ROBOTIC SENSORS IN PROGRAMMABLE AUTOMATION

Technical Note 183

March 1979

By: David Nitzan
Program Manager, Industrial Automation
Artificial Intelligence Center
Computer Science and Technology Division

333 Ravenswood Ave. • Menlo Park, California 94025
(415) 326-6200 • Cable: SRI INTL MPK • TWX: 910-373-1246

ABSTRACT

There exist both a social and an economic need for the advancement of automation in general, and of programmable industrial automation in particular. Characterized by flexibility and the ease of setup for new production tasks, programmable automation employs industrial robots. Today's robots possess "muscles" only; there is a need to develop intelligent robots that can detect faults and correct errors by using sensors and computer control. Intelligent robots may have contact sensors (sensing force/torque, touch, position, etc.), noncontact sensors (sensing visual images, proximity, range, etc.), or both. Such sensors are applicable to three basic functions: inspection, finding objects, and robot-control feedback. Fast reaction to multisensory data can be achieved by using a distributed network of microcomputers that process the sensory data in parallel and control each effector in a modular fashion. Application of sensor-controlled manipulation to material handling, inspection, and assembly tasks has been demonstrated in several laboratories in the United States, Japan, and Europe.

CONTENTS

I INTRODUCTION

A. The need for Advanced Automation

The need for advanced general automation is based on both social and economic imperatives. Let us examine them briefly.

1. Social Need

The social objective of advanced automation comprises two major categories: raising the standard of living of the entire population and improving the working life of the labor force.

Advanced automation can enhance the standard of living in several ways. Specifically, it can:

- Increase purchasing power by reducing the rate of inflation.
- Improve the quality of goods by minimizing human errors.
- Permit customized products to be sold at prices approximating those of mass-produced products by applying programmable automation, as we shall see later.
- Reduce waste and scrap by computer-aided design and manufacturing.

Advanced automation can improve working life in the following ways:

- Creation of a better work environment.
- Enabling human workers to perform jobs that are less harmful, dangerous, strenuous, unpleasant, or dull.
- Reduction in the number of working hours per week to counter the unemployment problem that may result from automation advancement.

2. Economic Need

The economic justification for advanced automation falls into two main categories: combating inflation and the competition for world markets.

In recent years, labor costs in the industrialized countries (USA, Japan, and Western Europe) have been rising faster than labor productivity. For example, based on statistics provided by the U.S. Department of Labor, the average rise in wages from 1965 to 1970 in these countries was 9.1%, while the average gain in labor productivity was only 6.3% (see data in [Merchant, 1975]). This difference has had an adverse effect on inflation.

The next question we must ask is: which sector of industry is most in need of advanced automation? Since automation in the processing of homogeneous products (such as steel, oil, and chemicals) is well advanced, we are concerned in this paper exclusively with discrete-part manufacturing.

Consider the effect of production rate (units per year) on the relative cost per part. If the rate is low (less than, say, 50 units per year), automation is not justified and the cost per part is inherently high. On the other hand, mass production (more than, say, 100,000 parts per year) justifies the high cost of hard automation, the type of automation that employs special-purpose equipment; consequently, the cost per part is much lower than for "one-of-a-kind" parts.

Between these two extremes lies batch production, which justifies little automation, has poor machine utilization, and is very labor-intensive. It was estimated [Cook, 1975] that 50% to 75% of the total value of discrete-part manufacturing falls into this category. The economic incentive exists, therefore, to develop a new automation technology for batch production. It will conceivably bring the cost per part close to the equivalent figure in mass production. This technology is called programmable industrial automation.

B. Programable Automation

Programmable automation is characterized by flexibility and ease of training [Nitzan and Rosen, 1976]. Flexibility is the capability of a machine or a system to perform automatically a variety of production tasks, not just a single task. Ease of training is the facility with which factory personnel can program (or teach) the system to perform the desired tasks.

The two most important programmable automation systems existing today are the Numerical Control (NC) machine tool, which is used to fabricate parts, and the industrial robot, which is used to handle and assemble parts.

C. Today's Industrial Robots

Within the last 15 years, programmable manipulators, known as "industrial robots," have been introduced into industry to replace workers doing undesirable jobs. The origin of these robots can be traced to two sources: (1) the use of teleoperators, that is, human-controlled "arms" and "hands" for handling radioactive materials, and (2) NC machine tools.

An industrial robot consists of an "arm" (driven by hydraulic, pneumatic, or electric power) and a "hand" (driven by pneumatic or electric power). The hand is usually a gripper, but may also be a suction cup or a tool, such as a welding gun or a paint sprayer. The arm has several degrees of freedom realized by rotary and sliding joints, mostly rotary. A six-joint manipulator is shown schematically in Figure 1, where θ_1, θ_2, s_3, θ_4, θ_5, and θ_6 are joint coordinates, O_H is the position of the hand center, and REACH, LIFT, and SWEEP are hand coordinates that determine orientation of the hand in room coordinates. Three degrees of freedom are sufficient to position the hand center within the working space; three additional degrees of freedom are sufficient to orient the hand in any direction.

An industrial robot can be programmed to perform a variety of tasks. During the training mode a factory worker, using a "teach box," leads it through a sequence of positions and records them in its memory. During a subsequent replay mode these positions are retraced by servoing the robot within a specified accuracy. For most jobs (such as material handling in die casting, forging, spot welding, and palletizing) industrial robots move from point to point, but for some jobs (such as arc welding and paint spraying) they move along continuous paths.

Industrial robot installations may be divided into three categories (Engelberger, 1975]):

(1) Robot is surrounded by work, as in forging and trimming, press-to-press transfer, plastic molding, and investment casting.

(2) Work comes to the robot, as in spot welding, arc welding, and palletizing.

(3) Robot goes to the work, as in loading and unloading a group of NC machine tools.

There are thousands of industrial robots in operation today, primarily in the United States, Japan, and Europe. Most numerous are the simple limited-sequence manipulators (such as Auto-Place) with two to five joints -- at a cost ranging from $3,000 to $15,000 each. Their joints are not servoed; instead, each moves between two fixed but adjustable stops. More sophisticated robots, storing multistep programs, have five or six joints that are all servoed with positional sensory feedback. The present cost of an industrial robot is between $50,000 and $75,000; the cost of the indexing equipment, however, may be as much as three times that of the robot itself. Hence there is a need to develop methods for path control of robots performing material-handling and assembly tasks without the use of jigs or indexing fixtures in batch manufacturing.

There are three principal reasons for the fact that more and more industrial robots are entering the manufacturing work force:

(1) An industrial robot is programmable; this attribute is essential for automation in batch manufacturing.

(2) For some jobs, industrial robots are more economical than human workers. This is especially true if they operate in more than one shift. A return-on-investment of one to two years is not uncommon.

(3) Government regulations for safety, health, and environmental protection are becoming stricter in many countries; industry has no choice but to replace human workers in certain undesirable jobs with machines.

On the other hand, today's robots are subject to some technoeconomical limitations as compared with human workers. This is a problem we shall now discuss.

II THE NEED FOR INTELLIGENT ROBOTS

A. Intelligent Robot

A manufacturing process is basically a stochastic process; namely, it entails events that vary randomly with time. To cope with these events a person uses his arms and hands (with or without a tool), his senses, and his brain. To emulate these human capabilities the robot must be intelligent. By "intelligent robot" is meant a robot consisting of three components: a servoed manipulator (arm and end-effector), sensors, and computer control. Today's robot is not intelligent -- at least, not yet.

B. Adaptive-Control Functions

An intelligent robot uses sensors and a computer to perform two main adaptive-control functions:

(1) Correction of inherent errors in the position and orientation of workpieces, and of the robot end-effector.

(2) Detection of potential faults and minimization of their effects.

Let us examine these adaptive-control functions.

1. Correction of Errors in Position/Orientation

There are two basic approaches for correcting errors in the position and orientation of workpieces in manufacturing. The first approach is to fix them on pallets or in jigs. This is the method being used with today's robots; it is usually costly and may be justified in mass production, but not in batch production. The second approach is to transport workpieces randomly on moving conveyors or in bins, and to find and acquire them, as required, at each work station. This second approach is less costly and therefore common in industry, especially in batch manufacturing. At present this approach is being implemented by people who make use of their muscles, senses, and brains to find, acquire, transport, and position the workpieces. If we want robots to perform these tasks, we should provide them with comparable means for sensing and decision making.

There are also two approaches to minimizing position and orientation errors of the hand or of the robot's tool. The first approach is to achieve a high degree of accuracy by using a rugged high-resolution robot. This approach, which has been applied to NC machines, is very costly. A second approach is to utilize a less costly robot with lower accuracy, but to compensate for this drawback by using sensors. For example, a robot may be commanded by its computer to move until a touch sensor is activated. We found that the positional repeatability of one commercial robot obtained in this way is about 20 times better than the nominal repeatability.

Finally, let us consider the errors entailed in manipulation of workpieces, such as in loading and unloading them from fixtures in machine tools, presses, etc., or presenting and fitting them in assembly. Here too we distinguish between two approaches to overcoming these errors. The first approach is passive accommodation, for which no sensors are used. This approach amounts to simple open-loop control and, while usually fast, it does require initial alignment of the workpieces. The second approach, called active accommodation, employs sensory feedback under closed-loop computer control. For example, a robot may find a hole with a visual sensor and then insert a peg into that hole by moving its hand in a direction that minimizes the sensed components of force and torque. Although the second approach may be slower than passive accommodation, it does not require initial alignment. An optimal solution would probably be to combine these two approaches.

2. The Need for Sensors to Detect Faults

Let us now consider the need for sensors to detect faults. Here I include collision avoidance and in-process inspection. Sensors can be used to prevent, or minimize the effect of, collision between the hand, tool, load, or arm links and people, equipment, or workpieces.

In-process inspection of workpieces and subassemblies is necessary while they are being handled throughout the manufacturing process. Here we distinguish between inspection of dimensional tolerance, surface quality, and part integrity (Is the part missing? Is it disoriented? Is it the wrong part? Is it damaged?).

III CONTACT AND NONCONTACT SENSORS

We distinguish between contact and noncontact robotic sensors [Rosen and Nitzan, 1977]. Contact sensors sense force, torque, and touch while in physical contact with objects. Noncontact sensors sense visual image, range, and the presence of objects without making any physical contact. Let us examine each of these types of sensors.

1. Contact Sensors

Contact robotic sensors include force/torque sensors, touch sensors, and position sensors.

a. Force/Torque Sensors

Force/torque sensors are intended primarily to measure the three-dimensional components of force and torque being applied by a robot end-effector to a workpiece. Three types of force/torque sensors can be used for this purpose:

(1) A sensor measuring the (x, y, z) components of force and torque acting on the robot joints and added vectorially. For a joint driven by a dc motor, sensing is done by measuring the armature current. For a joint driven by a hydraulic or pneumatic drive, sensing is done by measuring the differential pressure. This method requires no special transducer, but is relatively inaccurate because the measured forces include those that are not transmitted to the end-effector, such as joint friction and the weights of the arm links.

(2) A wrist force sensor mounted between the last link of the robot and its hand. It consists of an elastic structure and transducers, such as strain gages, that measure the deflections of the structure sections due to the applied force.

(3) A pedestal force sensor that provides a base for assembly operation and measures the components of force and torque applied to a workpiece on the pedestal. It also consists

of an elastic structure and strain gage transducers. Such a device was built by Watson and Drake (1975), but did not have sufficient resolution for some assembly operations.

Because of the limitations of the other types of force sensors, we will consider here only the wrist force sensor.

Figure 2 shows the structure of a wrist force/torque sensor built at SRI International. It was milled from a 3-inch-diameter aluminum tube, forming eight narrow elastic rods whose deflection has no hysteresis. Two foil strain gages (shown as black rectangles) are cemented at one end of the rod. The strain here is increased by constricting the other end of the rod. The two strain gages are connected differentially to a potentiometer circuit whose output voltage is proportional to a force component normal to the plane of the strain gage. The differential connection of the strain gages provides automatic compensation for variation in temperature. Since the eight pairs of strain gages are oriented normal to the x, y, and z axes, the three components of force F and three components of torque M can be determined by properly adding and subtracting, respectively, the output voltages. Specifically:

$$F_x \propto P_{y+} + P_{y-}$$

$$F_y \propto P_{x+} + P_{x-}$$

$$F_z \propto Q_{x+} + Q_{x-} + Q_{y+} + Q_{y-}$$

$$M_x \propto Q_{y+} - Q_{y-}$$

$$M_y \propto -Q_{x+} + Q_{x-}$$

$$M_z \propto P_{x+} - P_{x-} - P_{y+} + P_{y-}$$

A simpler, uniaxial version of the above force may be used to measure the gripping force with which a robot hand holds a workpiece. This measurement is important to assure that the gripping force is above a slippage value and below a distortion value (especially if the workpiece is delicate).

A modular force sensor was described by Wang and Will (1978). It consists of interchangeable modules, each consisting of a pair of strain-gage blocks and an intermediate block along the x, y, or z axis. Unlike most force/torque sensors that are installed between the arm and the end-effector, this sensor is installed at the base of each gripper, thus reducing errors due to end-effector inertia.

b. Touch Sensors

Touch sensors are used to obtain information associated with the contact between the finger(s) of a manipulator hand and objects in the work space. They are normally much lighter than the hand and are sensitive to forces much smaller than those sensed by the aforementioned force sensors.

Touch sensors may be mounted on the outer and inner surfaces of each finger. The outer sensors may be used to search for an object and possibly determine its identity, position, and orientation. Outer sensors may also be used for detecting unexpected obstacles and stopping the manipulator before any damage can occur. The inner mounted sensors may be used to obtain information about an object before it is acquired and about the location of grasping points as well as any workpiece slippage during acquisition.

Touch sensors may be classified into two types, binary and analog. A binary touch sensor is a contact device, such as a micro switch. Being binary, its output is easily incorporated into a computer controlling the manipulator.

Garrison and Wang (1973) built a gripper with 100 pneumatic binary touch sensors located on a grid with 0.1- by 0.1-inch centers. The sensors consisted of contact terminals, a thin metal sheet with elastic shallow spherical domes, and a flexible insulating rubber sheet on the outside. Physical contact was sensed whenever external pressure exceeded a preset threshold, causing activation of a snap-action switch consisting of a dome and a terminal.

An analog touch sensor is a compliant device whose output is proportional to a local force. This sensor is usually mounted on the inner surface of the fingers to measure gripping forces and to elicit information about the object between the fingers.

Hill and Sword (1973) built a manipulator hand with a wrist force sensor and analog touch sensors; the hand is shown in Figure 3. Seven outer sensing plates and a matrix of 3-by-6 inner touch sensors are mounted on each finger (or jaw). The force on each sensor acts against a compliant washer, displacing a vane that controls the amount of light received by a phototransistor from a light-emitting diode.

Contact sensing is still in a highly experimental stage. As yet there is no commercial line of contact sensors other than binary switches that have been proved in industrial application. At the same time, it has become quite evident to manipulator users and developers alike that contact sensing will be a valuable addition to programmed manipulation in the near future.

c. Position Sensors

The absolute position of a robot end-effector in the room coordinate system can be computed from the joint coordinates of the robot. The latter can be measured by various transducers, such as optical encoders, resolvers, and potentiometers. From the viewpoint of performing a given operation (e.g., grasping), however, the position of the robot end-effector relative to the workpiece is more important than the absolute position of the end-effector in room coordinates. Such relative position can be sensed by a linear potentiometer whose output voltage is proportional to the extension of a movable end-effector in contact with the workpiece. This position sensor, for example, can be used to verify that a bolt has been inserted and threaded properly in a bolt hole [Rosen and Nitzan et al., 1978].

2. Noncontact Sensors

Noncontact sensors are useful in identifying and finding parts in sensor-controlled manipulation and for visual inspection. The major categories of noncontact sensors that have been used with robot systems are electro-optical imaging sensors, proximity sensors, and range-imaging sensors. These sensors are described separately below.

a. Electro-Optical Imaging Sensors

Electro-optical imaging sensors provide the most commonly used "eyes" for industrial robots and visual inspection. Solid-state TV cameras, interfaced with a computer, constitute the least expensive and most easily available imaging sensors. These cameras scan a scene, measure the reflected light intensities within a raster of, say, 128 x 128 picture elements (pixels), convert these intensity values to analog or binary electrical signals, and feed this stream of information serially into a computer within 1/60 of a second. These signals may either be stored in the computer's memory for subsequent processing or processed in real time "on the fly," with consequent reduction of memory requirements.

A uni-dimensional solid-state camera, using a linear-diode array that varies from 16 to 1872 elements, is also available commercially. This device can perform a single linear scan and is very useful for sensing objects that are in relative motion to the camera, such as workpieces moving on a conveyor belt. An example is shown in Figure 4, where a connecting rod moves past the viewing station and top-view and side-view linear scans are performed by two linear-diode arrays, each scan initiated by a repetitive signal from a position sensor (incremental encoder) coupled to the moving conveyor belt.

b. Proximity Sensors

A proximity sensor is a device that senses the presence of an object without making physical contact. Major types of proximity sensors now available commercially are based on radio frequency, a magnetic bridge, ultrasound, a permanent-magnet hybrid, photoelectric cells, and the Hall effect. These noncontact sensors have widespread use, such as for high-speed counting, protection of workers, indication of motion, sensing the presence of ferrous and nonferrous materials, level control, reading of coding marks, and noncontact limit switches.

The modern photoelectric proximity sensor appears to be well adapted for controlling the motion of a manipulator. This sensor consists of a solid-state light-emitting diode (LED), which acts as a transmitter of infrared light, and a solid-state photodiode, which acts as a receiver; both are mounted in a small package. As shown in Figure 5, the sensing space of a proximity sensor is approximately the intersection of two cones in front of the sensor. This sensor is not a range finder, as the received light is not only inversely proportional to the distance squared, but is also directly proportional to the target reflectance and the cosine of the incidence angle, both of which may vary spatially. However, if the reflectance and incidence angle are fixed, the distance may then be inferred with suitable calibration. As a rule, a binary signal is generated when the received light exceeds a threshold value that corresponds to a predetermined distance. Such devices are sensitive to objects from a fraction of an inch to several feet in front of the sensor.

c. Range-Imaging Sensors

A range-imaging sensor measures the distances from itself to the raster points in the scene. Although range sensors are used for navigation by some animals (e.g., the bat), hardly any work has been done so far in applying range imaging to manipulator path control. This situation may change in the future, however, as the technological and economical difficulties currently involved in the use of range-imaging sensors are overcome.

Various range-imaging sensors have been applied to scene analysis in a number of research laboratories. These sensors may be classified into two types -- one based on the trigonometry of triangulation, the other on the flight time of light or sound.

Triangulation range sensors are further classified into two schemes -- one based on a stereo pair of TV cameras (or one camera in two successive locations), the other on projection of a sheet of light by a scanning transmitter and recording of the reflected-light image by a TV camera. Alternatively, the second scheme may involve transmitting a light beam and recording the direction of the reflected light by a rocking receiver. The first scheme suffers from the difficult problem of finding corresponding points in the two images of the scene. Both schemes have two main drawbacks: missing data for points seen by the transmitter but not by the receiver (and vice versa), and poor accuracy for points that are too distant.

The above drawbacks are eliminated by the second type of range-imaging sensor, using a laser scanner, which is also classified into two schemes: one based on transmitting a laser pulse and measuring the arrival time of the reflected signal [Johnston, 1973], the other based on transmitting an amplitude-modulated laser beam and measuring the phase shift of the reflected signal [Nitzan et al., 1977]. A simplified block diagram of the latter sensor is shown in Figure 6. The transmitted beam and the received light are essentially coaxial.

IV APPLICATIONS

A. Application Functions

Sensor applications may be broadly divided into three functional areas: finding parts, performing visual inspection, and controlling manipulation.

1. Finding Parts

Where fixed or hard automation is justified for high-volume mass production, workpieces must be positioned and oriented with considerable precision, usually at high cost for special jigging. For material-handling and assembly operations in the unstructured environment of batch manufacturing, on the other hand, it will probably be necessary to "find" workpieces -- that is, to determine their positions and orientations and sometimes to identify them as well. It is possible to preserve workpiece orientation throughout the manufacturing process by suitable jigging or special palletizing. As mentioned previously, however, the cost entailed in this approach may not be justified, especially if batches are small or product modifications are frequent.

For example, present-day industrial robots cannot cope with the problem of picking up parts one at a time from a bin containing many randomly oriented parts. Such bins are used in many factories for temporary storage and for transportation of parts from station to station. It is highly unlikely that a very expensive replacement for this function would be implemented in most factories, especially for small batches or for parts that emerge from processes that inherently produce disorder, such as tumble polishing and plating.

To enable them to determine the identity, position, and orientation of parts and to perform visual inspection as needed, existing robots must be augmented with visual sensors. Such sensors may be TV cameras and/or several types of scanning systems that measure intensity and range data.

2. Visual Inspection

Here we are concerned solely with an important subset of visual inspection: the qualitative and semiquantitative type of inspection performed by human vision, rather than by measuring instruments. Such inspection of parts or assemblies includes part identification; detection of burrs, cracks, and voids; examination of cosmetic qualities and surface finish; counting the number of holes and determining their approximate locations and sizes; ascertainment of completeness of assembly; and so on. Sensory methods being developed for augmenting industrial robot systems can also be applied effectively to inspection requiring accurate mensuration.

A large number of factory workers are engaged in visual inspection. Explicit inspection, performed by workers whose sole job is to inspect parts, subassemblies, and assemblies, has been estimated to comprise approximately 10 percent of the total labor cost for all durable goods [Nevins et al., 1976]. This cost is second only to that of assembly operations, which is approximately 22 percent of the total cost. It is fair to assume that the majority of these inspection tasks are done visually.

Implicit visual inspection is performed by assemblers to ascertain that the assembled workpieces are indeed the correct ones, are complete, and have not been damaged. This task represents a small but essential part of the assembly function. Adding defective or wrong workpieces to an assembly that may have acquired a considerable value will yield costly scrap or require expensive correction later.

The wide variety of significant characteristics that are routinely examined visually by humans indicates the complexity of the processing that must be performable by automated systems. It is evident that a large library of computer programs will have to be developed to cope with the numerous classes of inspection. To avoid lengthy and costly programming for every new inspection job, this library must be made available to the broadest possible spectrum of manufacturing firms.

3. Controlling Manipulation

Manipulation of workpieces and tools for material handling and assembly jobs involves many basic operations, such as grasping, holding, orienting, inserting, aligning, fitting, screwing, turning, and so on. In a completely structured environment it may be possible to perform all these operations in a feed-forward manner, with no need for sensory control or correction.

It is instructive to note that human manipulation, being imprecise, depends almost entirely on sensory feedback to control both simple and complex manipulative operations. In general, the human worker makes use of both noncontact (visual) sensing and contact (touch) sensing. There is little doubt that a blind worker can perform many manipulative tasks using his tactile sense alone. However, he cannot easily perform most inspection functions, cannot readily cope with unexpected intrusions into his work space, and would generally require a far more structured situation to equal the performance of a sighted worker.

On the basis of these observations, it appears useful to consider the use of both contact and noncontact sensors in manipulator control and to try to determine where each sensor is most appropriate. One approach is to use noncontact sensors for coarse resolution, contact sensors for fine resolution. For example, in acquiring a workpiece that may be randomly positioned and oriented, a visual sensor may be used to estimate the relative position and orientation of the workpiece, say, to one tenth of an inch. From this information the manipulator can be positioned automatically with a precision matching that of the visual sensor. The somewhat compliant fingers of the manipulator hand, bracketing the workpiece, will now be close enough to effect closure, relying on touch sensors to stop the motion of each finger when a specified contact pressure is detected. After contacting the workpiece without moving it, the compliant fingers have flexed no more than a few thousandths of an inch before stopping. The touch sensors have thus performed fine-resolution sensing and have compensated for the lack of precision of both the visual sensor and the manipulator. This task, which is quite common, illustrates the relative merits of each sensory modality and the advantages of using both. It appears likely that the combined use of these sensors and the associated computer hardware and software will soon be cost-effective.

Other common applications for contact sensors, which implicitly entail fine resolution or precision sensing, include:

- Collision avoidance, using force sensors on the links and hand of a manipulator. Motion is quickly arrested when any one of preset force thresholds is exceeded.
- Packaging operations in which parts are packed in orderly fashion in tote-boxes. Force sensors can be used to stop the manipulator when its compliantly mounted hand touches the bottom of the box, its sides, or neighboring parts. This mode of force feedback compensates for the positional variability of the box and the parts and for the small but important variability of the manipulator positioning.
- Insertion of pegs, shafts, screws, and bolts into holes. Force and torque sensors can provide feedback information to correct the error of a computer-controlled manipulator. Again, one may first use visual sensors with relatively coarse resolution to find the hole, bring the peg to an edge of the hole (perhaps partially inserted), and then align the peg with the hole by moving it in a direction that minimizes the measured binding force and torque.

Reducing the field of view makes it possible to increase the resolution of noncontact sensors to a level approaching that of contact sensors. This method may be too slow because of the large number of fields required to cover a given area of interest and the excessive amount of computation. But this limitation may be overcome by using a fixed, wide field of view to find the target and a movable, narrow field of view to obtain high resolution. Alternatively, mounting a small optical sensor on the hand of a manipulator can provide reasonably high resolution over the small but important field of view close to where it is most needed. It would then be possible for the manipulator to follow especially identified lines, or to control its motion in accordance with the location of holes or fiduciary marks. It is likely that this "eye-in-hand" mode of operation can be applied successfully to situations that require semiprecise positioning.

B. Hardware/Software Modular Subsystems

To implement the application functions described above in an efficient and orderly manner, the sensor-equipped robotic system should be organized in a modular fashion. Such a system has been developed at SRI, for example. The major components of this system are described next.

1. Manipulator Path Control

We have developed analytical techniques and a library of software subroutines to control the trajectory of a manipulator, including:

(1) Transformations between Cartesian and joint coordinates of a Unimate arm.

(2) Transformations between different Cartesian coordinate frames, such as frames attached to the following: the Unimate base, its end-effector, a work station, a moving conveyor, and a workpiece.

(3) Smooth path control for arbitrary trajectories.

2. Machine Vision

To sense the image of a stationary object, we have used a solid-state TV camera (General Electric 100 x 100-element Model TN-2200); to sense the image of a moving object, we have used either a TV camera with a flash-lamp strobe or a linear-diode array. The gray-level data acquired from the image sensor have been converted into binary (black and white) to reduce the amount of image data, thereby increasing the speed and reliability of processing and lowering the cost of equipment. We have developed analytical techniques and a library of software subroutines for many vision functions, including:

(1) Finding "blobs" (connected regions) and extracting a set of distinguishing features of each blob "on the fly." These features are independent of the position and orientation of the workpieces.

(2) Training the vision subsystem to recognize an object on the basis of a subset of these features by simply showing the object to the TV camera in different orientations.

(3) Determining the identity, position, and orientation of objects that are either stationary or on a moving line.

We have designed and constructed a practical visual module for development and execution of application programs entailing vision tasks. These tasks include: training-by-showing for part recognition; determination of the identity, position, and orientation of parts, holes, and the like; visual inspection of the integrity of parts and assemblies. The Vision Module is shown in Figure 7. It consists of three major components: a GE Model TN-2200 solid-state TV camera with 128 x 128 elements; a DEC LSI-11 microcomputer with a 28K-word memory; an interface preprocessor between the TV camera and the microcomputer. The microcomputer memory stores the entire vision library previously developed at SRI, as well as certain application programs and newly developed capabilities. These capabilities include the ability to specify a "window" in the image, as well as to compute the area and first and second moments of each blob within the window more rapidly than by invoking the connectivity analysis of these blobs.

The interface preprocessor performs five functions:

(1) Providing an adjustable threshold for converting the analog intensity data to binary image data.

(2) Buffering the binary image data of one frame (128 x 128 bits) to permit the LSI-11 to read in the image data later at a rate independent of the camera speed.

(3) Converting the binary image data into run-length code, i.e., the image coordinates of the transitions from black to white and white to black (this general purpose feature may save up to 25% of the processing time in the LSI-11 microcomputer).

(4) Providing a second adjustable threshold that, together with the first one, can be used to generate intensity histograms (we intend to analyze these histograms and apply them to automatic adjustment of the first binary threshold).

(5) Triggering a flash lamp in sync with the frame scan of the TV camera.

3. Communication in Distributed Processing

We have been developing a distributed computer system consisting of a PDP-11/40 minicomputer and several LSI-11 microcomputers, as shown in Figure 8. Each microcomputer will perform an isolated modular function, such as: controlling a manipulator and its end-effector and nonvisual sensors; visual sensing and processing; controlling auxiliary equipment (x-y table, part presenter, etc.) and its nonvisual sensors; safety monitoring; voice control. Since most tasks are controlled by more than one microcomputer, we have developed a system of hardware and software for intercomputer communication. The communication hardware consists of 16-bit parallel input-output interface units (DR11-C interface in the PDP-11/40 and DRV-11 interface in the LSI-11). The communication software is organized hierarchically, so that a user can invoke any of the high-level functions provided without having to know either the details of the communication protocol or that some functions are being executed remotely. This software organization protects programs in one computer from the effects of changes made in another computer and simplifies programming for new jobs.

Our current distributed system consists of a supervisory PDP-11/40 minicomputer and two LSI-11 microcomputers, one controlling the Unimate and its end-effector and the other consistuting part of the Vision Module.

C. Examples

I now wish to illustrate the application of contact and noncontact sensors to programmable automation tasks in material handling, inspection, and assembly. Most of these examples have been demonstrated in the laboratory rather than in actual manufacturing.

1. Material-Handling Tasks

- At the University of Nottingham [Heginbotham et al., 1972], the identity, position, and orientation of flat workpieces were determined, one at a time, from a top-view image obtained by a TV camera. The camera and a manipulator were mounted on a turret in the same fashion that lens objectives are mounted on the turret of a microscope. After the identity, position, and orientation of each workpiece had been determined, the manipulator rotated into a position coaxial with the original optical axis of the camera lens and acquired the workpiece.
- At Hitachi Central Research Laboratory [Hitachi Review], prismatic blocks moving on a conveyor belt were viewed one at a time, using a vidicon TV camera. A low-resolution image (64 x 64 pixels) was processed to obtain the outline of each block. A number of radius vectors from the center of the image area to the outline were measured and processed by a minicomputer to determine the identity, position, and orientation of each block. The block was then picked up, transported, and stacked in an orderly fashion by means of a simple suction-cup hand whose motion was controlled by the minicomputer.
- The following laboratory experiments have been demonstrated at SRI in the past five years [Rosen and Nitzan et al., 1973, 1974, 1974, 1975, 1976, 1976, 1977, 1978]:
    - Part Recognition. Using distinguishing features (such as area, perimeter, and moments of part images), the vision subsystem was trained to recognize four casting types in seven different stable states, as well as eight different components of a water pump, by simply showing each part to the subsystem.
    - Acquisition of Moving Parts. Various parts (connecting rods, water pumps, etc.) were placed randomly on a moving conveyor belt. As each part passed under the TV camera, the identity, position, and orientation of that part were determined by the vision subsystem. Based on this information, a Unimate robot was commanded

to track the part, pick it up, and transport it to its destination.

- Packing Water Pumps in a Box. Using feedback from a proximity sensor and a triaxial force sensor in its end-effector, the Unimate acquired individual water pumps from approximately known positions and packed them neatly in a tote-box.
- Packing Moving Boxes. Boxes were placed randomly on a moving conveyor belt, the vision subsystem determined the position and orientation of each box, and the Unimate packed castings into the box regardless of the conveyor speed.
- Bin Picking. An end-effector with four electromagnets and a contact sensor was built and used to pick up four separate castings from the top of a jumbled pile of castings in a bin. The Unimate transported the four castings and set them apart on a backlighted table. The vision subsystem then determined the stable state, position, and orientation of each casting. With this information, the Unimate gripper acquired each casting and transported it to a destination determined by the stable state of the casting.
- Stenciling Moving Boxes. Boxes were placed randomly on a moving conveyor belt and the vision subsystem determined the position and orientation of each box. Using this information, the Unimate placed a stencil on the upper right corner of each box, sprayed the stencil with ink, and then removed the stencil.

* The above acquisition-of-moving-parts experiment was extended at the General Motors Technical Center to handle noisy cases, i.e., acquisition of parts with reflectance similar to that of the conveyor belt. The belt was illuminated by two sheets of light (forming a V-shaped side view) intersecting across the belt plane. The line of intersection with the belt was seen by a linear-diode-array camera from above. Upon arrival of a part the line of intersection was shifted horizontally and disappeared from the camera's view. The end points of the lines seen by the camera constituted the outline of the part.

2. Visual Inspection Tasks

* The following laboratory experiments have been demonstrated at SRI [Rosen and Nitzan et al., 1974]:

  - Water Pumps. Washing-machine water pumps were inspected to verify that the handle of each pump was present and to ascertain which of two possible positions it was in.
  - Lamp Bases. A group of lamp bases was inspected to verify that each base had two electrical-contact grommets and that these grommets were properly located.

* At the General Motors Technical Center [Baird, 1976] a computer vision system was developed for locating and inspecting integrated-circuit chips for automatic alignment during the various manufacturing processes. The system is now operating successfully in a factory.

D. Assembly Tasks

* At Olivetti [d'Auria and Salmon, 1975], a force-displacement sensor has been used to monitor a screwdriving operation fastening two parts together by the Olivetti SIGMA robot system. As force is applied against a spring, the displacement is sensed by an LVDT transducer. In normal operation, the arm moves to a position above the thread while the spring is relaxed, then down to the thread, compressing the spring, and finally driving the screw down, releasing the force applied to the spring. Abnormal displacements indicate no hole in the top part, no bottom part, or improper threading. This information is sent to the computer, which then aborts the operation.

* At Hitachi Central Laboratory [Goto et al., 1974], a compliant wrist equipped with a strain-gage force sensor provided feedback in controlling the insertion of a 1/2-inch-diameter polished cylinder into a hole with 7 to 20-micron clearance. The insertion time was less than 3 seconds.

* The following experiments have been demonstrated at SRI [Rosen and Nitzan et al., 1976, 1978]:

  - Pop Riveting. Force sensing was applied to detect a hole and trigger one-sided riveting in that hole. A pneumatic Chobert riveting gun was mounted on a leaf spring to which a strain gage had been cemented. A search algorithm was executed by the Unimate until the riveting gun mandrel entered a hole, causing the leaf spring to deflect. This deflection was sensed by the strain gage, the pop riveting gun was triggered, and the mandrel was pulled to expand the rivet in the hole.

- _Assembly System_. Assembly techniques, employing minimum jigging, was demonstrated by assembling a compressor cover. The assembly station, controlled by a PDP-11/40 minicomputer, is shown in Figure 9. It included: the Unimate (simulating a limited-sequence manipulator) under LSI-11 microcomputer control; a gripper mounted on a force-controlled accommodator; a programmable x-y table with passive accommodation; a 100 x 100-element solid-state TV camera above the x-y table; auxiliary equipment. After calibration the x-y table, guided by the vision subsystem, moved the compressor housing to a fixed position at which the Unimate placed the cover on the housing. The Unimate then bolted the cover as the vision subsystem positioned the x-y table for each bolt insertion and subsequent inspection. Bolting has also recently been performed by an Auto-Place Series 50 limited-sequence arm.

- _Visual Servoing_. A vertical bar of light was projected on a workpiece to visually servo a manipulator along an edge or groove of the workpiece. The image sensed by a TV camera in the manipulator end-effector is a bright streak whose shape is a function of the workpiece surfaces. Techniques and software were developed to extract the centerline points of such a streak and fit straight-line segments to them. Algorithms were developed experimentally to improve the servo response. This hardware/software system was used for visual servoing of a Unimate in two operations: simulating spot welding along a flat workpiece edge on a moving line, and tracking of a stationary "tub" groove in three dimensions.

REFERENCES

1. M. L. Baird, "An Application of Computer Vision to Automated IC Chip Manufacture," Proc. 3rd Int. Joint Conf. Pattern Recognition, 1976, pp. 3-7.

2. A. d'Auria and M. Salmon, "SIGMA - An Integrated General Purpose System for Automatic Manipulation," Proc. 5th Int. Symp. Industrial Robots, IIT Research Institute, Chicago, IL, September 1975, pp. 185-202.

3. R. L. Garrison and S. S. M. Wang, "Pneumatic Touch Sensor," IBM Technical Disclosure Bulletin, Vol. 16, No. 6, November 1973.

4. T. Goto et al., "Precise Insert Operation by Tactile Controlled Robot HI-T-HAND Expert-2," in Proc. 4th Int. Symp. Industrial Robots, Tokyo, Japan, Nov. 1974, pp. 209-218.

5. W. B. Heginbotham, P. W. Kitchin, and A. Pugh, "Visual Feedback Applied to Programmable Assembly Machines," Proc. 2nd Int. Symp. Industrial Robots, IIT Research Institute, Chicago, Illinois, May 1972, pp. 77-88.

6. J. W. Hill and A. J. Sword, "Manipulation Based on Sensor-Directed Control: An Integrated End Effector and Touch Sensing System," Proc. 17th Annual Human Factor Society Convention, Washington, D.C., October 1973.

7. "Hitachi Hand-Eye System," Hitachi Review, Vol. 22, No. 9, pp. 362-365.

8. M. E. Merchant, "The Future of CAM Systems," in AFIPS Conf. Proc., Vol. 44, pp. 793-799, 1975.

9. J. L. Nevins, N. Kamrany, and D. E. Whitney, "Productivity Technology and Product System Productivity Research," Report R-928, Vol. 2, prepared by Draper Laboratory under National Science Foundation Grant No. APR75-15334, 1976.

10. D. Nitzan and C. A. Rosen, "Programmable Industrial Automation," IEEE Trans. Computers, Vol. c-25, pp. 1259-1270, December 1976.

11. C. Rosen and D. Nitzan et al., "Exploratory Research in Advanced Automation," Reports 1 through 5, NSF Grants GI-38100X and GI-38100X1, Stanford Research Institute, Menlo Park, CA 94025 (December 1973 through January 1976).

12. C. Rosen and D. Nitzan et al., "Machine Intelligence Research Applied to Industrial Automation," Reports 6 and 7, NSF Grant APR75-13074, SRI International, Menlo Park, CA 94025 (October 1976 through September 1977).

13. C. A. Rosen and D. Nitzan, "Use of Sensors in Programmable Automation," _Computer_, pp. 12-23, December 1977.

14. S. S. M. Wang and P. M. Will, "Sensors for Computer Controlled Mechanical Assembly," _The Industrial Robot_, pp. 9-18, 1978.

15. P. C. Watson and S. H. Drake, "Pedestal and Wrist Force Sensors for Automatic Assembly," _Proc. 5th Int. Symp. on Industrial Robots_, IIT Research Institute, Chicago, Illinois, September 1975, pp. 501-511.