil inside, acts as a heat dissipater, and transmits torque from the engine to the internal parts. Skillful blending of the properties of geometry and materials to exploit the multiple functions of mechanical parts is a basic skill of mechanical designers. In addition, since each element is essentially redesigned for each new product, it is not possible to work from a library of proven elements with proven design rules that guarantee manufacturability. Thus another set of skills is required to convert product designs into process plans for fabrication and assembly.

The top-down VLSI design process was enabled in the late 1970s by Carver Mead and Lynn Conway whose textbook[10] showed how to use library elements with proven behavior and design rules to create cookbook designs. Before that time one had to be an electrical engineer to design VLSI. From that time until around 1990 one could be a logic designer and create VLSI. A downside to the Mead-Conway approach was that the library elements (called "standard cells") were not adjusted in shape when they were placed on the chip. The practical result is that such designs take up more space than space-optimized designs produced by EEs.

Harder agrees that the period from about 1978 and 1990 was a golden age in VLSI design. This period has come to an end, more or less, for several reasons. First, so many elements are now required for advanced processors that the loss of space created by library elements can no longer be tolerated. Larger chips are more vulnerable to manufacturing failures, especially tiny particles of dirt.[11] Process yield is the focus of manufacturing, and low yields mean loss instead of profit. Second, smaller elements and more closely spaced conductors require the skill of EEs to design and debug properly.[12] No longer can one blindly convert logic to circuits and expect them to work. Obeying the design rules will result in a chip that is too large, while pushing the rules requires understanding the currents and fields. Element design and system design are no longer independent, and VLSI design is taking on the character of mechanical design.

So we have come full circle, and it again requires EEs to design VLSI.

VLSI systems are among the most complex things designed, and logical or system level errors can occur. They represent the same kinds of errors that occur in other system designs: lack of coordination, imperfect interface specifications between subsystems, lack of a comprehensive database with all product data available to all designers, incompatible versions of the design being used by different groups of designers, and so on. System level tools to

---

[10] Mead, Carver, and Conway, Lynn, Introduction to VLSI Systems, Reading, Mass. : Addison-Wesley, 1980.

[11] This point was made by Gene Meieran of Intel, whose expertise lies in VLSI process design and manufacturing.

[12] Typical line widths and feature sizes are now 0.5μ and by 1998 will be 0.35μ. Current 16 MB memory chips have features sizes around 0.365μ. It is expected that 0.15 - 0.18μ is the target for 2001. At that point, it is unlikely that further progress will be possible using light as the exposure medium and masks as the format for design embodiment. See [Stix].

handle these problems are either not available or are just becoming available. The best tools handle individual steps or aspects of the design. It could be, too, that HP's design culture values point tools over system integration tools, but that cannot last because complexity will rise even as product development time must fall.[13]

One other interesting point emerged. Since DRAMs are at the cutting edge, their existence depends on ever better fabrication tools and cleaner fab facilities. Companies that make DRAMs therefore drive the advancement of manufacturing equipment. Microprocessors are next in design rule demands, and can use equipment developed for DRAMs after it has been proven in the DRAM environment. Processors therefore owe their existence to the skill of system designers and can wait for manufacturing technology to fall in their lap, thanks to DRAM production. The Japanese have chosen to produce DRAMs, which benefit from their skill in clean precision operations, while the US has chosen to produce processors, which benefit from superior system design capabilities as well as more aggressive design for the computer systems that use the processors. Thus a certain synergy has emerged, with the Japanese generating the DRAMs and spontaneously generating the associated manufacturing capability which the US buys in order to create the processors and associated computer systems that consume the DRAMs.

Summary

HP is a good example of a company that uses a number of its products as part of its own design and manufacturing infrastructure. This practice stands in contrast to the majority of US companies who mainly buy those capabilities. HP is also very systems-conscious and appears to be evolving an approach to product design that values a systematic approach. In addition, electronic commerce is emerging as a mainstay of HP's business methods. Evidence of this emerged in these interviews, and more recent information indicates that the trend is strengthening.

References

[BW] "They've Got the Whole World on a Chip," Business Week, April 23, 1994, pp 132-35.

[Geppert] Geppert, Linda, "Not Your Father's CPU," IEEE Spectrum, Dec., 1993, pp 20 - 23.

---

[13] Typical microprocessors today have about 3 million transistors. [Geppert] Industry experts estimate that some kinds of chips will have as many as 50 to 100 million transistors by 2000. The growth in transistors per chip greatly exceeds the growth in productivity of chip designers. [BW]

[Liker et al] Liker, J K, Kamath, R R, Wasti, N, and Nagamachi, M. "Integrating Suppliers into Fast-Cycle Product Development," to appear in Managing Technology Through Organization: US and Japanese Approaches, to be published by Oxford University Press, 1995.

Status of Solid Modeling: PTC Drives Technology, D H Brown Associates, Port Chester NY, July, 1994

HP's SolidDesigner Breaks with the Mainstream, D H Brown Associates, Port Chester NY, July, 1994

[Stix] Stix, Gary, "Toward 'Point One,'" Scientific American, Feb., 1995, pp 90 - 96

[Whitney, Nevins, De Fazio and Gustavson] Whitney, D E, Nevins, J L, De Fazio, T L, and Gustavson, R E, "Problems and Issues in Design and Manufacture of Complex Electro-mechanical Systems," C S Draper Lab Report R- 2577, Dec., 1993. A copy of this report is available as:
(http://elib.cme.nist.gov/made/papers/Whitney/paper.html)

figure 1

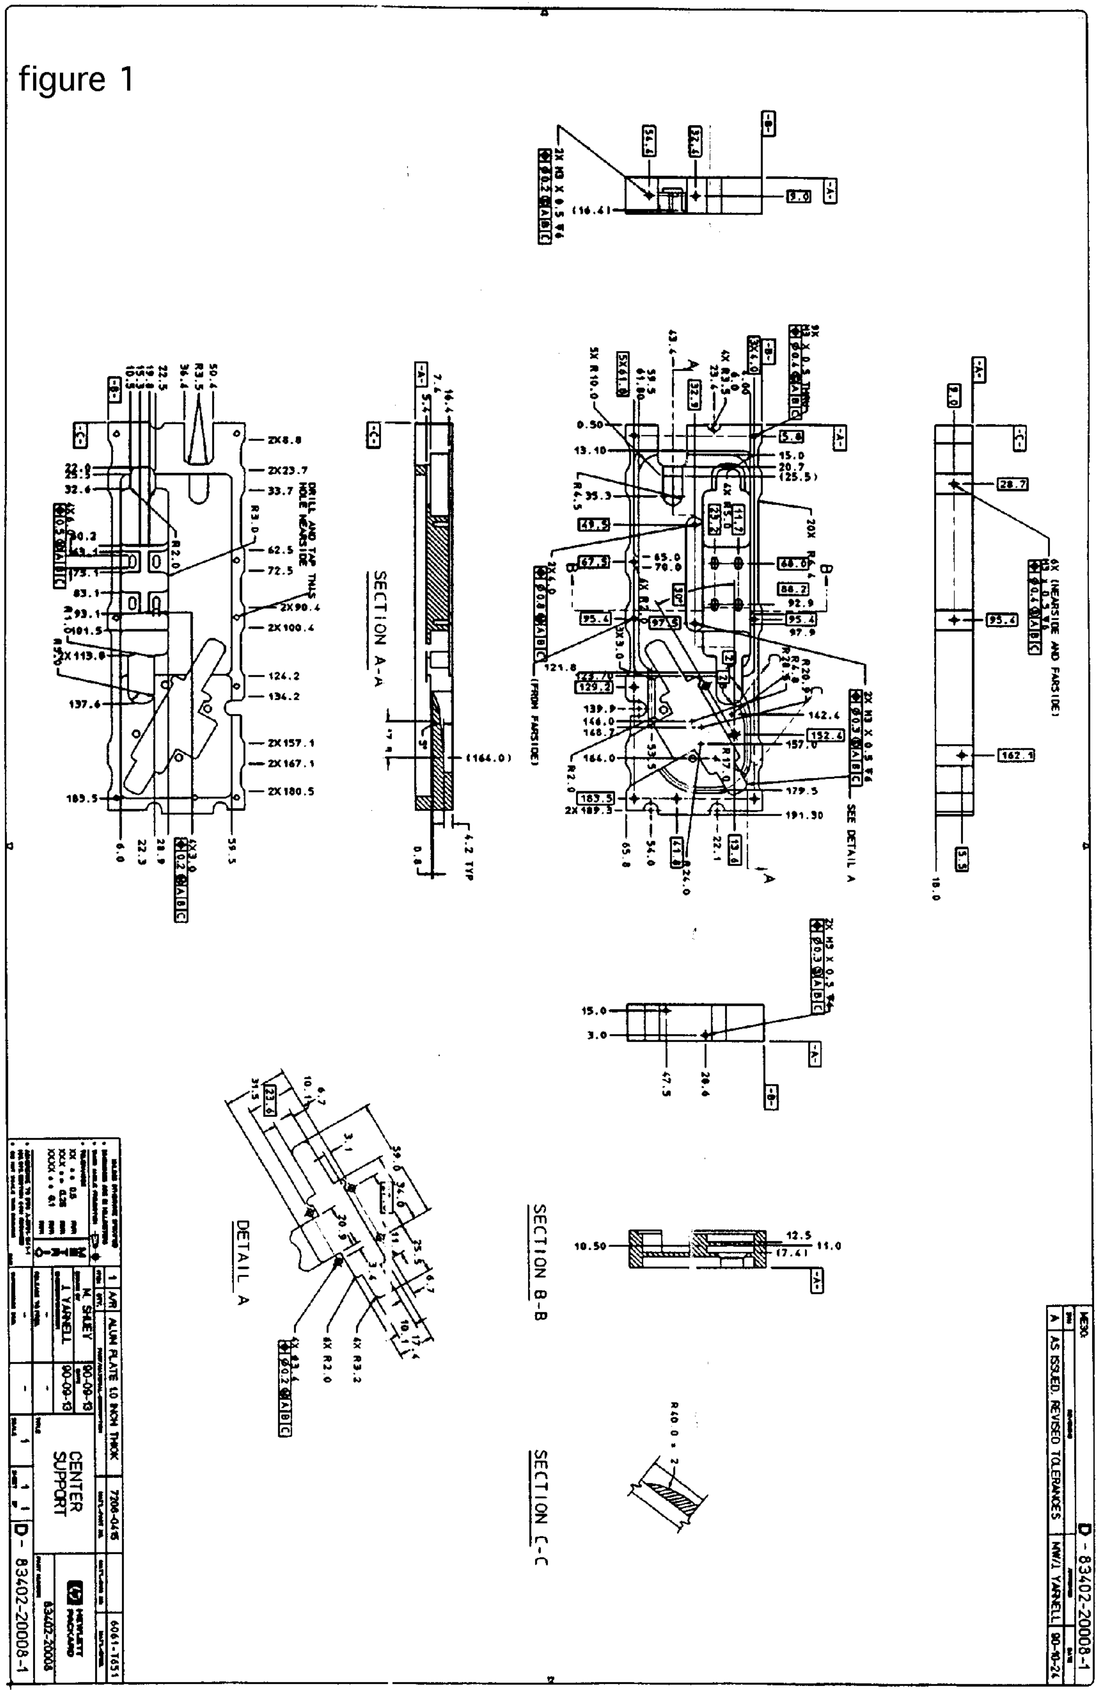
SECTION A-A
SECTION B-B
SECTION C-C
DETAIL A
SEE DETAIL A
DRILL AND TAP THIS HOLE NEARSIDE
6X (NEARSIDE AND FARSIDE)
(FROM FARSIDE)
4.2 TYP
A AS ISSUED, REVISED TOLERANCES MW/J YARNELL 90-10-24
D - 83402-20008-1
1 A/R ALUM PLATE 1.0 INCH THICK 7206-0415 6061-T651
XX ± 0.5
XXX ± 0.25
XXXX ± 0.1
M. SHUEY 90-09-13
J. YARNELL 90-09-13
CENTER SUPPORT
HEWLETT PACKARD
83402-20008
SCALE 1
SHEET 1 OF 1
D - 83402-20008-1

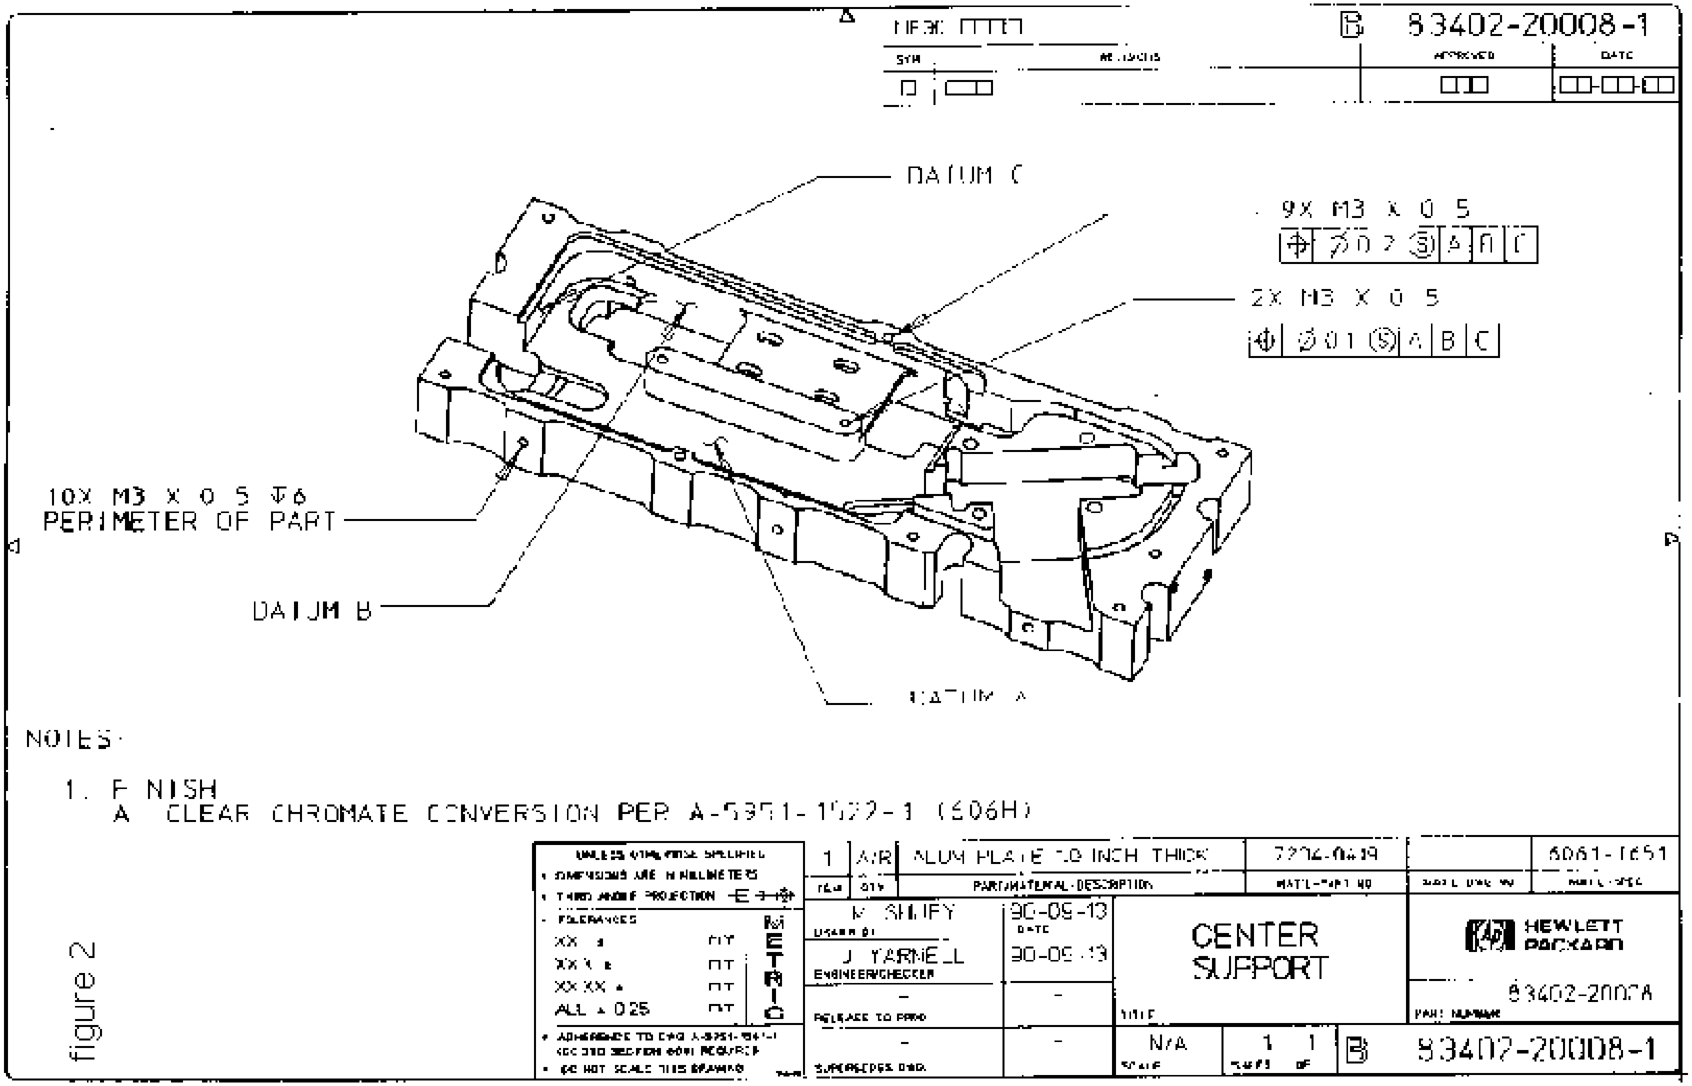
83402-20008-1
DATUM C
9X M3 X 0 5
⌖ ⌀0 2 Ⓢ A B C
2X M3 X 0 5
⌖ ⌀0 1 Ⓢ A B C
10X M3 X 0 5 ↧6
PERIMETER OF PART
DATUM B
DATUM A
NOTES:
1. FINISH
A CLEAR CHROMATE CONVERSION PER A-5951-1522-1 (606H)
UNLESS OTHERWISE SPECIFIED
DIMENSIONS ARE IN MILLIMETERS
THIRD ANGLE PROJECTION
TOLERANCES
ALL ± 0.25
METRIC
DO NOT SCALE THIS DRAWING
1
A/R
ALUM PLATE .50 INCH THICK
ITEM
QTY
PART/MATERIAL DESCRIPTION
MAT'L PART NO
MAT'L SPEC
6061-T651
M SHIFY
DRAWN BY
90-09-13
DATE
J YARNELL
ENGINEER/CHECKER
90-09-13
RELEASE TO PROD
SUPERSEDES DWG
CENTER SUPPORT
TITLE
HEWLETT PACKARD
83402-20008
PART NUMBER
N/A
SCALE
1 OF 1
SHEET
B
83402-20008-1

figure 2