Scientific American FEB 1978 IAS 3/2/89

Computer-controlled Assembly

High-volume products are assembled by people or by special-purpose machines. An experimental programmable robot suggests that robots would be cost-effective for the assembly of products in lower volumes

by James L. Nevins and Daniel E. Whitney

Continuing inflation, competition from other countries and record deficits in international trade have created a widespread awareness in the U.S. of the need to increase productivity in manufacturing, which means decreasing the man-hours, materials, energy or capital required to produce industrial goods of all kinds. An additional stimulus for increasing productivity arises from the desire to improve the quality of life, including the life of workers now engaged in stultifying, repetitive and sometimes hazardous tasks. For most of the past century growth in manufacturing productivity has been maintained by the substitution of power-driven machines and new technological methods for labor. Today there are pressures to use power, materials and capital more efficiently. The conventional remedies need to be reexamined and new solutions need to be sought.

Although there are many ways of increasing manufacturing productivity—financial, fiscal and social—we shall focus here on the possibility of raising productivity through the application of science and advanced technology to an old field: assembly. Technology has brought about radical changes in many areas—power generation, transportation, chemical manufacturing, communications and data processing—but it has had only a minor effect on the way the broad spectrum of consumer goods, from electric toasters to automobiles, are actually assembled. Alert to the need for raising productivity, Germany, Japan, Norway and other countries have begun long-term government financing of research on manufacturing methods. In the U.S. the National Science Foundation and a few industrial firms joined forces several years ago to support similar studies, although on a much more modest scale.

Attacking assembly alone will not be sufficient. The fraction of the manufacturing labor force engaged in assembly operations varies widely from industry to industry. It is seldom less than 10 percent, and even in the automobile industry, which has the volume to justify a heavy investment in mechanization, roughly a third of the total work force is engaged in assembly. Thus massive shifts in productivity can come about only if great changes are made in the entire production system, including assembly but extending well beyond it.

At present manufacturing is based largely on experience; it is really an art form. Equipment designers and factory managers prefer to repeat what has worked in the past, which can be taken as evidence that they are struggling with a vastly complicated situation. Such changes as are introduced tend to be small ones. Wholesale shifts in technique are expensive and risky.

At the Charles Stark Draper Laboratory in Cambridge, Mass., with support from the National Science Foundation and industry, our goal is to contribute to a base of manufacturing knowledge from which new theories, experimental techniques and assembly methods can emerge rather than to develop devices of limited value one at a time to meet specific factory problems. With guidance from our industrial partners we have developed economic models of the role of assembly in manufacturing. In this way our studies can be focused on the problems most in need of solution at the same time that we are acquiring the basic knowledge for solving them.

One lesson learned from this approach is that certain problems addressed in the past are of little interest to industry. For example, highly complicated robot-arm computer systems, guided by television "eyes," that will pick up parts arriving in any orientation and mate them to other randomly oriented parts, although intellectually challenging, do not tackle the core problems of assembly. What happens when parts touch each other? How can close-tolerance parts be mated when there is no way to see into the hole where wedging and jamming actually take place? What should a computer-controlled assembly machine "know" in the way of genera assembly skills and what should it b taught "on the job" in order to perform particular task on the factory floor What industrial products lend them selves to robot assembly? One must b able to answer these questions and sim lar ones on a firm factual basis if assem bly technology and other manufactur ing technologies are to advance.

At present assembly is performed b people and, when the productio volume is high enough, by specia purpose machines. People are readil taught new tasks, and they adapt t changing conditions (such as slight var ations between one part and the next), t different models of the same product o the same assembly line and to maj changes in product design. They mak skillful use of sight and touch both t move objects around and to carry o the fine tasks required for assembly. A a result they have little need for speci tools and parts holders (generally calle jigs and fixtures). On the other han people are subject to fatigue and the i ability to perform a task exactly t same way time after time. These limit tions often lead to substantial problem in quality control.

Special-purpose machines are ve efficient, give reproducible performanc and are not subject to fatigue, but the consist almost totally of jigs and fixtur built to perform one task, or a close related series of tasks, on one produc They cannot easily be altered to accon modate different models on the produ tion line or changes in product desig Such machines do not have sensors t guide or monitor the assembly proces although they can perform simple tes on the assemblies. Each machine is us ally one of a kind, expensive and labor ously tuned to the accuracy necessary t handle one set of parts. Slight variatior in the parts can cause such machines t jam as much as a third of the time. Th high cost of special-purpose machin and their inability to handle new tas or products limit them to applications where millions of identical units are made each year for several years. It has been estimated that such high-volume production characterizes only 5 percent of all goods manufactured.

Most products are manufactured in batches with wide style variations, in quantities too small or a design life too short to justify investment in a special-purpose machine. In addition many items are not designed with sufficient attention to assembly problems, partly because assembly phenomena are not well enough understood to allow precise requirements to be placed before the product designers. Moreover, assembly-line workers do more than just put parts together. They often make spot repairs and perform many vital inspection tasks. It is for these reasons that most assembly is still done by people. It is unlikely that completely unmanned factories will ever exist, because people will always be needed to supervise and repair the machines.

Our research group at the Draper Laboratory has been examining the hypothesis that low-volume batches of a product (a few hundred thousand items or fewer per year) and mixed batches of similar models of a product can be assembled profitably by machines that are adaptable and programmable. An adaptable machine is one that can perform an assembly task as it accommodates itself to relative position errors between the parts. Such errors arise from the usual tolerances within which all manufactured parts are allowed to vary and from the lack of perfectly repeatable performance by the assembly machine itself. It is these errors that cause parts to jam rather than go together smoothly. Adaptability is therefore central to successful assembly. A programmable machine is one that can be taught, with minor alterations, to perform a new assembly task or that can perform several tasks in sequence. This capability is essential if assembly ma-

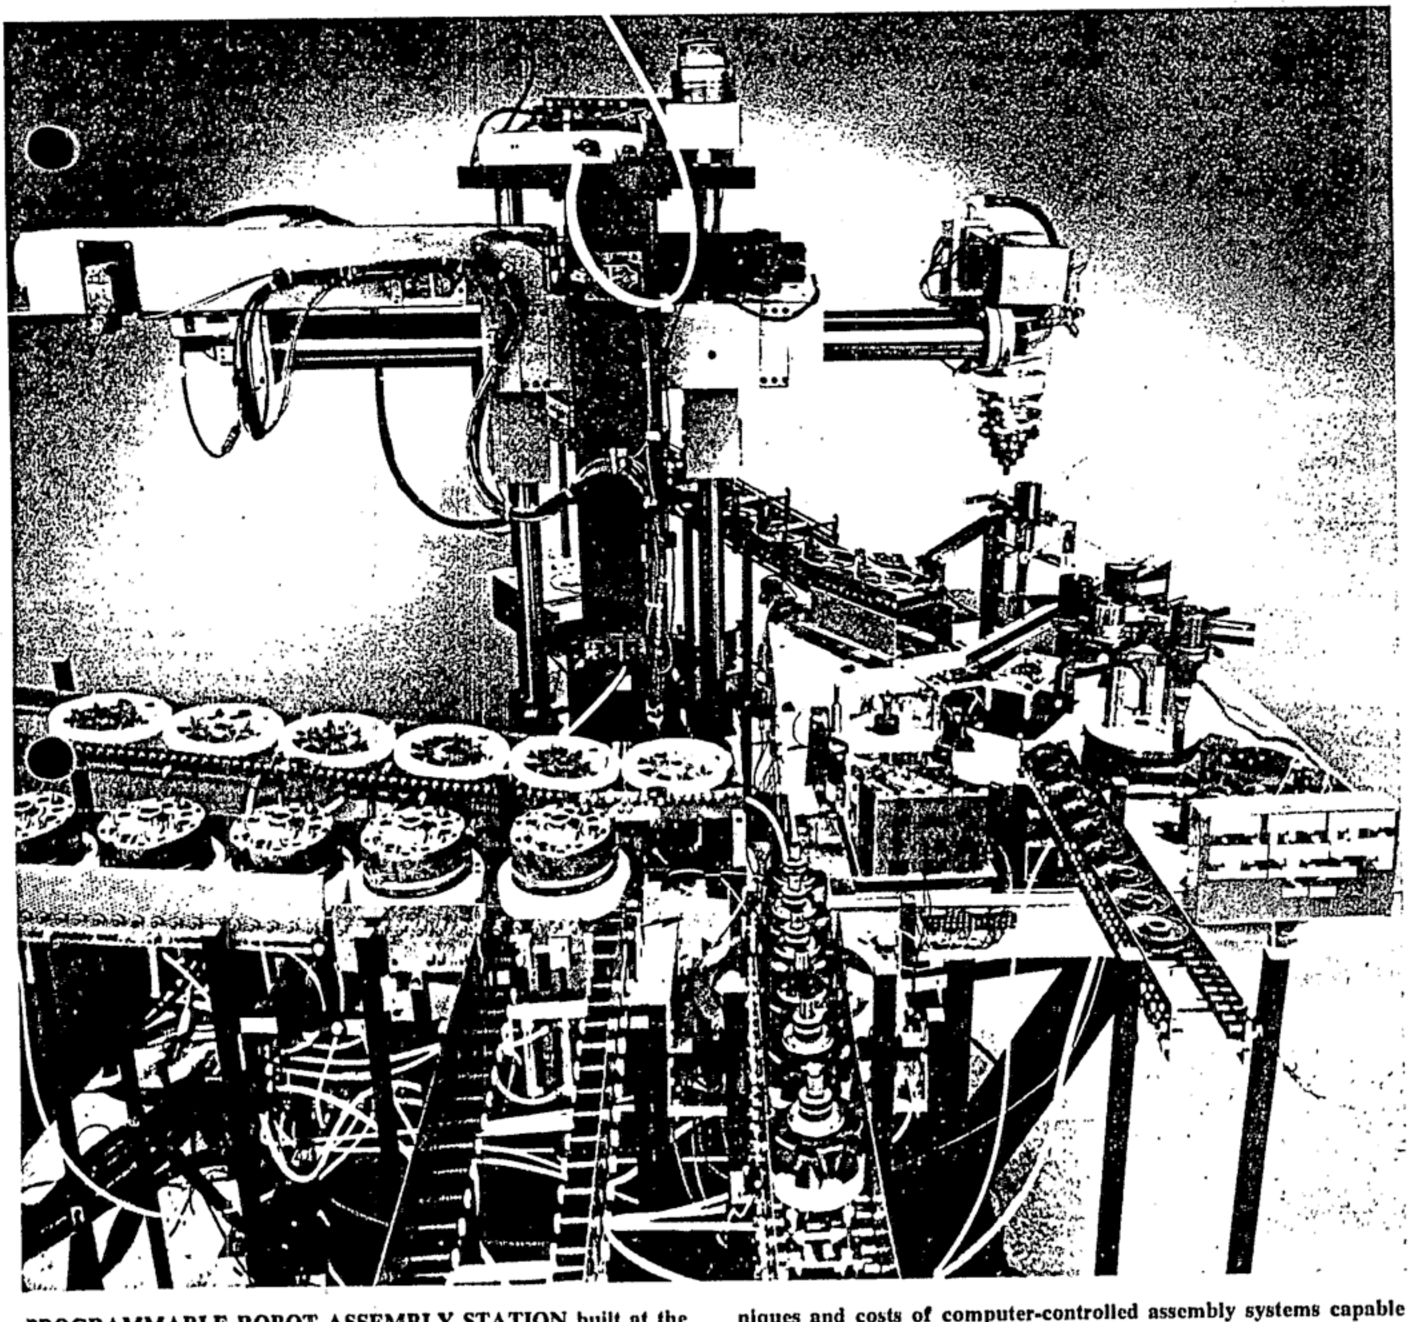

PROGRAMMABLE ROBOT ASSEMBLY STATION built at the Charles Stark Draper Laboratory can assemble the 17 parts of a commercial automobile alternator in two minutes 42 seconds. At the far right is a control box through which the robot can be taught a sequence of moves that can then be recorded in the memory of a minicomputer. The robot serves as a test bed for exploring theories, techniques and costs of computer-controlled assembly systems capable of being reprogrammed for various comparable tasks. The alternator was selected for the assembly experiment because it is an actual industrial product and thus requires mating of component parts that have standard industrial clearances. Alternator was also chosen because it is a "stack" product: all the parts can be added from a single direction.

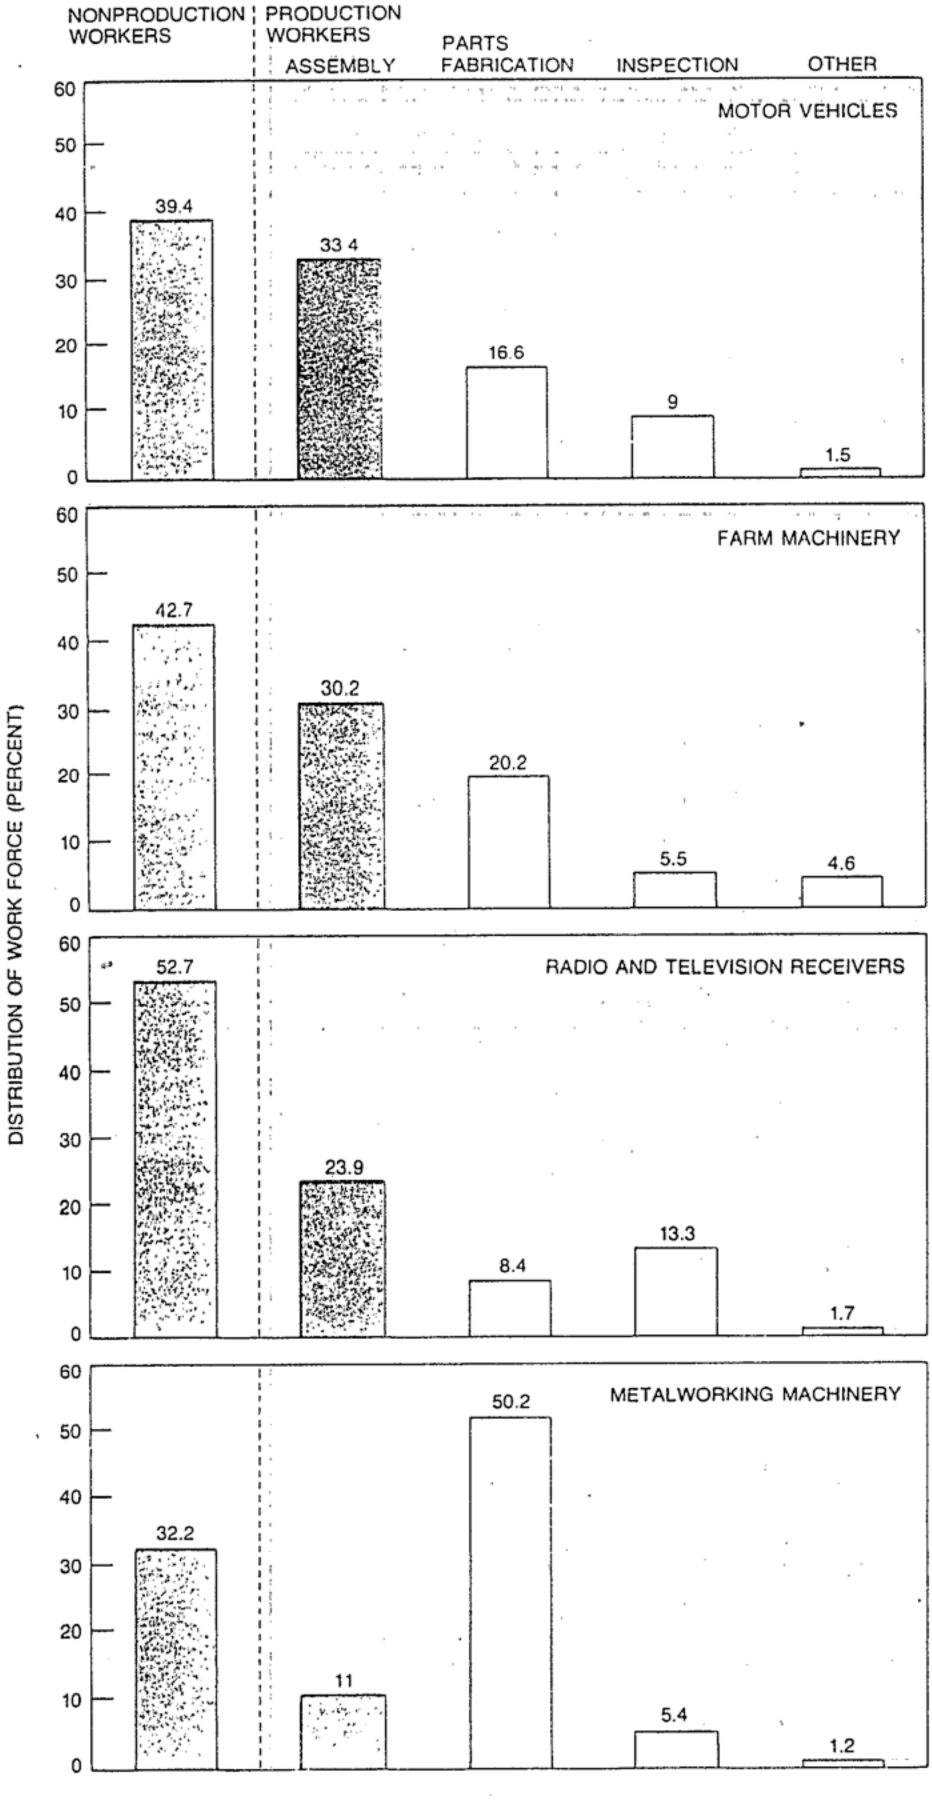

DISTRIBUTION OF LABOR in four major durable-goods industries demonstrates the large role that manual-assembly labor still plays, even in the highly mechanized motor-vehicle industry. The different distributions are characteristic of the industries, for example the intensive use of labor in fabricating precision parts of metalworking machinery and in inspecting radio and television receivers. The data for the chart are taken from the 1970 U.S. Census.

chines are to be economic for low-volume manufacturing.

Before setting out to design an adaptable and programmable assembly machine one must answer a number of questions. What assembly tasks will the machine face, or what tasks are appropriate? What should the machine's performance capabilities be in terms of speed, size and accuracy? Should it be capable of all the possible motions a human arm can execute, or of more or fewer—including both gross motions and fine ones? Should the machine be provided with efficient but inflexible fixtures or with an elaborate reteachable sensing capability to enable it to find the exact locations of parts? If the machine is to be programmable, it cannot be built from the start to do its assigned task; how then is it to be "taught" what to do? Should one machine be expected to assemble an entire product or should the assembly tasks be distributed among several machines that pass partially completed work along as in a conventional assembly line?

These questions have plausible but conflicting answers. Our approach has been to divide the problem into two segments: parts-mating phenomena and assembly systems. The mating of parts involves all the events that occur as parts touch and go together. Such events are governed by the geometry of the parts, particularly including the amount of clearance (or free space) between them after assembly, by the degree to which they are misaligned laterally and angularly when they first touch and by the influence of contact and frictional forces between them as they slide together. In order to understand the mating of parts we have studied idealized tasks, formulated hypotheses and verified them experimentally.

When we began this work five years ago we had several further assumptions, some of which have survived and others have not. We postulated machine systems with several work stations, each of which incorporated an "arm" something like today's conventional industrial robot [see "Robot Systems," by James S. Albus and John M. Evans, Jr.; SCIENTIFIC AMERICAN, February, 1976]. Such devices repeat a sequence of taught moves, combined with the opening and closing of grippers, that enable them to transfer objects from place to place. Current industrial robots do not have sophisticated controls or sensors to allow modified behavior in case of difficulty, although they can detect trouble and stop before damage is done. Neither are they accurate enough to perform assembly. Many of these limitations are being overcome or soon will be by research and redesign. At present, however, the robots are too big and too expensive for most assembly applications, and it is questionable whether assembly applications require all the motions they are capable of.

We further postulated that the arm at each work station should be capable of executing both gross motions for the transfer of parts and fine motions, measured in fractions of a centimeter, for the assembly of parts. In order to guide the fine motions we proposed placing in the arm's "wrist" a sensor capable of detecting both forces and moments. We concluded early that the spatial information necessary to put close-fitting parts together could be sensed as physical forces much more readily than it could be acquired by viewing mechanisms such as television.

Whether parts will mate successfully in any given instance depends on the relative error between the parts as they touch for the first time. If the parts touch with some relative error, a contact force will arise, causing the parts or the grippers and jigs to deform slightly. The deformation has the effect of altering the path along which the arriving part moves, with results that are either adverse, beneficial or neutral. If the relative error is small, mating can proceed without difficulty, but adverse effects are increasingly likely as the errors get larger. If the addition of expensive and inflexible fixtures is to be avoided, our mating mechanisms must be able to tolerate errors larger than those tolerated by existing machines. This requirement means not that the contact forces merely be allowed to work their will on the motions but rather that a strategy for producing beneficial responses and

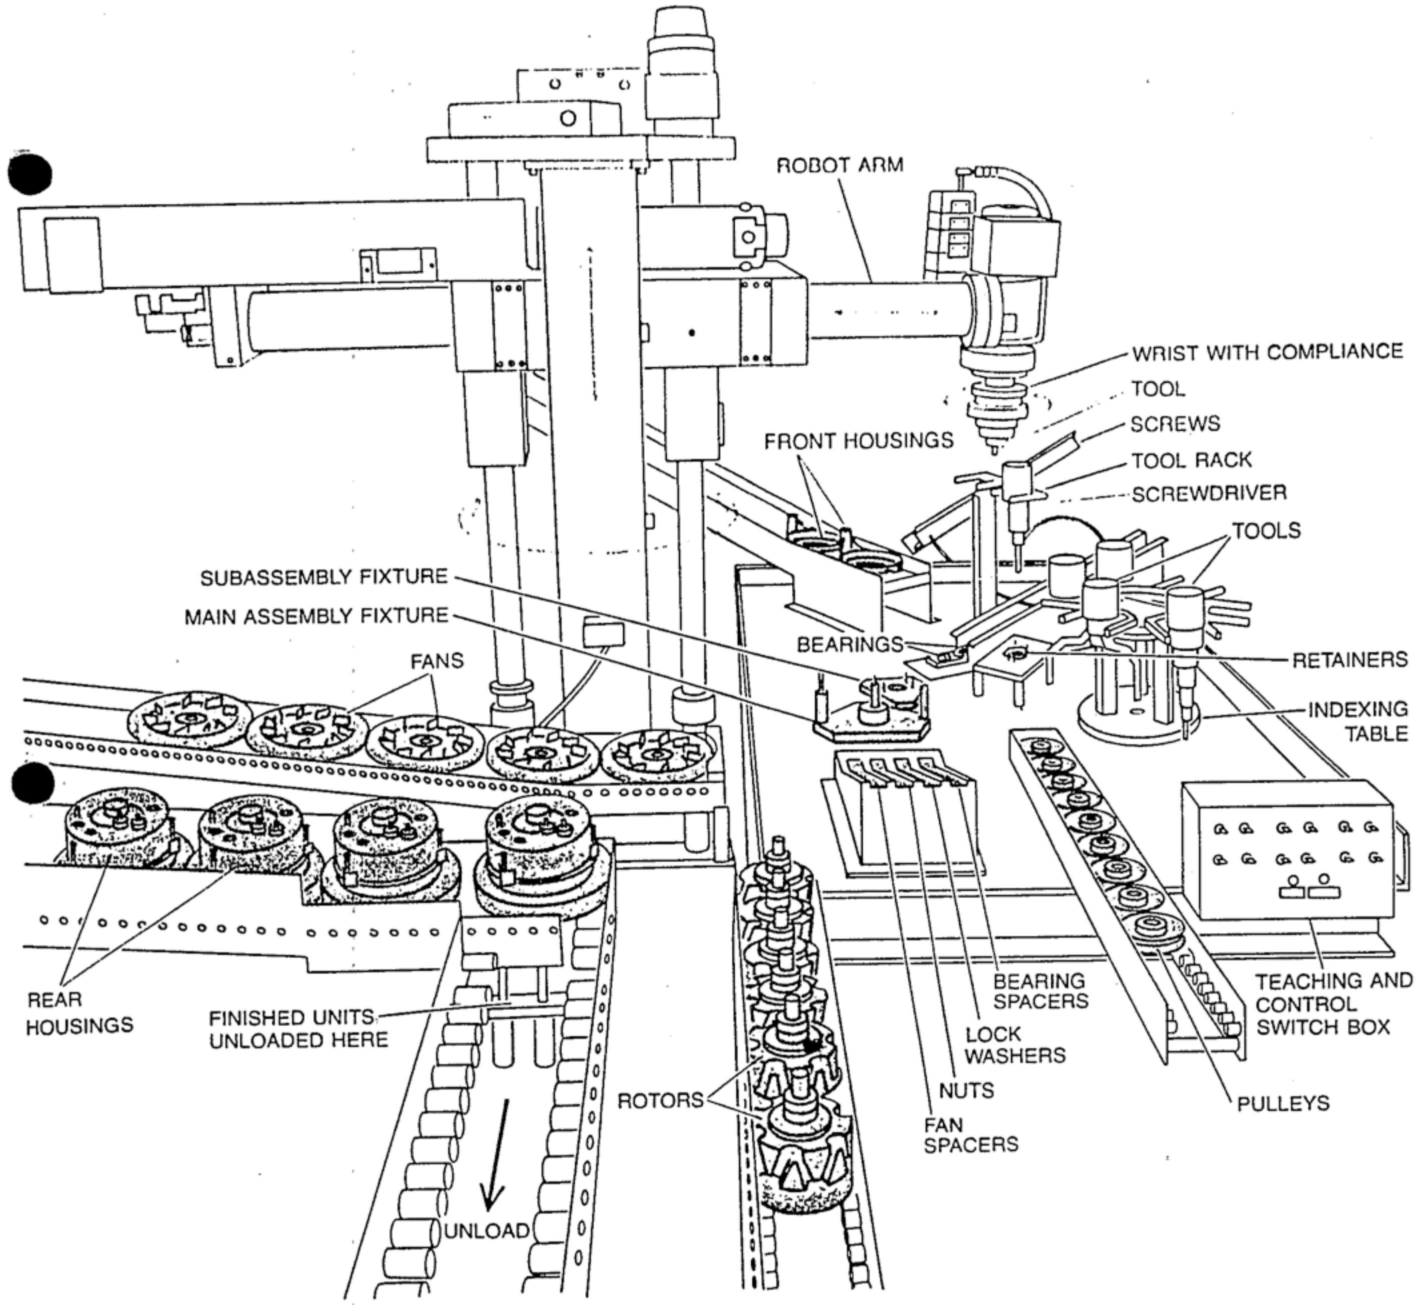

LAYOUT OF ROBOT ASSEMBLY STATION shown in the photograph on page 63 places tools and parts within easy reach of the robot's arm, which has four degrees of freedom (*colored arrows*). The assembly task requires six different kinds of tools, held on a table that "indexes," or turns, to supply the proper tool for each operation. The alternators's 17 parts are fed by gravity from 12 feeders. (The 17 parts include three screws, which have only one feeder, and three long bolts, fed together with the rear housing.) The assembly is performed on two different fixtures, one for the main assembly, the other for a subassembly. The robot is operated by a computer that drives the four joints to designated stopping points at designated speeds. The points, speeds and tool operations are programmed with the aid of a control box and a simple keyboard language. The language names and sequences the points and tool operations. A major feature of the robot is a wrist-and-gripper mechanism that responds compliantly so that parts can be inserted into close-fitting openings without jamming.

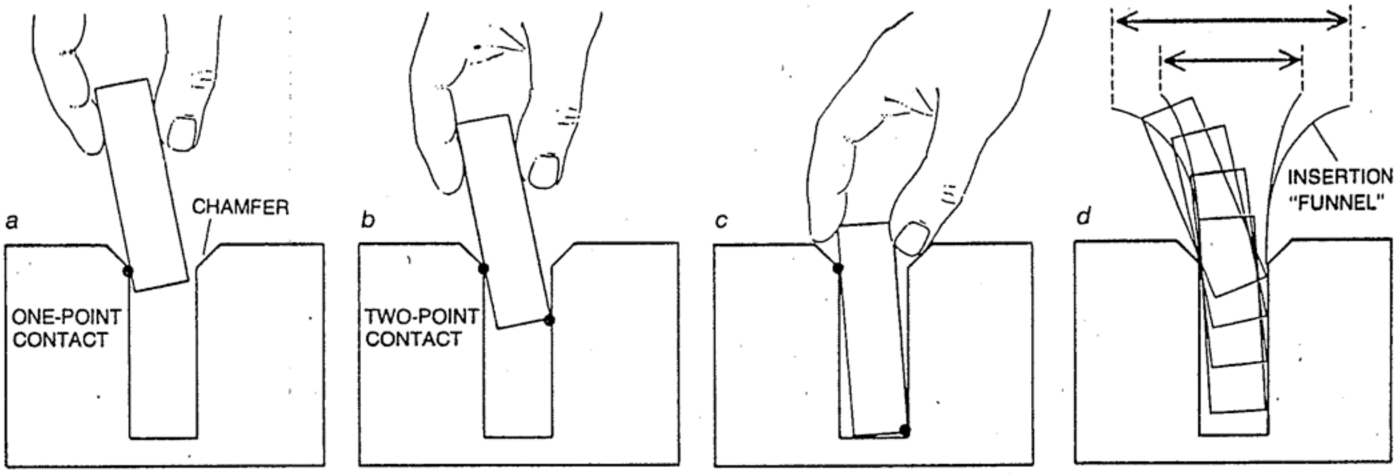

INSERTION OF A PEG IN A HOLE, a typical assembly task, is basically a problem in positioning. Holes are usually chamfered (beveled around the edge) to aid insertion. As the peg slides down the chamfer and enters the hole (*a*) it touches one side of the interior first (one-point contact). If the angular misalignment is large (*b*), the peg will soon touch the opposite side of the hole as well (two-point contact), with danger of jamming. In manual assembly vision can help to find the chamfer, but after the peg enters the hole one must rely on the ability to sense the resisting forces in order to maneuver the peg to the bottom (*c*). The geometry of the peg and the hole keeps the peg within the insertion "funnel" (*d*), the path that is traced by the top of the peg at successively deeper stages of the two-point contact. The smaller the clearance between the peg and the hole, the narrower the insertion funnel and the more difficult the insertion task.

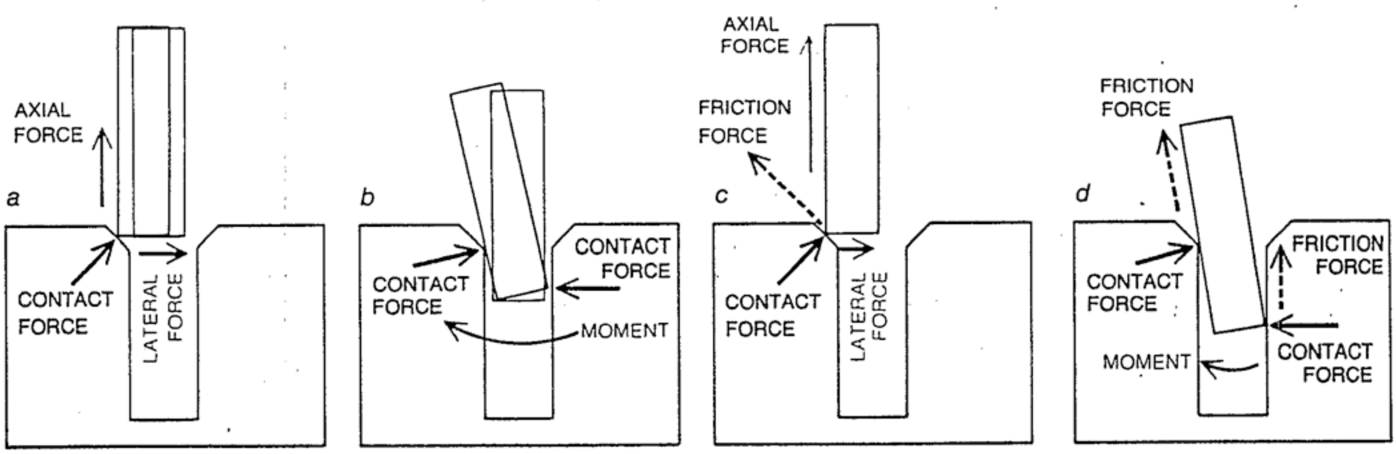

CONTACT FORCES BETWEEN PARTS can be used to guide corrective motions of the "wrist" of an assembly robot's arm. In the absence of friction (*a*) the contact force at the chamfer is sensed as two equal reactions, one vertical and the other lateral. The lateral force can serve as a cue to the desired corrective motion (*colored arrow*). Later (*b*) contact forces create a moment around the tip, which provides a cue to the desired corrective motion (*colored arrow*). When there is friction (*c*), the upward reaction at the chamfer is exaggerated, reducing the useful lateral reaction. Friction also reduces the useful information about moment (*d*). The ratio of the friction force (*broken arrows*) to the contact force, in other words the coefficient of friction, is about 0.2 for steel parts and 1.0 for aluminum ones.

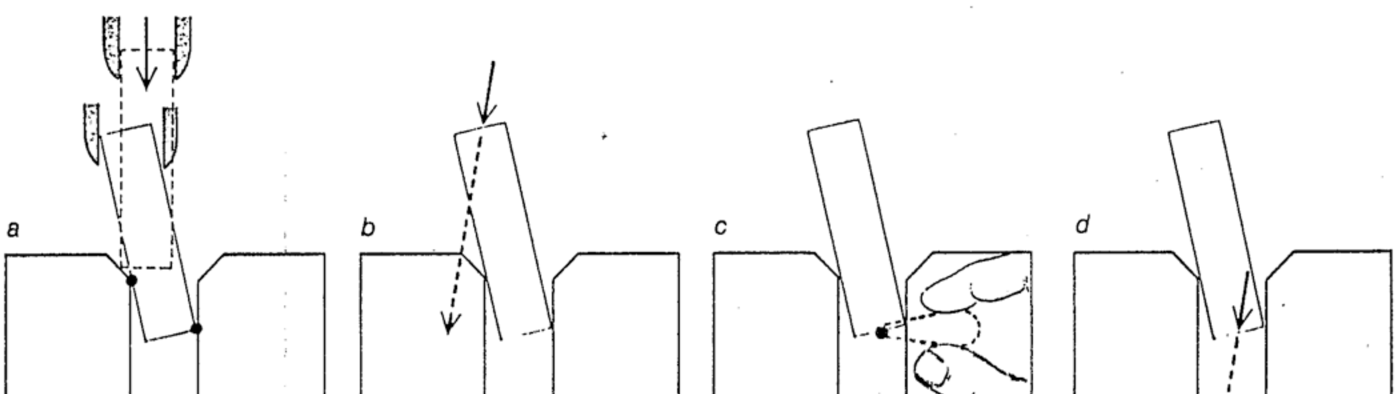

IF COMPLIANT GRIPPERS are used to hold a peg, a lateral error becomes an angular error as the peg slides down the chamfer and enters the hole (*a*). Continued application of force at the top of the peg (*b*) creates a torque that can lead to jamming. If the peg could be grasped compliantly at its tip (*c*), insertion could be accomplished in spite of substantial angular error. The same force that would lead to jamming if it were applied at the top of the peg would tend not to cause jamming if it could be applied at the bottom of the peg (*d*).

avoiding adverse ones be determined in advance.

One of us (Nevins) proposed that the force sensor on the wrist could be used to measure all the components of the vector of the contact force. The sensor would be designed to measure three components of force along three mutually perpendicular axes, *x*, *y* and *z*, together with three components of torque around the same axes. Thus we proposed to perform assembly as a blind person might, by measuring the forces and executing appropriate motions in response to what was felt, thereby correcting errors in steps until the parts were successfully mated. One of us (Whitney) then formulated a general strategy for force feedback that generates a vector of motions (three *xyz* translations and three rotations about the *xyz* axes of the tip of the entering part) in response to the sensed vector of the contact force.

We analyzed and experimentally verified this technique on two different computer-controlled robots and in the process learned a great deal about harnessing sensory feedback for assembly. The technique has many of the characteristics of a closed-loop control system, which operates as follows. The wrist force sensor detects the contact-force vector. The assembly strategy, residing in the arm's controller, will call for a certain response motion whose direction depends on the direction of the contact force and whose magnitude depends on the magnitude of the force. This motion, when it is executed, will change the deformation of the parts and the grippers, thereby changing the force. New force gives rise to new motions, which give rise to new force and so on in a loop.

Care must be taken in designing the strategy so that the right amount of motion is called for in response to the felt force. Too much motion will cause the arm to react as a person does when he touches a hot surface; too little motion will let large contact forces build up to a damaging level. The less stiff (more compliant) the parts and the grippers are and the lighter the arm's moving components are, the easier it is to obtain rapid, stable and effective responses with low contact forces. When low stiffness and rapid response motion cannot be built into the apparatus (because, for example, it is too heavy or the workpieces it is holding are), the only remedy for avoiding large contact forces is to make all the closed-loop motions slowly. This alternative is an unattractive one from an economic point of view.

One conclusion of this work was that assembly is best understood in terms of the forces and moments acting on the tip of the part, where it touches its mate during assembly. Another conclusion was that devices capable of fine assembly motions must be small, light and fast. Arms capable of gross motions cannot meet these criteria, indicating that in future assembly systems fine- and gross-motion devices, like the human hand and arm, will be separate entities.

A third conclusion was that friction can mask the force data needed to guide the strategy and can cause the parts to jam rather than slide together. Our coworker Sergio N. Simunovic has demonstrated mathematically that there are really two distinguishable phenomena: wedging and jamming. Wedging involves deformation of the parts themselves and occurs only if the entering part, such as a peg, enters a hole at such a sharp angle that it touches opposite sides of the hole before it has penetrated very far. More specifically, wedging occurs when the ratio between the depth of insertion and the diameter of the hole is smaller than the coefficient of friction:

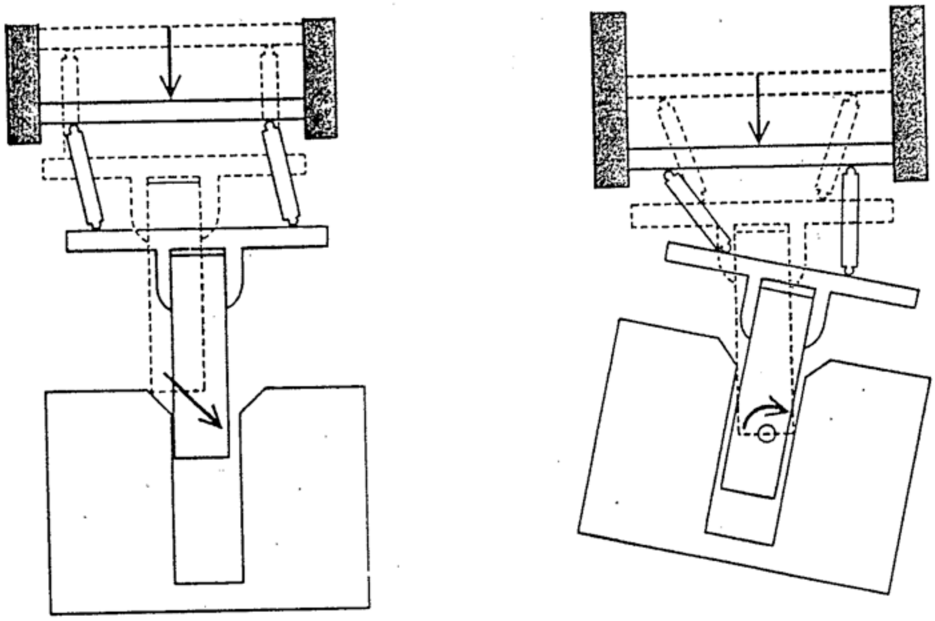

GRIPPERS WITH LINKAGE SUPPORTS can provide the proper response to an error either in lateral position (*left*) or in angular position (*right*) but not in both. The first linkage allows a laterally misaligned peg to slide down a chamfer and enter a hole without rotating. Second linkage allows a misaligned peg to become aligned by rotating at the tip without moving sideways.

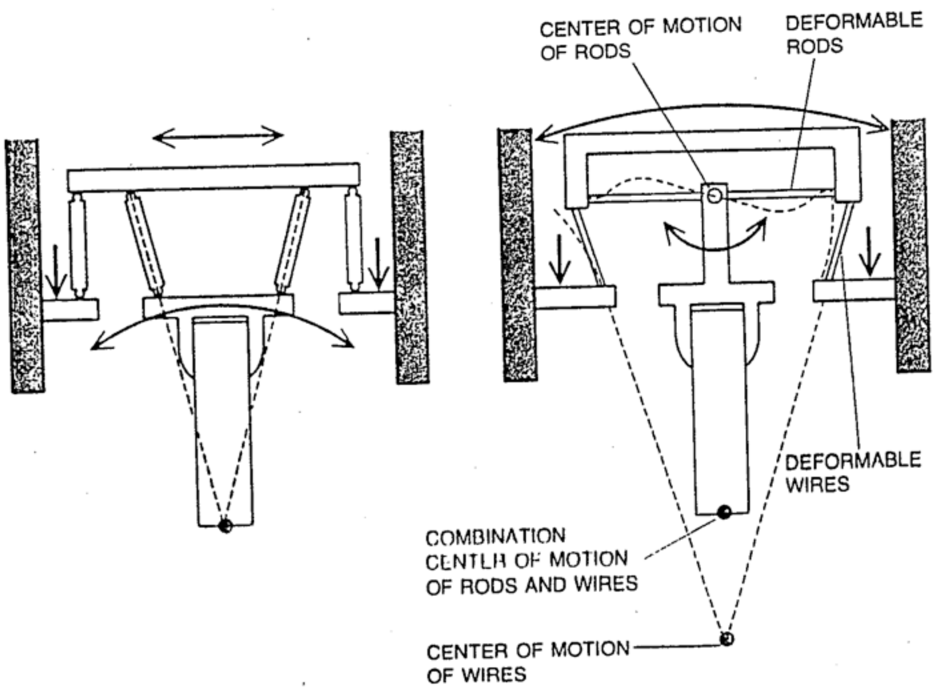

GRIPPERS WITH COMPOUND COMPLIANCE can correct for both lateral and angular misalignment. The "remote center compliance" at the left combines the two types of linkage depicted in the upper illustration. If a lateral force is exerted at the tip of the peg, it will translate (move sideways) without rotating. If a torque is exerted at the tip, the peg will rotate without translating. The device shown at the right can correct for misalignments by a suitable arrangement of deformable rods and wire springs, which are used in place of linkages. The relative stiffness of the rods and wires determines the location of the combined center of motion.

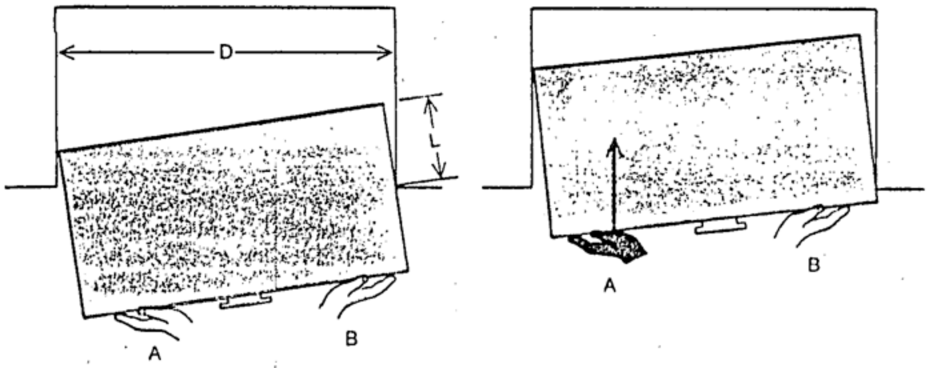

DIFFERENCE BETWEEN WEDGING AND JAMMING was clarified during the development of compliant-gripper mechanisms. When, for example, a bureau drawer becomes wedged (*left*), it is literally locked. Any further application of force will deform the drawer or the bureau or both. Theory shows that wedging arises when the drawer is inserted at such an off angle that the ratio of L/D is less than the coefficient of friction (μ) when two-point contact first occurs. The only remedy is to pull the drawer out and start again. If, however, the ratio L/D is larger than μ at the time of initial two-point contact (*right*), wedging cannot result, although further movement can be impeded by jamming. The remedy is to break the two-point contact by pushing at *A*, thereby changing the direction of both the applied force and the applied moment. Compliance devices on preceding page apply forces in accord with this theory.

the ratio is typically less than 0.2. There is no known remedy for wedging except to pull the peg out and try again.

Jamming, on the other hand, arises principally from the relation between the insertion force and the frictional forces. It can be remedied by changing the direction of the vector of the applied insertion force. Simunovic has derived a quantitative strategy that has been verified experimentally; it calls for sensing the applied forces and moving the arm under computer control to satisfy certain relations between the components of applied force. This approach is an effective and general one, but it could be expensive to implement in an assembly device. We and many other investigators have nonetheless found it useful in coping with balky drawers and window sashes in the home.

We conducted another illuminating experiment, using conventional industrial robots with no force feedback and relying only on their ability to repeat a taught sequence to within half a millimeter. Donald S. Seltzer of our group taught a robot to insert a crankshaft and a gear into the housing of a small gasoline engine, with clearances of about .05 millimeter, a tenth of the robot's accuracy. At first we were surprised by his success, but closer examination revealed the explanation. When a peg is partway into a hole, it can wobble back and forth a good deal farther than the clearance itself. Because of an incidental and unappreciated compliance in the arm and the grippers this wobbling was allowed to occur and the pieces went together. The contact forces were undoubtedly large.

Can all simple peg-and-hole insertion tasks be accomplished in this way? The answer is "Yes, but...." An engineer wants to be sure his machine will work. In this instance he wants to know how much wobble he can expect and how to arrange the compliances so that jamming will not occur. He cannot depend on the accidental compliance provided by the grippers. In order to improve our understanding of these effects we undertook an experimental and theoretical program of geometrically analyzing assembly tasks, combining the results with the jamming analysis. The aim was to create a passive compliant wrist that could execute the fine motions required for close-clearance insertions without the use of active sensors and actuators. The investigation was guided by Paul C. Watson, Samuel Drake and Simunovic.

Traditional time-study methods distinguish "easy" and "difficult" manual insertion tasks qualitatively, but we needed a description that was quantitative. We found it by determining how much a peg could wobble in a hole as a function of the depth of insertion, the diameter of the hole and the clearance. The results of the analysis are expressed in terms of ratios so that they will apply to all sizes of pegs and holes. The difficulty of the task is expressed in terms of the clearance ratio (the clearance divided by the diameter). We found that most parts of a given kind (washers or bearings, for example) are designed for a particular clearance ratio almost independently of their size.

This finding enables us to predict the difficulty of insertion for many industrial assembly tasks. A clearance ratio of .001 is typical of a fairly difficult insertion. In this instance a peg inserted to a depth equal to one diameter can wobble back and forth about .06 degree. The same peg, just entering the hole, can wobble about 1.5 degrees. We compared this result with the angular accuracy required to start screws into threaded holes without mismatching the threads on opposite sides of the hole and determined that a larger wobble was permissible. We concluded that machines accurate enough to perform most insertions can also install screws.

To test the geometric and friction analyses Drake, with the aid of Seltzer, conducted a series of controlled experiments with carefully measured round metal pegs and holes. The peg and the hole were mounted in a milling machine so that during the insertion of the peg precise relative errors could be imposed. The peg was attached to a sensor that measures forces and torques along three axes at the tip of the peg. To provide a known amount of compliance the top of the peg was attached to the sensor by a thin metal rod.

The experiment showed the need to improve the compliant rod in two ways. With the peg pivoted from its top there is a tendency for lateral error to become angular error as insertion proceeds. This tendency gives rise to two-point contact, and further insertion is possible only if the top of the peg can move laterally. The compliant member must be flexible enough to allow such lateral movement, otherwise large contact forces will be exerted on the tip of the peg and on the walls of the hole. When the rod is made too flexible, however, it tends to buckle and collapse.

The first step in solving such problems involved constructing a linkage device called a remote-center compliance, which allows the peg to rotate about its tip if it is angularly misaligned with the hole. A second linkage cascaded with the first allows the peg to move from side to side to correct a lateral error without introducing unwanted rotation.

The mathematical analysis of jamming can show that jamming is least likely, and insertion forces are lowest, when the peg is allowed to rotate about its tip. The combination of desirable properties—lateral motion, when it is needed, and rotation about the tip—have been built into a compliant mechanism of improved design that is not subject to buckling collapse. The new device, the first practical passive compliance, is now part of our laboratory assembly machinery and is being studied by several industrial firms. With it we have performed once-difficult assembly tasks, such as putting a bearing into a housing with a clearance ratio of .0004 in a fifth of a second, starting from a lateral error of one millimeter and an angular error of 1.5 degrees. The new device is clearly applicable to special-purpose automatic equipment and to programmable assembly systems.

The only other fine-motion device of which we are aware that can achieve close-clearance insertions is the Hi-Ti-Hand, developed by Hitachi, Ltd., in Japan. It is an active motorized unit that inserts pegs into holes by deliberately

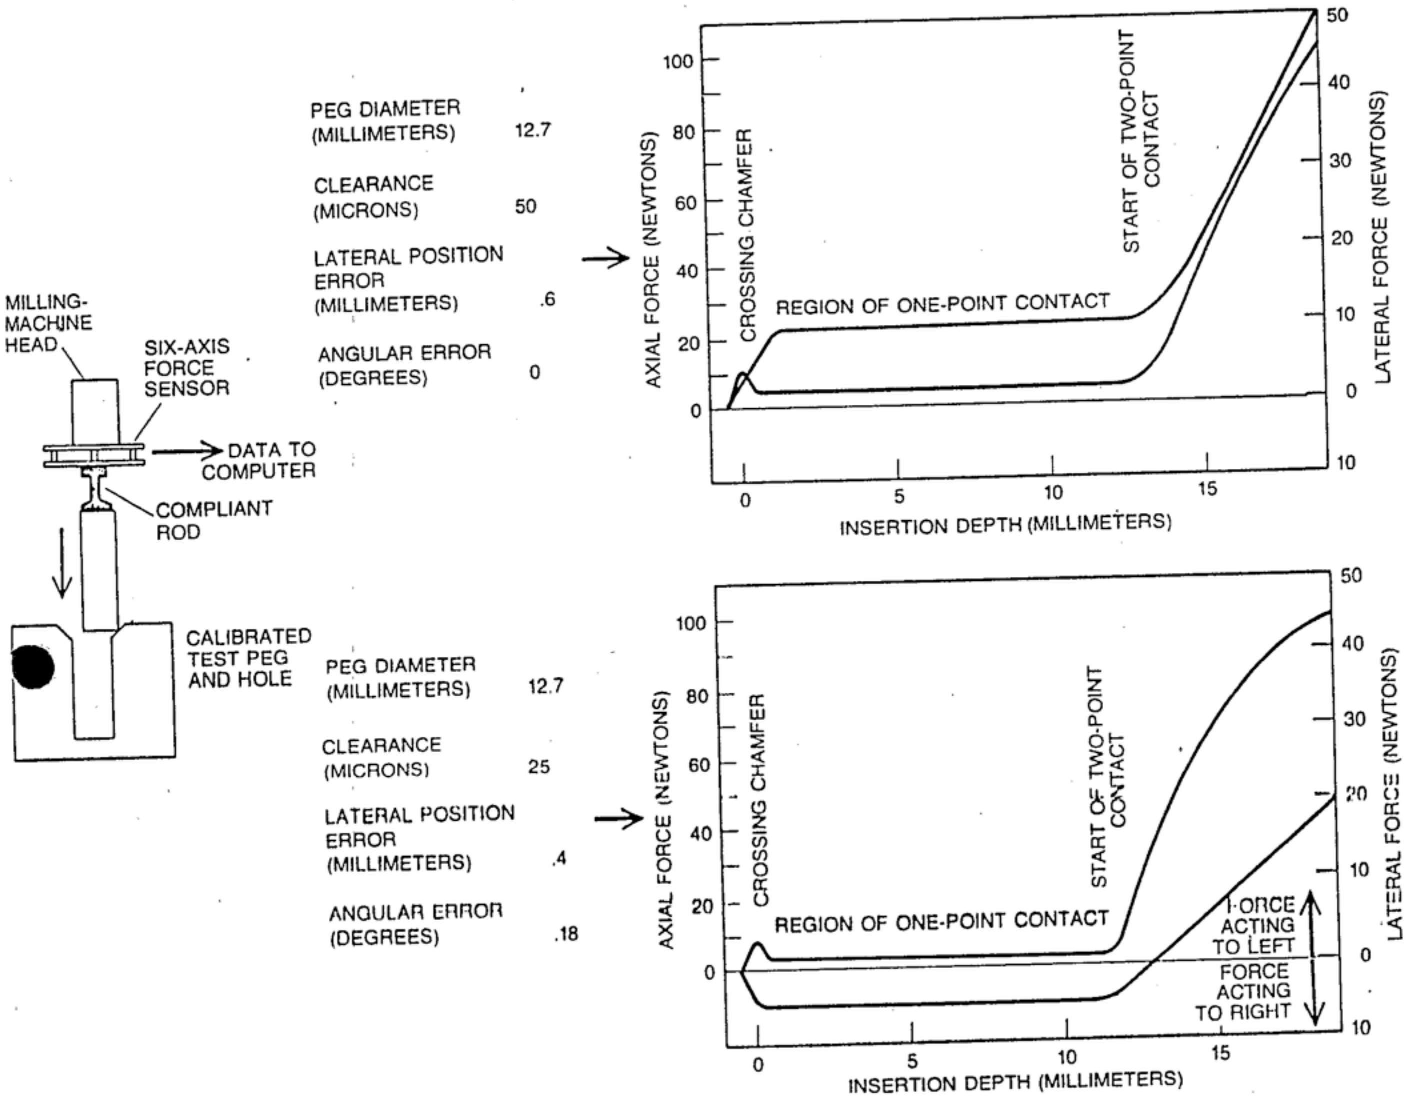

PART-MATING EXPERIMENTS depicted in this illustration disclose the forces that are developed when a steel peg is inserted into a hole under carefully calibrated conditions of clearance and misalignment. The forces and moments are detected and sent to a computer by a six-axis force sensor attached to the head of a milling machine. A simple compliant rod provides a limited amount of compliance between the force sensor and the test peg. In this pair of experiments the axial forces increase sharply by about the same amount after two-point contact takes place. The behavior of the lateral force in the second experiment (first pushing to the right and then to the left) is characteristic of combined lateral and angular error. The results of the part-mating experiments conformed to theoretical predictions.

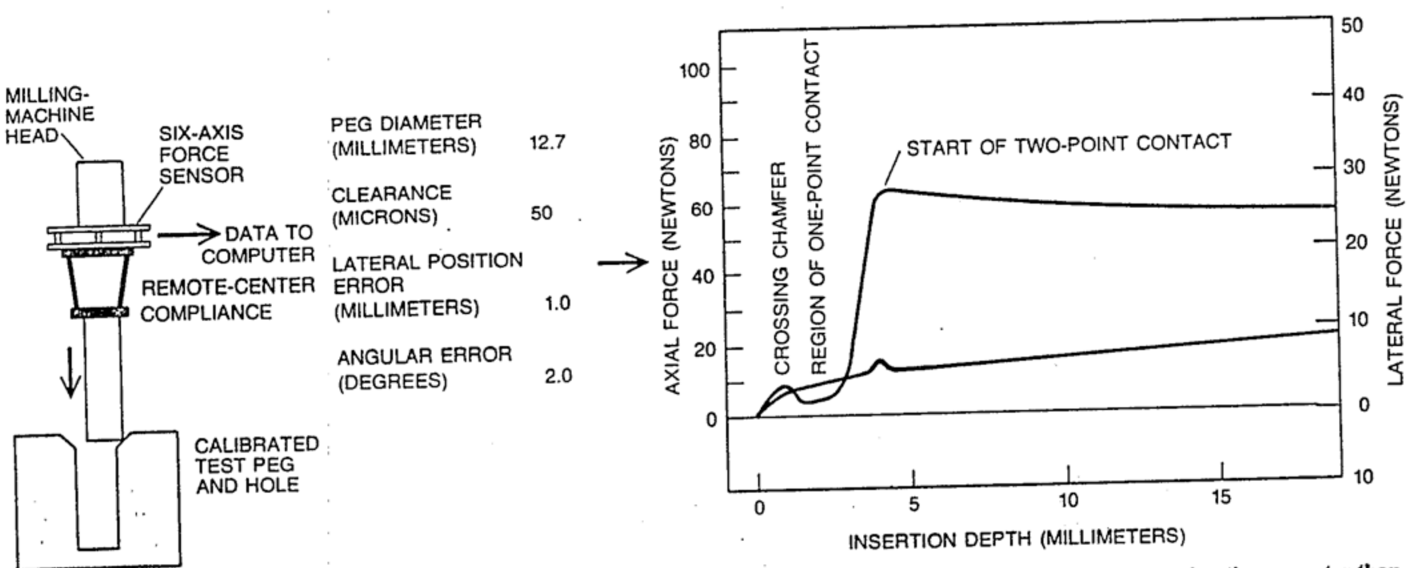

MUCH LARGER MISALIGNMENTS between a peg and a hole are tolerated and generate smaller axial and lateral forces when the peg is held by the newly developed remote-center compliance instead of by a rod. In this experiment the initial error in lateral position is about twice the value of the error in the compliant-rod experiments. The angular error of two degrees is more than five times greater than it is in the second rod experiment. The virtue of the new compliance scheme is that it allows the peg to rotate about its tip as it meets resistance entering the hole. The peg does not jam even though the two-point contact between it and the hole is made near the top of the hole.

moving the top of the peg from side to side and sensing with force sensors when, as a result of two-point contact, no further side-to-side motion is possible. In this way the device finds the boundaries of an imaginary funnel outside the hole defined by all the positions limited by two-point contact. Another version of the Hi-Ti-Hand pushes the peg down until two-point contact is made and then searches for a free area inside the funnel. The control system for detecting contacts and changing the peg's direction of motion is subject to the same constraints we established in our force-feedback studies. The gripper holding the top of the peg must be provided with compliant springs, and search speeds must be kept low to limit the contact forces. Current models of the Hi-Ti-Hand are able to perform only vertical insertion tasks and need several seconds to accomplish the kinds of insertion our remote-center compliance device does in a fifth of a second.

One other commercial system, developed by Olivetti, S.p.A., also incorporates force-sensing for assembling parts. The Olivetti system measures the forces along three axes as pieces come in contact with each other and makes what amounts to a sequence of binary decisions (part present or absent, hole present or absent) until mating is either complete or abandoned. The system has two arms and has been used for the machining and assembly of small parts.

The passive compliance assembly mechanisms designed in our laboratory represent the confluence of three endeavors: the assembly strategy based on sensing force vectors, the geometric analysis of assembly tasks and the theory of wedging and jamming. It seems to us that comparable studies of tasks other than insertions will lead to still other devices for aiding assembly. Let us now take up the kinds of problems that must be considered in designing a complete assembly system.

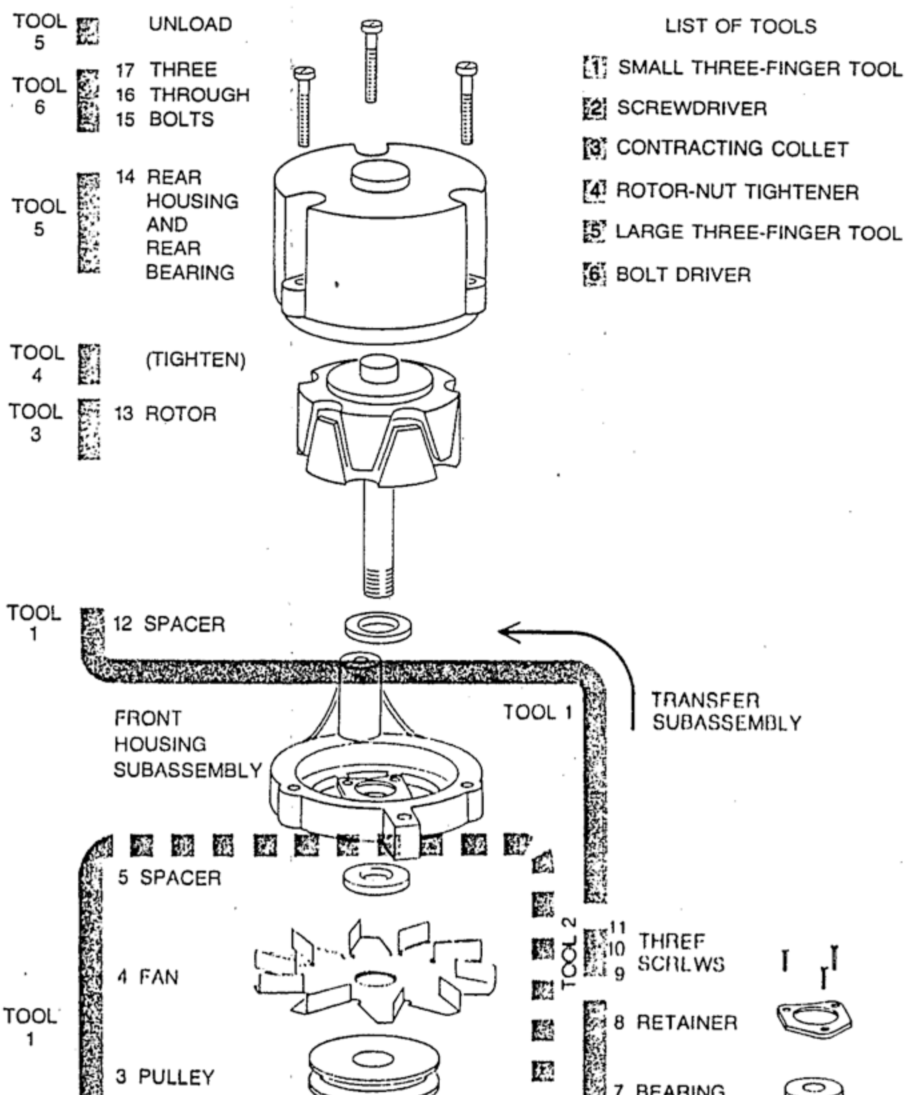

EXPLODED VIEW OF THE ALTERNATOR shows the sequence in which its 17 parts are assembled by the authors' programmable robot and identifies the tools that perform each task. The center rod in the main assembly fixture is a spring-loaded plunger. The collar at the base of the rod contains a hexagonal cavity that firmly holds part *1*, the nut. When the rotor (*13*) is inserted, the plunger is depressed, enabling the rotor's threaded shaft to engage the nut. The rotor is then spun by tool *4* in order to thread the shaft tightly into the nut. The authors calculate that time now required by entire operation, two minutes 42 seconds, could be reduced to one minute five seconds if changes were made in design of tools, fixtures and the alternator.

Assembly systems are collections of assembly machines capable of putting together one or more products. To design machines that can be widely applied in industry and to combine them into efficient systems requires an understanding of what industrial products are like, what assembly tasks arise, how small the clearances between parts are, how many models of one product are made and so on. It is also necessary to understand how assembly lines are arranged efficiently so that, for example, one machine does not lag behind the others. Finally, choosing products amenable to automatic assembly calls for thorough economic studies of what performance the new methods can provide and at what cost compared with present methods.

The first design decision is whether to use a manual system, a special-purpose machine or a programmable assembly machine, or perhaps some combination of all three. This problem is essentially an economic one as long as technically feasible designs can be formulated. It is common practice in industry to make such evaluations (between manual systems and special-purpose machines at present, of course) on a case basis, but we decided to devise a mathematical formulation of the important factors. Such a formulation makes it possible to name variables, to establish the relations between them and to explore the conditions, if any, that favor the use of programmable assembly machines. We particularly wanted to find out whether our original hypothesis about the applicability of programmable assembly machines in low-volume manufacturing was correct. This effort, carried out in collaboration with Paul Lynch, resulted in our economic model of assembly alternatives.

Before we had formulated our model we were beset by conflicting opinions and rules of thumb. A common opinion, the result of years of studying particular cases, is that an automated assembly process is more economic than manual assembly only if at least a million assembled units are needed each year for several years. How then could "low-volume programmable automated assembly" be economic?

Another opinion is that robot assembly arms should have six powered joints, since fewer joints would prevent the arm from making all possible rotational and translational (back and forth) motions as it positioned an object. Arms with fewer than six joints would inevitably be specialized; the manufacturer would have to make several stock types or he would always be building expensive special units rather than one mass-produced inexpensive arm that could "do everything." But should the customer pay for joints he does not need when mass-production quantities of robot arms are not likely to be achieved for many years? Without an economic model and a quantitative study of assembly tasks such arguments can continue indefinitely and inconclusively.

Let us take up the problem of economic modeling first; it will help to clarify the question of how many joints an arm should have. Our approach is necessarily simplified, but the main considerations are not greatly changed if the complexity is increased. Industry normally judges an investment (the purchase of a piece of equipment) in terms of how much money can be saved and how rapidly the savings pay back the investment. A straightforward way to evaluate the economics of an automatic assembly system is to compare the cost of assembling one unit manually with the cost of assembling the same unit automatically. The major cost in manual assembly is the hourly wage. This cost can readily be combined with the hourly production rate, the number of parts in the unit and the time it takes a person to add each part to the unit, from which the cost of assembling the complete unit can be found.

Since machines do not collect wages, a different approach must be taken with them. There are many such approaches, but a simple one will illustrate the principle. The company's management establishes a minimum period over which the machine must pay for itself, with

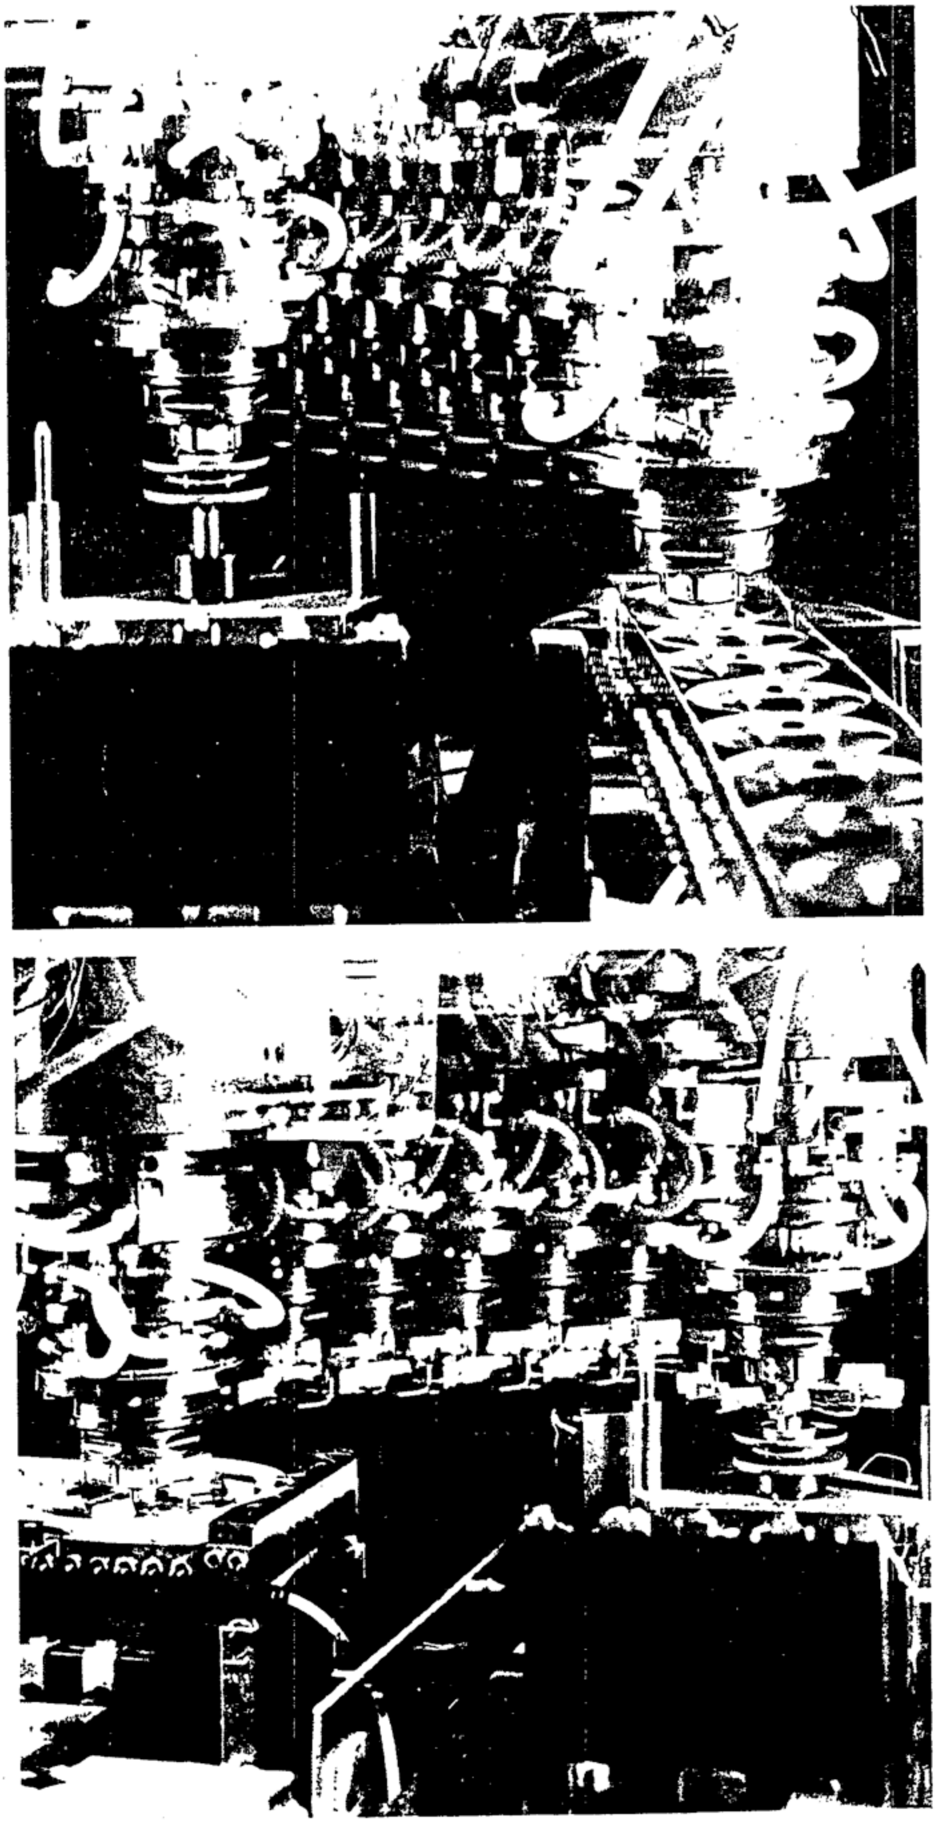

TYPICAL ASSEMBLY OPERATIONS are shown in these multiple-flash photographs taken in the authors' laboratory. At the top the robot arm is performing the third step in the assembly sequence in the illustration on the opposite page. With the aid of tool *1* (small three-finger tool) the robot is picking up a pulley from the feeder line at the right and is lowering it onto the plunger of the main assembly fixture. In the photograph at the bottom the robot, using the same tool, is picking up a fan from the feeder line at the left and placing it on top of the pulley.

shorter periods being the more desirable. Each of the units assembled during the selected payback period is allocated an equal share of the initial investment, thus creating what amounts to a "cost to assemble one unit." That cost will be higher if the payback period is made shorter. If the cost is less than the present or projected cost of manual assembly, it will be economic to buy the machine. Although other factors—taxes, interest rates, maintenance costs, salvage values and so on—enter in, this is the essence. It follows that if an economic model is to predict whether or not a given type of machine will be economic, the model must be able to predict the amount of the initial investment. This requirement in turn calls for modeling the amount of work that must be done to assemble each unit.

For the three broad alternatives—manual assembly, special-purpose-machine assembly and programmable-robot assembly—we have measured the amount of work simply by counting how many parts make up one unit. The investment for special-purpose machinery is assumed to be proportional to the part count, since each part calls for a separate dedicated work station. For small parts a station might cost $6,000; for large parts it might cost $200,000. This cost breaks down into three categories roughly as follows: machine parts and materials, 25 to 35 percent; design and construction, 30 to 50 percent, and adjusting the machine after it has been set up, 15 to 45 percent. The number of stations for a given job does not depend on the required production volume unless the volume is so large that several machines are needed. Hence economies of scale are realized if production volume is large, but the "cost to assemble one unit" is large if production volume is small.

The investment for a programmable machine, in contrast, is expected to depend significantly on the required production volume, because one station can perform several assembly tasks on one unit. Thus low production volumes can be met, with low investment, by a few stations each of which puts many parts on one unit. Larger volumes can be met by many stations each of which puts on a few parts. To be sure, grippers and feeders are needed to handle each part, so that some of the costs do not depend on production volume. Simple equations relating the cost of a robot station, the average time required to attach one part to a unit and the cost of tools and fixtures can predict the required investment and hence the unit assembly cost for a given payback period.

The models of the three assembly methods relate most of the recognized quantities: the work content of the task, the cost of the various techniques and their speed, the cost of labor and the required annual production. When we insert reasonable numerical estimates for these factors into our equations, we get curves of the type shown in the illustration on page 74. Although the curves represent a theory so far unverified, they do appear reasonable in several respects. First, they reproduce the conclusion that special-purpose machines become economic compared with manual assembly when production runs exceed a million units per year. Second, the curves show the economic benefits of programmability: the curve for programmable machines falls below that of special-purpose machines because fewer work stations are needed at low production volumes. This creates a range of production volumes below mass production where programmable machines have economic promise. Since the machines are reprogrammable, it is feasible to apply them to the assembly of products manufactured in families of models. The production volume of each model may be small, but the aggregate will be large enough to justify programmable assembly.

The equation for the costs of programmable assembly also resolves an old argument: whether to achieve economy by making assembly arms that are fast or to achieve it by making assembly arms that are cheap. The equations show that only the cost of the arm multiplied by the time needed to add one part is important. An equal percentage change in either factor has the same effect. Further study of specific arm designs has also shown that the most effective way to reduce the product of the two factors is to make a small, lightweight arm that has as few joints as possible.

This brings us back to the unanswered question of how many joints an arm should have, which is part of the larger question of matching the arm to the job. To obtain some basic information on assembly tasks our colleague Anthony S. Kondoleon took apart and reassembled a number of typical products, such as a bicycle brake, a refrigerator compressor, an electric jigsaw, a small electric motor and an electric toaster-oven. A key question concerned what might be called choreography: In what sequence are the parts attached to the product and from what directions do they arrive? The arrival direction is determined by the part's location on the product, how it fits into a hole, over a pin or into a slot, and what other parts block its path.

Kondoleon found that in all but one of the items he inspected about 60 percent of the parts arrived from one direction, 20 percent more arrived from a second direction opposite to the first, 10 percent arrived in a plane perpendicular to the first two planes and the remaining 10 percent arrived from various other directions. This means that the products are essentially stacks. They are made of cast, machined or molded materials and consist of a main body or frame to which parts are added. The only exception in the admittedly small sample of products Kondoleon studied was the toaster-oven, which was made of sheet metal, plastic and wires. Here no one arrival direction predominates and the interlocking wires and sheet-metal pieces make the product difficult to assemble.

Kondoleon also catalogued the various kinds of assembly tasks required for each product, their relative frequency and the directions of approach of the parts associated with those tasks. Simple peg-hole insertion tasks were not only the most frequent overall but also the most frequent along the dominant direction of approach. The installation of screws was the second most frequent task, but interestingly enough most of the screws were inserted in directions perpendicular to the dominant one. Many other tasks were identified, but none was dominant in any particular direction.

Since the dominant directions form an *xyz* right-angled coordinate system, it can be concluded that an arm with just enough joints to move along three perpendicular directions can perform all the gross motions required to move parts together for typical "stack" products. Of course, additional degrees of freedom, although they could be limited to fine motion, would still be needed to carry out the total assembly task. When this general finding is combined with the compromises in arm cost and speed indicated by the economic analysis, our conclusions are that stack products offer a promising area in which to apply programmable assembly and that arms whose gross motions follow rectangular or cylindrical coordinates would be well suited to the task.

As part of our program we have recently built and operated an adaptable programmable assembly system in our laboratory. It consists of a computer-controlled industrial assembly robot with four degrees of freedom. The wrist of the robot contains an advanced version of our passive-compliance insertion device. The system has been taught to assemble a Ford automobile alternator, using six interchangeable tools. The alternator has 17 parts, all of which can be inserted from one direction with the aid of two assembly fixtures.

The system was built in order to obtain data on adaptable programmable assembly and to test our parts-mating concepts on an actual industrial prod-

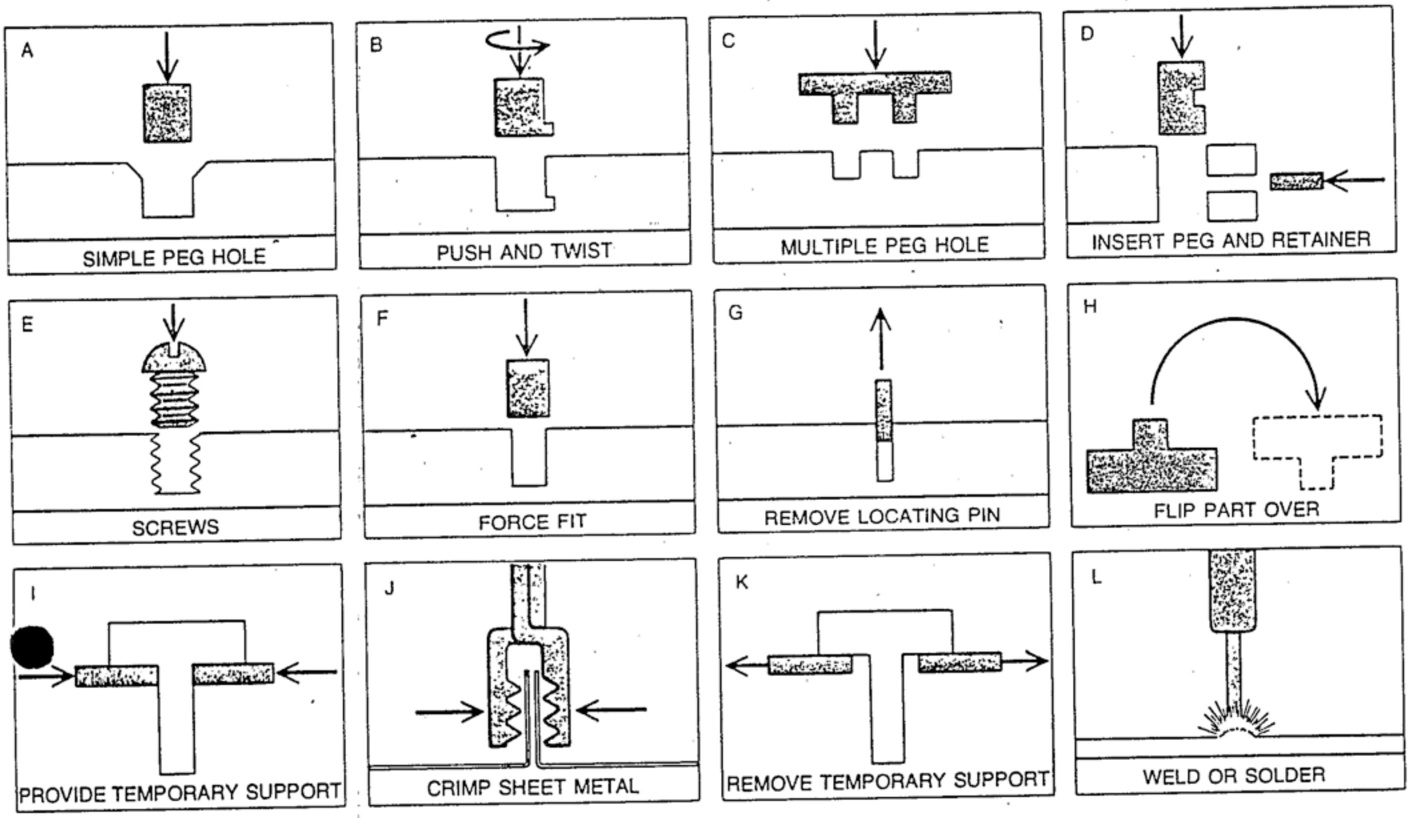

TYPICAL MANUFACTURING TASKS were identified by taking apart and reassembling a variety of products and their components, including a refrigerator compressor, an electric jigsaw, an induction electric motor, a toaster-oven, a bicycle brake and the automobile alternator used in robot assembly project. All the items could be assembled with various combinations of 12 operations depicted.

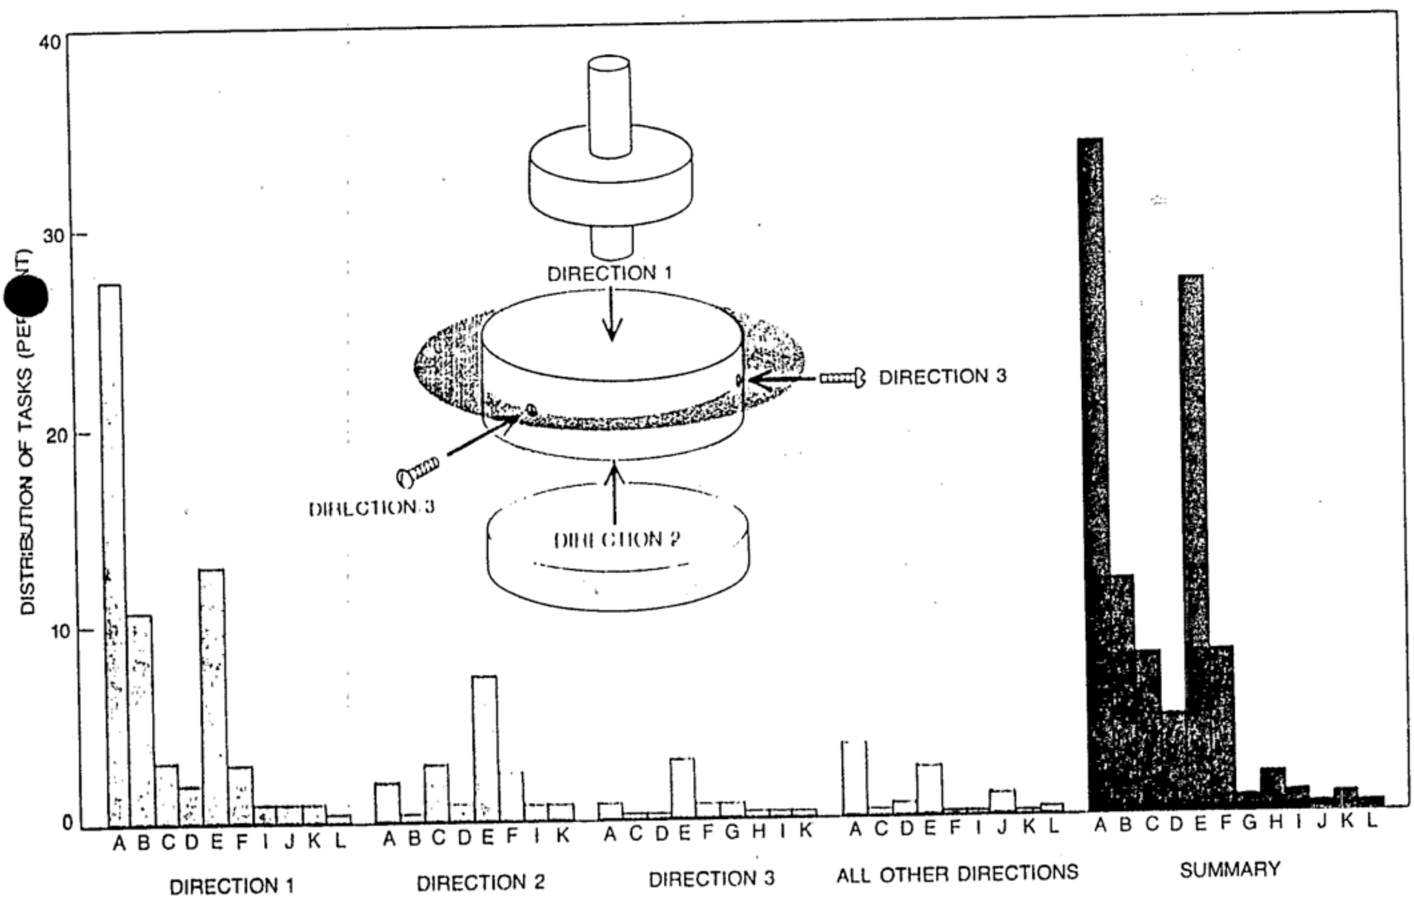

DIRECTION OF ATTACHMENT OF PARTS was analyzed in the study that classified the types of operation required. The inset diagram defines the three principal directions of attachment. Direction *1* is dominant, followed by direction *2* and direction *3*. Direction *3* is any direction perpendicular to the other two. The bar graph correlates attachment direction with type of task involved according to the identification in the illustration at the top. Simple peg-and-hole tasks (*A*) outnumber all others, followed by insertion of screws (*E*).