S35-37
188
188135

Development of a Semi-Autonomous Service Robot With Telerobotic Capabilities

J.E. Jones
American Welding Institute
Louisville, TN 37777
and Colorado School of Mines
Golden, CO 80401

D.R. White
MTS Systems Corporation
Minneapolis, MN 55424

AU 304538
CT 840271
M 8576630

1.0 Abstract

The importance to the United States of semi autonomous systems for application to a large number of manufacturing and service processes is very clear (1). Two principal reasons emerge as the primary driving force for development of such systems: enhanced national productivity and operation in environments which are hazardous to humans. Completely autonomous systems may not currently be economically feasible. However, autonomous systems that operate in a limited operation domain or that are supervised by humans are within the technology capability of this decade and will likely provide reasonable return on investment.

The two research and development efforts of autonomy and telerobotics are distinctly different, yet interconnected. The first addresses the communication of an "intelligent" electronic system with a robot while the second requires human communication and ergonomic consideration. This paper discusses the work in robotic control, human/robot team implementation, expert system robot operation, and sensor development by the American Welding Institute, MTS Systems Corporation, and the Colorado School of Mines -- Center for Welding Research.

2.0 Introduction

In almost every industry in the United States, there are opportunities for application of robotics. Robot systems which are appropriately utilized to perform work which is manually difficult, tedious, or hazardous can distinctly increase overall productivity. Such systems could replace human workers in chemically, radioactively, or biologically hazardous and uninhabitable environments such as nuclear power plan maintenance, space fabrication/maintenance and servicing, paint and coating work, chemical production, construction, and mining. In many cases, robotic systems augment rather than replace human workers. By freeing the human to do those parts of operations that require more complex intellectual activities, the robot system allows the full capabilities of both the machine and the human to be realized. Thus, the development of such robotic systems will have a significant impact on national productivity and offers the opportunity to develop new markets for U.S. goods and services.

Robot systems which can operate completely independent of human intervention and control over a wide range of processes and activities will require a substantial technological capability which is not currently available. However, systems which can operate in a semi-autonomous mode, without human interaction over a limited range of activities, are within the technology developments of this decade. Specifically, the requirements of such a system include: (1) rapid, effective, and natural communication between the system and the human supervisor; (2) efficient control in a multiple sensor environment; (3) coordinated control of multiple manipulators; and, (4) redundant and cross linked sensor systems.

2.1 Communication

Communication between the robot system and the human supervisor must be through multiple channels. The types of communication that are needed include visual, tactile, and voice. Visual communication is most important for the human supervisor and should be designed to allow the human to receive information about the status of the system without causing loss of eye-contact with the image of the operating process. Such devices as holographic display and text overlay on image displays will be needed to provide this capability. Stereo imaging to provide depth perception and three dimensional image processing capability will allow the human supervisor to "help" the robot with difficult manipulation tasks. In addition, the same sensing system can provide special information for improving the robot systems' telepresence abilities.

Tactile communication must be provided in a bi directional mode. The human supervisor needs to be provided with surface active devices that provide a "feeling" of the pressure exerted by an object on the manipulator work piece. In addition, the human control device should allow the human to communicate required movement to the robot system through tactile sensing of movements of the human fingers, hands, arms, and/or feet. Voice capability will provide a very rapid and natural communication link between the robot and the human. Connected speech recognition, which is context driven, is the most rapid type of communication used by humans. Artificial intelligence techniques can be used to interpret the signals from the robot system, and then generate natural language synthesized speech. This bi directional voice communication link will likely be the most important one in the system.

2.2 Multiple Sensor/Manipulator Environment

A semi-autonomous system will be required to operate with multiple sensors. A necessary problem with multiple sensor environments is that much of the information from the sensors will be redundent. This redundency solves some problems in that the system will be able to compensate for malfunctioning or degraded functioning of one, or perhaps more that one, of its' sensors. However, redundent information that conflicts must be rectified. Consequently, intelligent conflict managment schemes will need to be included in the control software of such systems.

The multiple manipulator environment will require that manipulators be synchronously operated when doing tasks which require the use of multiple robot arms and be operated independently otherwise. To accomplish this task, hierarchical expert system control will be required. The control system will require the coupling of high speed numeric processing, and optimized symbolic processing. Lower level processors will use the capabilities of parallel processing both for sensor data understanding as well as for manipulator control on the servo or motor function level. Supervisory control of the system will be accomplished by higher level expert systems built in a layered design. Thus the higher level control software will not be "aware" of the detailed manipulator movement, but of the overall process activity. The task of getting the manipulator to move where the higher level controllers require will be left to the lower level controllers.

The remainder of this paper will discuss two systems which have been developed that incorporate certain aspects of the needed technologies to make a semi-autonomous robot system with telerobotic capabilities operate. Voice communications and expert system control have been used to develop a voice activated robot system and a highly accurate proximity sensing and tracking system has been developed and integrated into a second robot system. These two systems represent proof of engineering and feasibility of the use of advanced computer technology to provide highly capable semi-autonomous robot systems.

3.0 Development of a Highly Accurate Robot Proximity and Tracking System

The technologies described here, although not developed specifically for space application, represent advanced sensing and closed loop control technologies of the type required if semi-autonomous robotic "assistants" are to be used for applications such as flight service units. A number of highly accurate tracking and measuring technologies have been developed for robotics and automation applications. However, highly accurate systems usually depend on the use of stiff and massive components to overcome the dynamic effects of accelerations on the structure of the manipulator. This allows the world coordinates of the end effector to be calculated strictly from the transducers attached to each machine member. Non-orthogonal axes are typically either ignored or the system is "calibrated" and transducer values are adjusted automatically by the machine's processor before they are reported to the user. This approach works reasonably well unless there is a requirement for taking data points at rates which exceed about two per second. At these higher rates, the accelerations associated with moving from one position to another quickly introduce unacceptable inaccuracies into the measurement scheme. In a space environment where small inaccuracies into the measurement and motion control can rapidly lead to unstable situations, this is unacceptable.

Using a different technology, measurements may be taken at rates of up to 100 points per second without a loss of accuracy and used for proximity sensing, path planning, and end effector control. This is accomplished through the use of a non-contacting gauge to detect the surface being sought or followed, and an interferometric Space Location System (SLS) to detect the position of the gauge (and/or end effector) in space. In this manner, an unstressed reference frame is maintained which bypasses the effects of the dynamics on the end effector and robot system.

3.1 Proximity Sensing

In the current application, the manipulator is a wrist with a laser proximity sensor based on triangulation. A laser beam is projected onto the surface being followed, and the location of the image of the resulting spot on a linear CCD array is calibrated to correspond with the distance to the surface nearby. The linear array can be read at the rate of 1000 Hz, and arrays of up to 3000 elements are available. This enhances the rate at which the array can be read over 2D arrays (typically about 60 Hz) and also provides improved resolution, since the average 2D arrays are 512 elements. The sensor is a "look to the side" gauge, which allows the sensor to track close to the work surface, while maintaining a reasonable angle for the incident laser beam. Since accuracy and repeatability are functions of the standoff and angle of incidence this is highly desirable. The standoff should be maintained at about .5 inches. The range of the gauge is 0.40 inches, with a resolution of .0005 inches. If the range of the gauge is exceeded, (due to a malfunction in the control of the robot motion, for example), the system stops operating until it is corrected by human intervention, to prevent possible collision with obstacles.

In the current measurement application this sensor is very useful, because exact knowledge of the end effector's distance from the workpiece is needed. For other applications, a less accurate sensor with a larger range or a combination of two or more proximity sensors may be more desirable. Sensors based on ultrasonic, electromagnetic transducers or other emerging technologies could be substituted. Although some change in positioning accuracy might [illegible], the manipulator tracking system would compensate for this, and enhanced path independent collision avoidance capabilities would result.

In effect, this separates the concept of "accuracy" into two areas: control accuracy and measurement accuracy. Control accuracy is related to the accuracy of the end effector. Measurement accuracy is related to the sensors accuracy. In-flight service application, needs for control accuracy will almost always be very high. In some conditions the need for measurement accuracy may be considerably less. Multiple or interchangeable sensors of the types described earlier would be most suitable for these changing conditions. A tracking and control system of the type used here for highly accurate end effector control ("control accuracy") enables the robot to make optimal of use of any type of proximity sensor data.

3.2 Tracking and Space Location

The proximity sensor gives an accurate picture of the distance of the manipulator from the surface it is following. In order to know the location of the robot, the space location system (SLS) is used. The SLS uses multiple tracking interferometers to determine the exact location of the manipulator. The SLS also provides the necessary feedback for the system's endpoint control scheme.

The fundamental technology of the SLS is a set of tracking interferometers. The basic principle underlying these is the interference of light that has traveled over two different paths. In non-turbulent atmospheric conditions, accuracies on the order of a fraction of a wave length of light (100 nm) can be measured between the laser source and a retro-reflector (referred to as a "retro"). These retros have the property of reflecting incident light back into its incoming path, thus providing a return signal for the interferometer. As the retro is moved away from the source, a detector counts the number of times that the optical interference between the outgoing and return beams changes from maximum to minimum. This count is then multiplied by a conversion factor which is based on the frequency of the laser being used in the interferometer. Note that the interferometer only measures the change in the position of the retro along the axis of the beam. To get absolute measurements, the retro must be moved to a reference point and the interferometer zeroed. The unit will then read the distance travelled from the reference point.

The six lasers used by the SLS interferometers track retros mounted on a target which is a fixed to the manipulator. Tracking is achieved by using the return beam not only for the measurement of the distance from the tractor to a retro, but also as an indication of whether or not the retro is moving. As the return beam from the retro enters the tracker, a portion of it is re-directed to impinge on a quad-cell sensor. This sensor detects minute deviations in the location at which the light is impinging, which is directly proportional to a change in the position of the retro. This information is then used as an error-signal in a servo loop to redirect the outgoing beam in the direction the retro has moved. In this manner, each of the six beams of the SLS are made to track a retro as the target moves through space.

The SLS is responsible for providing the interferometer-measured distance between each tracker and the retro which it is tracking. To calculate all six degrees of freedom of the target, we need three distances to one point on the target, two distances to another point and one to a third. The three distances to the first point will "fix" the targe in XYZ space, but it will still be free to rotate about that point. The two distances to the second point will remove two of the remaining angles of rotation of the target, but it will still be free to rotate about the line between the first and second points. Finally, the single distance to the third point on the target confines the target to be in a single position and orientation.

Clearly, the geometries of the tracker assembly and target assembly must be known to a very high degree of accuracy if the overall SLS accuracy is to be maintained through subsequent calculations. In this system a complete self-calibration procedure was developed to ensure that the necessary accuracy regarding information such as assembly geometries and to ensure that the initial interferometer distances could be obtained automatically using the SLS itself.

This self-calibration procedure consists of moving the retro-reflector to several locations in the work space, none of which are precisely known, changes in interferometer readings are measured, and interactive calculations to locate the center of each retro are made.

Other approaches to the problem of dynamic laser-based measurement of manipulators have been explored. In one effort (2), the system is limited to three degrees of freedom (DOF). In another technique requiring use of angular encoders the accuracy of the encoders may prove a limiting factor, but it is proposed as an alternative approach allowing six DOF measurement. Costs may be reduced using the method of Gilbert (3), however, diminished accuracy will also result.

3.3 Path Planning and Control

In the current system, path planning algorithms were designed for tracking complex surfaces in a small work space and may have little relevance to many of the problems likely to be encountered by semi-autonomous flight service units. However, the concept behind the system, its interaction with the interferometric tracking system and the requirements for end effector control are useful to consider. A planning utility converts locations in space or on a surface to the coordinate system of the robot. Segments formed in this way are linked together to form a complete path. In a system with highly developed vision capabilities this would be compared with sensor input to ensure that the path will be free of collisions. In the current system this is compared with an existing representation of the workcell and workpiece to ensure that the path is collision free. A path for the end effector from one point to another in robot coordinate space is then calculated by the planner utility through the use of a heuristic search technique.

Regardless of the complexity and rigor of the path planning methods used, a typical control scheme which does not take into account the dynamics of the machine being controlled will incur significant dynamic inaccuracies. In the current application these problems are resolved by ensuring that the control scheme is concerned with "end point control". This entails using information from the SLS not only to obtain measurement accuracy, but also to provide feedback on the true position of the robot manipulator for motion control purposes. The advantage of end point control lies in its ability to use both types of position feedback, conventional and end point, since each has inherent advantages at different frequencies.

Three position feedbacks are used. The first is from velocity transducers connected directly to the motors. The second is from incremental encoders along each of the system members. The third and final position feedback is from the SLS. The SLS and encoder feedback are summed to form a composite feedback with great precision from the SLS at low control frequencies, and stabilizing effects from the encoders at higher frequencies. It is important to note that without this encoder feedback at higher frequencies, the control loop would have severely limited frequency response. It is this combination of the feedback signals which serves to extend the useable band width of the position control loop and improve the measurement speed capability of the system. Given the appropriate machine resolution, this scheme allows the proximity sensor and manipulator to be positioned dynamically to within .025mm, at high rates of speed.

The system described here integrates intelligent path planning methods, adaptive control techniques, and real time tracking and proximity sensing to perform its tasks at a high speed, with great positional accuracy and with a minimal risk of collisions. In addition, the sensors used are independent of the mechanical design of the system and can be scaled up for use on very large parts, and across long distances.

4.0 Development of a Voice Activated Robotic Assembly System

Throughout the history of computers, the keyboard has been one of the most common problems with the human/computer interface. The human-to-computer link has always required a set of switches (keys) to be manipulated by the human. Such a link depends on the speed with which a human can manipulate those keys, and error check such manipulation, while continuing to think about the information to be communicated. The result is a generally slow, cumbersome, and inefficient communication link.

4.1 Voice Recognition

Humans do not communicate with each other in a natural environment with a keyboard. Instead, a variety of communication techniques are employed -- one of the most important being verbal. Thus, a most efficient technique for human communication is the one most practiced: speaking and listening. Voice recognition by computers has been studied and developed for many years. However, it has been clear throughout the history of voice recognition that such activity was quite compute intensive and required fairly large quantities of high speed memory capacity. It is only recently that computer speech recognition can be accomplished on computers other than large capacity mainframe systems. Recent advances in digital signal processing integrated circuit development coupled with the increased memory and speed capacity of small computers has made possible the use of speech recognition by microcomputers.

Two general categories of computer speech recognition are currently used on microcomputer systems. For purposes of their discussion, these two types will be referred to as: Voice Print Dependent (VPD) and Voice Print Independent (VPI). In each case, the basic computer hardware configuration is identical. Figure 1 is a schematic representation of the hardware associated with a speech recognition system.

The system operates when the input transducer (microphone, headset, telephone, etc.) receives a voice signal and transmits the resulting electronic analog data to the speech processing device. Appropriate filtering and conditioning of the signal occurs in a preprocessor. Then the signal is passed to a digital signal processing circuit. The analog signal is converted to a digital data array and passed to the template matching hardware. By the use of optimized template matching techniques, a specific template, from a set of presorted templates, is identified as having the closest match to the input signal array. The index number, or other identification of the matched template, is then passed to the host computer through a direct bus link (5).

In the template matching stage, the system can match to a set of templates which are presorted digitized signals of an actual voice. Those voice signal templates are distinctly characteristic of the speech patterns of the individual who input the voice signals. In this case, the system is a VPD type recognizer, and the template matching is done very rapidly. The system will search for a nearly identical digital template among the prestored set available. In the case of a VPI type recognizer, the template matching algorithm searches instead for a generalized template pattern which most closely matches the input signal digital array. This generalized template is related to specific pronunciation of a word, rather than a digitized signal from a certain human voice. Thus the system can recognize the word when spoken by anyone, but requires extensive processing to match a generalized template in comparison to matching an identical voice signal template. Such matching is much slower, requiring a greater complexity of processing and is difficult to optimize. Consequently, a VPI type recognizer was not employed in the current prototype system.

4.2 Speech Generation

In order for a human to have a complete voice communication link to a robot system, the computer must also be capable of generating responses as spoken words or phrases. The generation of speech by a non-human is a problem of greater complexity than that of speech recognition. Three separate processes must be accomplished to allow the computer system to generate intellectually useful speech. First, a set of signals within the computer is processed into a series of ideas or concepts. Second, that series of ideas is translated into an ordered set of words in the form of text. Third, a Text-to-Speech-Synthesis is accomplished and the speech is sent to an output transducer (speaker, headset, telephone, etc.). Figure 2 is a schematic representation of the computer speech generation system showing the output speech processing device.

While a human is capable of understanding and processing a relatively wide range of ideas or concepts, which are stimulated by external signals, computer systems can only operate in a limited knowledge domain. Expert systems developed from artificial intelligence based technology can operate well with such symbol manipulation. These systems also have a limited knowledge domain, but can be easily expanded into additional areas of the domain.

4.3 Signal Processing

The interpretation, (by the Signal Processing System), of signals received from the Robot Control Module (RCM) into ideas or concepts that will eventually be communicated to the human operator is an excellent application for expert system technology. The expert system is only required to operate within the limited knowledge domain reserved for operating the RCM. Thus the signals, as well as error and warning messages from the RCM, are the only external signals that the speech generation system currently operates on.

4.4 Text Generation

The translation, (by the Text Generation System), of conceptual information generated by the input of external signals is accomplished through the use of natural language linkage. Thus, a signal which indicates that an error has occurred with the robot, when combined with knowledge about the last command or sequence of commands sent to the robot, generates a concept of the specific situation. The understanding of that situation which can be communicated to the user. The output of this translation is a text string in the operators own language, which is passed to the Text-to-Speech-Synthesizer (TTSS).

4.5 Text-To-Speech Synthesis

For several years, computer speech generation has been the subject of research; much of which has been in the artificial intelligence area. Beginning with simple systems that played back prerecorded speech, computer speech generation progressed to synthesized words linked to forms phrases. The prerecorded speech systems lacked the flexibility to allow fast and complete responsiveness by the computer which was sufficient to allow the system to inform the user of all needed responses.

In order to synthesize speech from text, the TTSS must combine several functions. First, the phonetic pronunciation rules must be available to determine the correct pronunciation of each individually pronounced bit of speech called a phonem. The TTSS must establish the phonems in each word, and then determine the most optimum way to pronounce them.

Secondly, the TTSS must determine the use of each word by phrase parsing techniques. Then the pitch, accent, and speed of the speech for each word in the phrase is determined. The phrase parsing is accomplished using artificial intelligence techniques, which utilize a series of rules to transform natural language text into a representational structure thereby allowing each phrase to be individually analyzed for appropriate pronunciation of each word (6). The result is a verbal response, which includes words properly connected, with pitch and accent adjusted to sound like ordinary human speech. Thus the robot operator can use natural listening techniques rather than needing to "learn to listen" to the robot differently than to other people.

4.6 Voice Activated Robotic Joining

The voice activated robotic (VAR) joining system consists of two computer systems and an interfacing network and software which operates the system, as shown in Figure 3. The first computer system, called the voice control module (VCM), is equipped with a template matching type speech recognition system based upon the Texas Instruments TMS32020 integrated circuit. The VCM is of the voice print dependent (VPD) type and requires that the operator pre enter a series of voice prints which are stored as digital templates. Those digital templates are subsequently matched against a digitized voice signal to determine the specific word or phrase which is spoken by the operator. Two different models of VCM were used in this project; the first was a Texas Instruments Professional PC (TIPC); the second was a Texas Instruments Business Professional (TIBP) computer.

The second computer system used in the VAR is the robot control module (RCM). The robot system selected was a Unimation PUMA 560 set up for joining metals with a Linde Digi Mig welding wire feeder and Linde 450 power supply. A modified Digital Equipment Corporation LSI 11/02 computer was used as the RCM. The configuration of the RCM is shown in Figure 4 (7,8,9).

The voice control software was written in two high level computer languages. compiled BASIC and IQLISP, a dialect of the LISP language produced by Integral Quality, Incorporated. Certain specialized algorithms for voice recognition and text to Speech synthesis were produced by Texas Instruments Corporation and were written in the PASCAL and in the FORTH language. The software includes algorithms for the following activities: (1) voice recognition; (2) command interpretation; (3) command and response level interfacing of the VCM and RCM; (4) response interpretation; (5) textual response generation; (6) text to speech synthesis; (7) error trapping; and (8) command/response analysis.

The VAR system prototype, which was developed for this investigation, is designed to operate under normal fabrication environmental conditions entirely under voice control. Software and software/hardware interfaces were developed to function in each of the eight activity areas described above. Each of these activity areas required a specific set of algorithms and each was approached by algorithm development and program coding based upon the performance requirements of that area.

4.7 Voice Recognition and Interpretation Activities of VAR

The voice command entered by an operator is digitized and then matched to a digital template. The template indexing allows the command interpretation software to issue the command to the RCM in the appropriate format for the RCM. This indexing system then allows the command interpreter to accomplish natural language linking of the human operator and the RCM command structure (10,11,12). Natural language linking is an important aspect of improving the robot/human interface in moving toward the robot/human team concept. Thus, the operator can talk to the robot system using his/her own words, and the command interpreter issues the appropriate format command to the RCM. For example, the operator could say any of the following words: "Wait", "Stop", "Halt", "Quit", or others. The command interpreter would issue the command "ABORT" to the RCM which causes the current operation in which the robot was engaged to stop.

The command interpreter also performs the function of identifying faulty commands and taking appropriate action. If the operator issues a faulty command, the fault may be of several forms. A command fault exists if (1) the operator speaks a word or phrase for which there is no digital template match; (2) the operator issues an incomplete command, for example, the command to change the robot arm speed with no new speed indicated; and (3) the operator speaks a command to which the robot cannot respond, for example, a computer control movement command might be issued that cannot be executed because the system was in the manual mode.

If, for instance, the operator issues a spoken word or phrase for which there is no digital template match, then either the operator wants the robot to do something which it cannot do or has used an incorrect word or phrase for a specific command. In either case, the command interpreter will recognize the non-match condition, then send information to the voice response system which results in a spoken, natural language description being issued to the operator. In the case of an incomplete command, the command interpreter will note the missing information and generate a textual response. That is then sent to the TTSS which generates a spoken description of the problem for the operator.

Following recognition and interpretation of a command, the Command and Response Level Interface System formulates and transmits the command from the VCM to the RCM. The interface system then obtains a signal response from the RCM which may include a variety of information including: (1) normal response; (2) warnings to the operator; and (3) error messages. The interface accepts the response from the RCM and then processes that response and passes it to the response interpreter.

Since the response from the RCM includes data that is important to the operator, the response interpreter must analyze and interpret responses from the RCM and prepare the speech synthesized response to the operator. The interpreter examines the response and determines its form for further processing. If the response is normal for that command, then the interpreter prepares a text string which tells the operator that the system has heard the command and asks for verification. If the response indicates an abnormal RCM reaction, such as an operator warning or error message, then the response interpreter examines the response and compares it to known RCM responses. It then compares the response to the command that was issued. If the command has generated a response which is due to a standard entry error, then a text string is generated by the Textual Response Generator which explains the abnormal response. If the response is a non-standard one, then the textual response generator prepares a text string which duplicates the abnormal response to the operator.

The text generated by the response interpreter textual response generator is then passed to the text-to-speech synthesis (TTSS) system. Using several parameters that can be adjusted in the software, the TTSS system produces a spoken response to the operator. The speed, pitch, phonetic pronunciation, and other parameters can be adjusted in the system to produce speech which is most useful in the specific application of the robot system. The TTSS system makes use of a text-to-speech software/hardware system produced by Texas Instruments Corporation.

4.8 Vocabulary Development and Use

The digital template matching software is capable of storing up to 50 templates in a single vocabulary file. A main vocabulary was developed which has all of the commands which can be used to control the RCM. The individual operator can enter one or more voice prints for each of the robot commands. An auxiliary vocabulary of alpha and numeric characteristic is also available to the operator for entering special character strings to the RCM.

4.9 Future Considerations

The current prototype version of the VAR is designed to allow an operator to observe the assembly process and, by utilizing the main and auxiliary vocabularies, to alter the joining parameters and robot path. This system allows for hands-free operation of a robot and provides a capability for real-time adjustment of the assembly process. The long-range goals for speech control of assembly include future capability which will allow the operator to describe the current situation and allow the VAR to determine the appropriate actions to optimize the assembly process.

5.0 Conclusion

The technological requirements for a semi-autonomous robot with telrobotic capabilities are being developed and proven. Some of the technological capabilities have been demonstrated by the two systems described in this paper. Robot systems can perform complex tasks requiring some human intervention quickly and efficiently as a human/robot team. The degree of autonomy that such systems can achieve is dependent on the sensing and control, as well as on simulated deductive reasoning accomplished by the computer systems. Rapid and efficient voice communications in a natural language format can be applied to robot system control. In addition, highly accurate proximity sensing and tracking of robot activities can be achieved with current technology.

6.0 References

1. "A Competitive Assessment of the U. S. Flexible Manufacturing Systems Industry", Technical Report, International Trade Administration, U. S. Department of Commerce, July 1985.

2. Greenleaf, A., "Self Calibrating Surface Measuring Machine", SPIE Proc., Vol. 228, April 1980.

3. Gilbert, J.H. and Parker, G.A., "Laser Tracking System to Measure Robot Arm Performance", Dept. of Mech. Engr., Univ. of Surrey, Gilford, UK, Sensor Review, Oct. 1982.

4. "The Robotics Industry", Technical Report, U. S. Department of Commerce, April 1983.

5. Nof, S. Y. (ed.), "Handbook of Industrial Robotics", John Wiley and Sons, New York, 1985, 1358p.

6. Jones, K. S. and Wilks, Y. (ed.), "Automatic Natural Language Parsing", John Wiley and Sons, New York, 1985.

7. "PUMA Robot Technical Manual", Unimation Incorporated, Danbury, Conn. 1980.

8. Lechtman, H., Nagel, R. N., and Ploman, L. E., "Correcting the PUMA Robot with the MIC Vision System and Other Sensors", Robots VI Conference Proceedings, Detroit, 1982, pp. 447-466.

9. Melidy, A. and Goldenberg, A. A., "Operation of PUMA 560 Without VAL", Robots IX Conference Proceedings", Volume 2, Dearborn, Mich., 1985, pp. 18-79.

10. "Users Guide to VAL: A Robot Programming and Control System", Unimation, Incorporated, Danbury, Conn., 1980.

11. "Breaking Away From VAL", Unimation, Incorporated, Danbury, Conn. 1982.

12. "Hierarchical Control Structure Using Special Purpose Processors for the Control of Robot Arms", Proceedings of the PPIP 82, IEEE Computer Society Conference on Pattern Recognition and Image Processing, Las Vegas, June 1982.

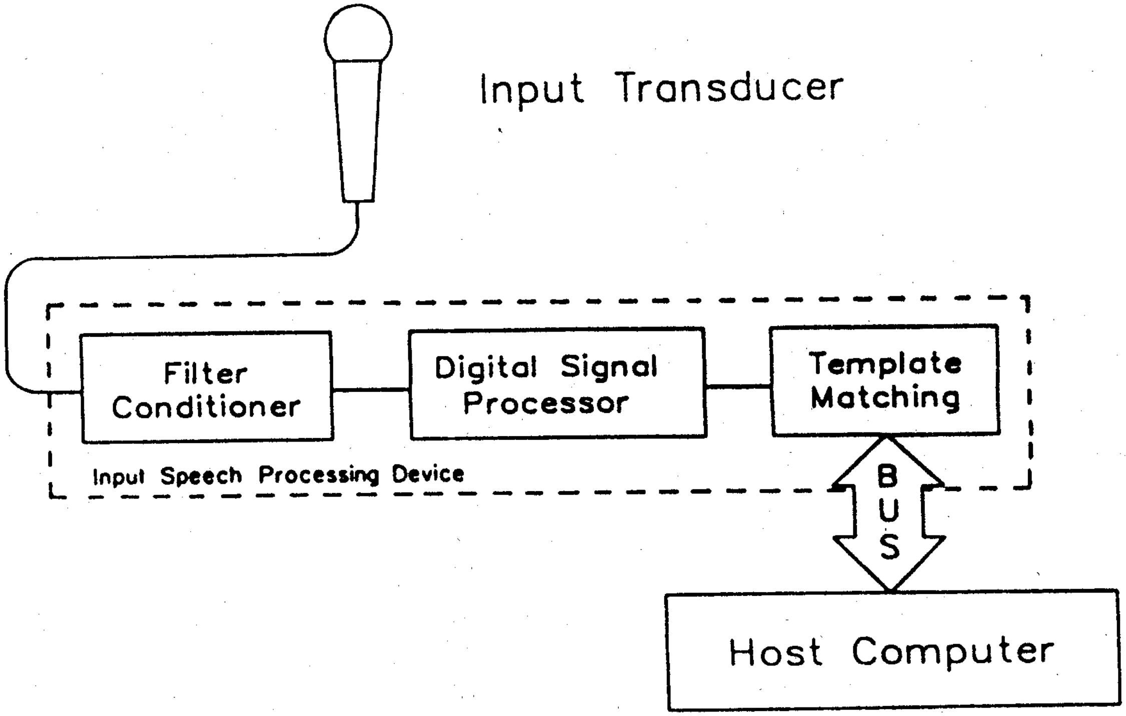

Figure 1. Schematic representation of the hardware system used for voice recognition on the VAR. Depicted are the three major processing components: A filter/Conditioner for the input voice signal; a high speed integrated circuit Digital Signal Processor; Template Matching hardware.

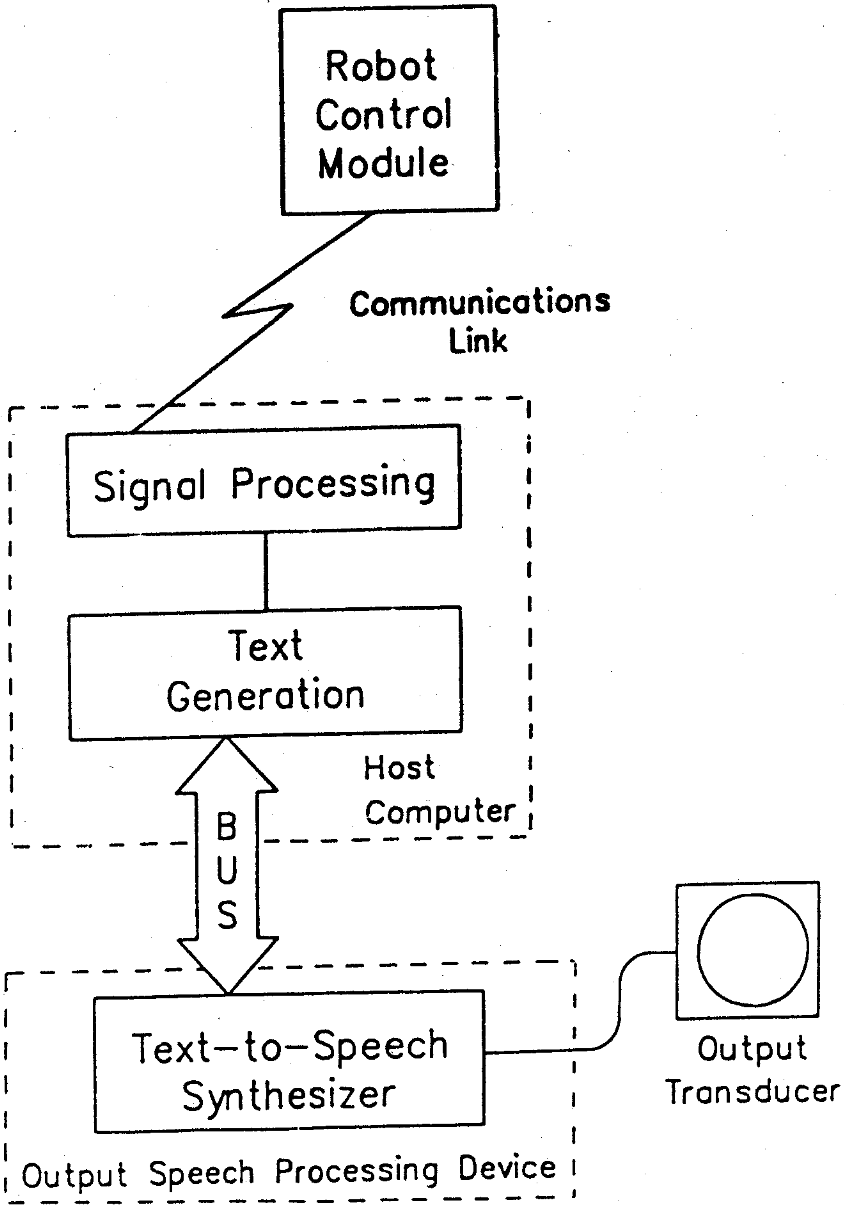

Figure 2. Schematic representation of the computer speech generation system. Depicted are the major component subsystems: Signal Processing from the RCM; Text Generation from signal processing generated concepts; Text-to-Speech Synthesis.

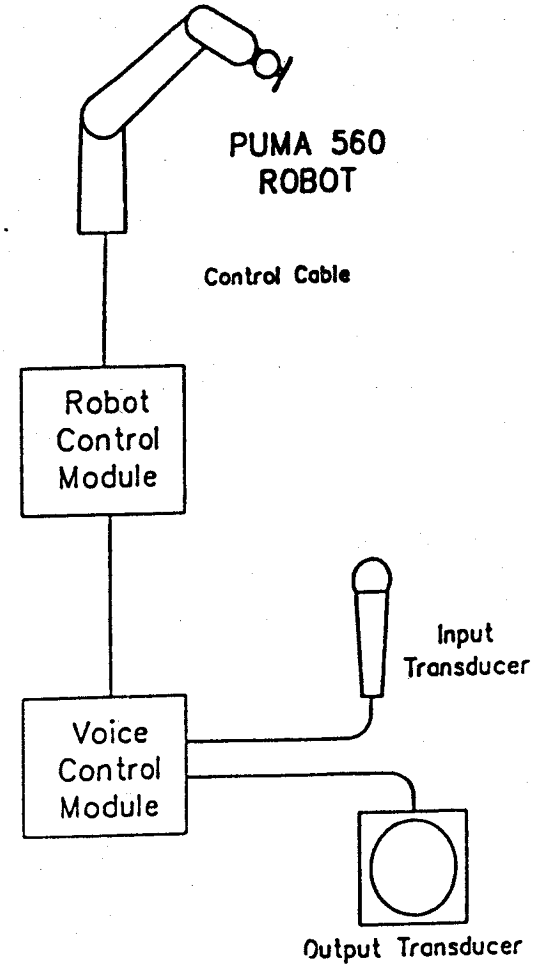

Figure 3. Schematic representation of the voice controlled robot system indicating the major components of the system.

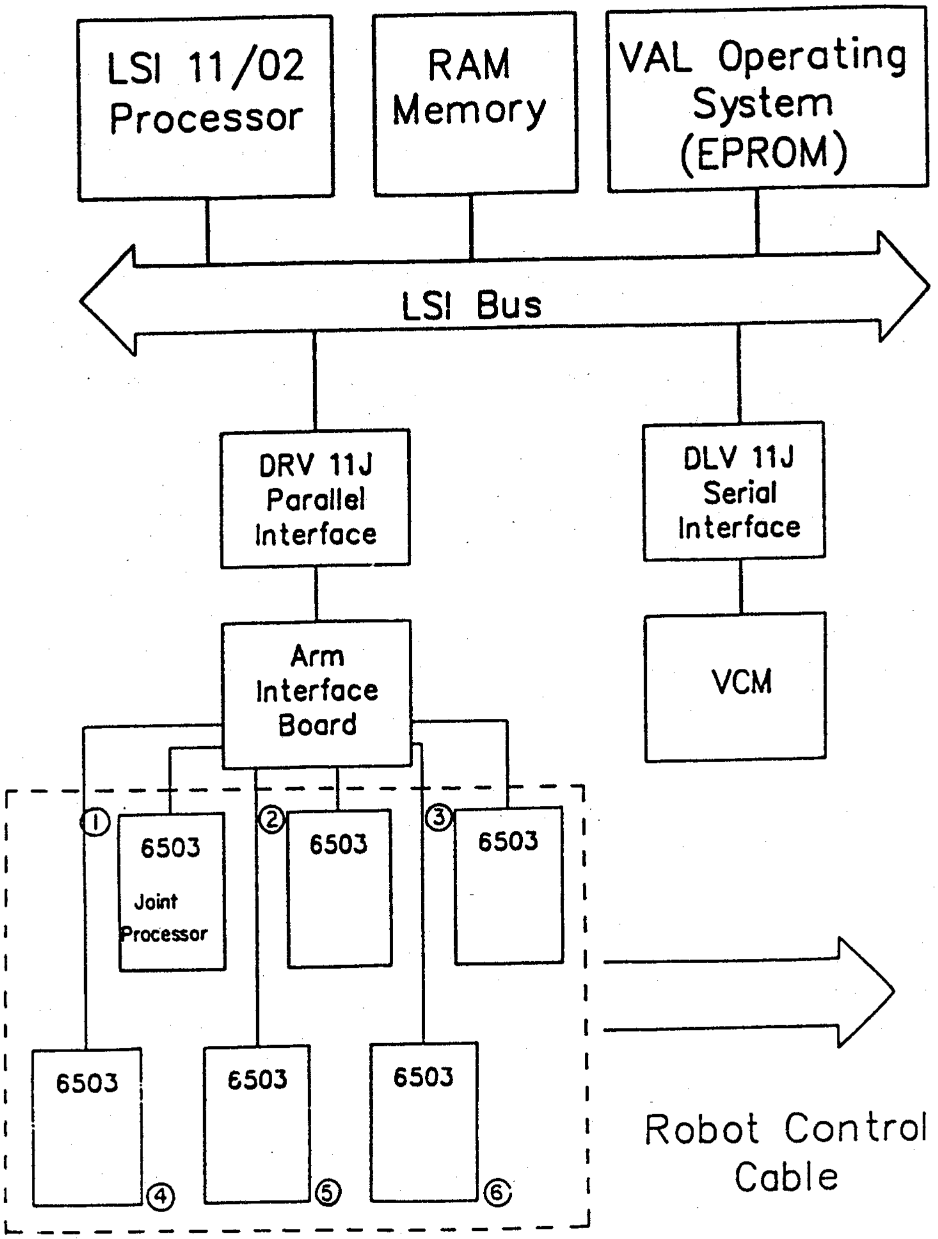

Figure 4. Hardware configuration of the Robot Control Module (RCM) indicating interfaces with the robot and the Voice Control Module (VCM).