

BYTE

the small systems journal

**A Mobile,
Cognitive Robot**



Newt:

A Mobile, Cognitive Robot




Figure 1: The robot Newt as it will appear when completed. There are three main subsystems shown in this diagram which was supplied by the author: motive, manipulator and sensory turret. The entire machine is controlled by an on board microcomputer. The motive power subsystem is built around two precision drive wheels actuated by stepper motors. The manipulator, a simple hand capable of grasping, lifting and rotating, is also actuated by stepper motors and includes various sensors. The sensory turret at the top is a platform which includes a main image sensor, which can be tilted to look straight down or up at a slight angle using a stepping motor; other sensors are shown in outline form, and are a subject for future experimentation.

Ralph Hollis
Dept of Physics and Astrophysics
University of Colorado
Boulder CO 80309


In the late 1930s, a young man named Rossum began manufacturing industrial robots in a small factory on the outskirts of Prague. This venture was immediately successful and would have virtually guaranteed a second industrial revolution had it not been for a singular tragic circumstance: The robot workers became irrational and revolted. They turned on their masters and burned the factory to the ground.

Fortunately, the preceding scenario is only a work of fiction by the Czechoslovakian writer K Čapek. Since Čapek's coinage of the word "robot" in his 1923 play *R U R* it has been the subject of a great many works of science fiction, including Isaac Asimov's *I, Robot* and the movies

"Forbidden Planet," "Gog," "Silent Running," "Westworld" and others.

Alas, the imaginative stories of the science fiction writers have far outstripped the efforts of the robot engineers. Progress in building real robot devices has been painfully slow over the past several decades, although a few individuals and small groups have produced some very interesting results.

It is not my purpose here to review these works, but rather to describe a project with which I have been personally involved for some time: the construction of a freely moving robot vehicle which will be capable of interacting with its environment in a rational way, and managing its own survival. The work is being done by a small group of people in the Duane Physical Laboratory of the University of Colorado. The project is being financed through personal funds, and in this sense qualifies as an "amateur" undertaking.



What are the requirements for such a robot? It must be able to explore the environment in some orderly manner, measure the attributes of objects and obstacles encountered, classify them according to some scheme, and incorporate them in its evolving internal world model. The world model must have a logical structure which allows modifications to be easily made; it must be compact with respect to memory space, and it must use a design which can be consulted in some reasonable way. The robot must be able to manipulate the world model (cognition), derive informed decisions from it, and carry out these decisions in physical action to achieve broadly defined goals.

Figure 1 is a general view of the robot as it will appear when completed. Practically speaking, it resembles less a mechanical man than, say, a shop type vacuum cleaner. The machine is cylindrical, about 36 cm (14 inches) in diameter and 76 cm (30 inches) in height, weighing approximately 27 kg (60 pounds). All rigid mechanisms lie within the cylindrical boundary or can be retracted within the boundary if necessary. This design greatly reduces the number and diversity of senses required. The robot will be fitted with a smooth cylindrical skin (not shown), removable in sections for easy inspection and maintenance. Modular construction is used in the robot wherever possible. There are three main subsystems: motive, manipulator and sensory turret. These are presided over by an 8080 based microcomputer with 8 K bytes of EROM and 24 K bytes of programmable memory. The entire system is powered by a 6 V, 84 amp-hour storage battery.

How the Robot Gets Around

The robot is given locomotion by two main diametrically opposed drive wheels. A third (unpowered) castoring wheel is located at the rear. This motive geometry has been successfully employed in several other robot vehicles constructed by other groups, and has been in use in the author's robot research since 1957. The vehicle's center of gravity is located well aft due to the placement of the battery and other heavy components, obviating the need for a front wheel. The wheels are constructed of aluminum with neoprene O-ring tires. The main wheels are driven by stepping motors through precision 3:1 gearing. A single command step rotates a wheel 0.6 degrees, so there is a total of 600 steps per revolution. If both wheels are moved in the same direction, the vehicle travels forwards or back-

wards in a straight line. If the wheels are moved in opposite directions, the vehicle executes a perfect rotation about its vertical axis. By stepping the wheels at different rates, a circular trajectory is approximated having a radius which can range from zero to infinity. The hardware and software necessary to drive the stepping motors are discussed in some detail later in this article, since they have very general application.

The robot navigates principally by open loop dead reckoning. That is, it depends largely on the precise control of the wheels during acceleration, deceleration, and constant speed motion to achieve accurate positioning of the vehicle. The robot's position and azimuth relative to a fixed origin are computed at the end of each motion segment by counting stepping motor increments and using trigonometry. The precision attainable is limited principally by wheel slippage, unequal wheel diameters, nonplanar floors, round-off errors, and step quantization. In practice, all these errors are small for short distance movements.

The robot's excellent open loop performance makes it unnecessary to have continuous closed loop servo systems, such as have been extensively employed in other robots. For example, consider the problem of having a robot view a small object on the floor a short distance away, and then go pick up the object. One approach would be to have the robot continuously view the object as it moves towards it, adjusting its motion in a continuous way to converge on the object and pick it up. In this approach, a rather crude motive system would suffice, since errors are always nulled. A heavy load, however, is placed on the sensor-computer system. An alternative approach, the one followed here, is to have the robot view the object once, and then compute exactly how to move to the object and pick it up. When the computation is completed, the machine simply carries out the proper motions in a blind fashion. There is a greatly reduced strain on the sensor-computer system, at the expense of having to build rather precise motive machinery. Of course, this approach assumes the relevant environment will remain fixed for the duration of the task.

In this general spirit, by using high precision open loop movements throughout the robot, only intermittent feedback through the senses is required to close the control loop. In this way, the overall complexity of the robot can be kept at a reasonable level, allowing the required computations to fall within the abilities of an on board microcomputer.

Epistemological Engineering

With artificial intelligence, robotics and applied cybernetics coming of age through the recent progress in the fabrication of computer systems, it is quite likely that the branch of philosophy called epistemology will have a much more explicit role in technology during the coming years. Epistemology is the study of nature and grounds of knowledge; understanding of epistemology is crucial to any attempt to realistically implement artificial intelligence systems. So in a future world of cognitive automata and advanced information systems, we may indeed find the new specialist who is the "epistemological engineer."

Scanning the Environment with Electronic Senses

The sensory turret (see figure 1) provides a general platform on which to mount various senses, some of which might be quite experimental and temporary in nature. It can pan a full 360 degrees, and a small section containing the main image sensor can tilt from approximately 30 degrees above horizontal to 90 degrees below horizontal. Both of these motions are controlled by stepping motors. The main image sensor has a motorized focus control. The geometry of the "hand-eye" system allows orthogonal views of objects held in the manipulator jaws by using both image sensors under conditions of controlled focus and lighting.

Each image sensor is a 32 by 32 element integrated array of photodiodes on a silicon chip measuring approximately 4 mm (0.1 inch) square. The array acts like a 1024 bit memory. Each element is precharged to a fixed voltage; then at some later time the voltage of each element is read out, decreased in proportion to the amount of light which has fallen on it. Several milliseconds are required to digitize the image. Up to 16 levels of gray can be discerned, which means that 512 bytes of computer memory are sufficient for a single image. Once the image is obtained, it is analyzed by appropriate software. The extremely small size of these solid state image sensors, and the simplicity of their associated electronics, make them ideal for robot use. Higher resolution devices such as charge coupled sensors and miniature television cameras are available, but their cost and complexity make them unattractive for such applications. Besides, the amount of data generated would be too unwieldy to handle with a microcomputer.

In addition to the main image sensor, several simple phototransistor light sensors are mounted on the tilting platform of the sensory turret. These enable the robot to locate and track point sources of light.

Several other proposed senses are to be mounted on the turret. These are shown in schematic form in figure 1. To the left of the image sensor is shown an ultrasonic ranging transmitter and receiver which should be extremely useful for finding the range to walls at distances as great as 10 meters (33 feet). The idea is to transmit short bursts of 40 kHz sound and measure the time required for an echo to be received in order to compute the distance. Much of the necessary circuitry is available in compact integrated form.

Just to the right of the image sensor is shown a long wavelength infrared radiation detector with which it might be possible to

locate sources of heat such as dogs, cats and humans.

Located on an axis perpendicular to the main image sensor are shown two microphones which form part of a sound location system. The intent is for each sensor to acquire a short sample of sound. The computer then attempts to find a phase relationship between them, thereby locating the direction of the source.

Many other possible senses can be imagined; the sensory turret is intended to provide a versatile platform and interface for experimenting with them. For example, using a helium neon laser, or perhaps a compact metal vapor laser, it may be possible to provide an optical range finding system with high resolution if a suitable detector can be found. Eventually it is hoped to provide limited forms of speech synthesis and phrase recognition (using off line electronics with an analog radio link). This would make it possible to give general spoken commands to the robot in contextual surroundings, and have verbal feedback to insure that the commands were being properly interpreted.

Sharing some space with the sensory turret system, and directly below the platform in figure 1, is the necessary power conversion electronics. These modules convert the battery voltage to ± 12 V and +5 V (regulated) for use by the electronics, and also house the battery recharging circuitry.

Manipulating Objects in the Environment

The number of tasks which a simple robot can do is greatly increased if it has some sort of hand with which to grasp and manipulate objects in the environment. The human arm and hand system, with its 27 degrees of freedom, is a marvelously versatile mechanism. Large industrial manipulators are fairly complicated and have six or seven degrees of freedom. For practical reasons, the present system is limited to a mere three degrees of freedom which, when combined with the two degrees of freedom of the motive system, should be sufficient for carrying out simple tasks.

The manipulator is able to grasp objects, move up and down along the front of the robot by means of a rack and pinion drive, and to rotate about a horizontal axis. All three motions are controlled by stepping motors. The manipulator has a parallelogram geometry which permits the jaw faces to remain parallel regardless of their separation. This feature simplifies the problem of picking up objects of varying sizes. When not in use, the jaws open wide enough to bring all

parts of the manipulator within the cylindrical boundary of the robot, where it is out of the way and does not cause problems when the robot turns and maneuvers in tight places.

An auxiliary 32 by 32 element image sensor with fixed focus lens is mounted directly in the manipulator assembly, providing a view of whatever object is between the jaws.

In addition to mechanical force sensors in the jaw faces, there are several infrared LEDs in one face with opposing phototransistors in the other face. These LED phototransistor pairs define beams of light between the two jaw faces. By scanning the manipulator up and down, and moving the entire vehicle forwards and backwards, these beams enable the robot to measure the height and depth of simple objects resting on the floor between the jaw faces, as well as to determine when objects are correctly positioned for picking up. The widths of objects are measured directly by counting the number of stepping motor pulses required to close the jaws on the object.

A tactile sensor is mounted on the front of each jaw to enable the robot to sense and locate large obstacles such as walls. The robot can align itself accurately either parallel or perpendicular to the wall, after performing a simple series of maneuvers. This action is convenient for many of the tasks the robot might have to carry out. For example, a simple strategy of maneuvers allows the robot to "square itself off" from a corner in a room, thereby defining a precise origin for its navigational coordinates.

The tactile sensors, double as electrical contacts used for plugging into ordinary wall outlets for the purpose of recharging the robot's battery (more about this later).

All electrical power and signals to and from the manipulator assembly are sent through folding ribbon cables which are not shown in figure 1.

Auxiliary Senses

In addition to the senses mounted on the manipulator and sensory turret assemblies, there are several auxiliary senses illustrated in figure 1. A pair of retractable feelers are provided in the rear of the vehicle and serve much the same purpose as the tactile sensors on the manipulator jaws. These rear feelers enable the robot to detect and square up from walls and corners if the manipulator is occupied with holding some object.

In addition to the rear feelers, a pair of retractable wall-follower feelers located on the sides of the vehicle enable the robot to

travel parallel to a wall by measuring the distance from itself to the wall. The rear feelers and the wall-follower feelers are both activated by double coil latching solenoids and contain "microswitches" consisting of movable masks and infrared LED-photo-transistor pairs.

For all of the robot's varied and specialized senses we have discussed so far, there is nothing to prevent the robot from accidentally running at high speed into a wall. To prevent just such a mishap, and to provide a "soft" broad area "sense of touch," the robot is equipped with a number of proximity detectors, several of which are shown in schematic form in figure 1. These detectors work on a simple, but elegant, principle. Infrared LEDs, amplitude modulated by a 20 kHz signal, send broad beams of light out from the robot. If a wall is nearby, some of this light is reflected back into phototransistor sensors located in the same module. The received signal is amplified and sent to a phase sensitive detector locked to the outgoing signal. The computer is notified if the signal exceeds some programmable threshold. These proximity detectors are quite insensitive to ambient light conditions, not very sensitive to the color or texture of the reflecting surface, and have useful ranges extending from several centimeters to perhaps one meter. With these detectors, the robot is free to travel at fairly high speed until a wall or other obstacle is encountered, and can then slow down and investigate it with caution.

It should be mentioned that all of the senses interact with the processor through a vectored priority interrupt system.

Robot Psychology

Up to the present time, the environment of the robot has been intentionally restricted. The robot is permitted to roam freely within a large "playpen" constructed of plywood. As various sensory functions are added, and the robot nears completion, simple objects such as blocks of wood, and larger obstacles constructed from plywood will be introduced, and restrictions on the environment will be gradually relaxed. While this is happening, the complexity of the software will greatly increase. It is expected to take several more years of effort before the robot reaches a form which can be considered "finished." Whether the robot will ever be able to cope successfully with a "general" environment, such as a typical research laboratory or home living room, is certainly an open question.

The software will be developed as a hierarchy of modules. The bottom strata of




Photo 1: The feeding trough. This is a photograph of an ordinary wall outlet with a white concrete block wall as a background. The robot must be able to seek out and plug into outlets such as this in order to periodically replenish its energy supply.

this hierarchy will be occupied by hardware oriented routines for motor control, interrupt servicing, sensory data acquisition and the like. These routines will be relegated to EROM, where they will reside as powerful appendages for the higher level software. In the highest reaches of the hierarchy will be the planning algorithms: heuristic programming which will take into account goals, subgoals, behavior, the structure of priorities, and what may generally be called the "psychology" of the robot.

If the robot is turned loose in a strange room, its immediate behavior will be to begin exploring the room in some systematic way, locating walls and the boundaries of objects with respect to a specific origin. As new features of the environment are encountered, they will be incorporated into the world model. Particular interest will be shown in the positions and heights of recognizable wall outlets, since these are critical for survival. As the model of the room nears completion, the original intense curiosity will subside, and will be replaced by an attitude of playfulness. The world model will be consulted to find out which objects (eg: wooden blocks) are small enough to be manipulated. The robot may choose to group together objects with similar characteristics, build simple structures by stacking blocks and so forth. When the energy supply nears exhaustion, the robot will interrupt its play, head directly for a convenient wall outlet, plug in, and recharge its battery. The recharging process might take several hours, during which all motors will be turned off, and all but one 4 K byte memory module will be powered down. After recharging, the

robot will unplug itself and continue on its way. When the robot eventually "tires" of its environment, it will perhaps leave the room, wander down the hallway, and look for other rooms to explore. An interesting experiment would be to observe how long it takes for the machine to recognize a room in which it is arbitrarily placed, but has previously explored. In addition to observing general behavior, one can give the robot general tasks to perform, such as picking up all small objects in a cluttered room and placing them in a box in the corner. (Mothers with small children, take note!)

Many of the computing requirements for the robot will exceed the capabilities of its on board microcomputer. For example, the analysis of complicated scenes viewed by the image sensors can, at present, be done only by a very fast computer with lots of memory. For this reason it is planned to equip the robot with a duplex radio telemetry link to a "black box" which connects with a telephone. The robot will then be able to initiate telephone calls to a large timeshared computer, and be able to communicate with it via the 2400 bps telephone lines. Control routines in the on board microcomputer will invoke large analysis programs stored on permanent disk files attached to the timeshared computer.

A visual scene can be transmitted in less than a second, and the large computer can extract features of interest and send them back to the robot in several seconds. It will also be useful to take advantage of the off line disk files to store large data bases such as portions of the world model. When the robot has finished its transactions with the

large computer, it simply "hangs up the phone" and goes about its business, perhaps reinitiating the hookup at a later time. It is important to point out here that in no sense is the large computer "in control" of the robot. The robot is simply using the services of the large computer to perform calculations that are too involved or too lengthy to do itself, much as you or I would use a computer to solve some problem. Also, it should be noted that if the large timeshared computer is busy, or if the computations are long, the robot is free to continue with other tasks until the requested analysis can be performed.

This usage of an off line timeshared computer will permit access to a large amount of general robot planning software written in high level languages by other groups. The telemetry hookup will, as a side benefit, enable the robot to make calls to human researchers (hopefully not in the middle of the night) to report malfunctions or unusual conditions in the environment by means of coded audio signals. (The robot night watchman?)

Finding Energy Sources

It may be necessary for the robot to survive for weeks at a time without supervision or interruption of its operation. To do so, it must manage its own energy supply. The battery voltage and current are sensed periodically and converted to digital form by a simple software driven analog to digital conversion system. This information, along with the battery's internal temperature (sensed with a thermistor probe), allows the state of charge to be determined. When this reaches some minimum level, the robot is obliged to renew its supply of energy by finding a wall outlet and plugging in.

From the outset, it must be said that the robot will depend heavily on preprogrammed strategies to accomplish this task. A standard wall outlet mounted on a white concrete block wall is shown in photo 1. The cover plate is of brushed stainless steel, and the receptacles are made of light colored molded plastic. From a distance of perhaps 10 meters (33 feet), with the turret image sensor focused on infinity, the outlet appears as a small spot. Any such spot, identifiable by the software as a few dark picture elements surrounded by a light field, is worth further investigation. If several spots are visible, a single one will be selected for closer examination.

Using the motive system and the sensory turret in combination, the direction and distance to the spot is then computed by triangulation. After this is done, the robot




Photo 2: Here is an oscilloscope display of the wall outlet shown in photo 1, as viewed by an image sensor. There are a total of 1024 picture elements in this sensor, each with 16 levels of gray perceptible to the computer. The picture has 32 columns and 32 rows of elements. The outlet appears as a dark rectangle with a light background. The two receptacles can just be resolved at this range.




Photo 3: As the robot moves closer to the wall outlet, more detailed features of the socket become apparent. This is an oscilloscope display of an individual receptacle, as viewed at close range by the image sensor. By applying pattern recognition techniques to images such as this and the image in photo 2, the robot can recognize outlets and determine the exact position and height of the electrical contacts, prior to its final maneuver to plug itself into the wall.

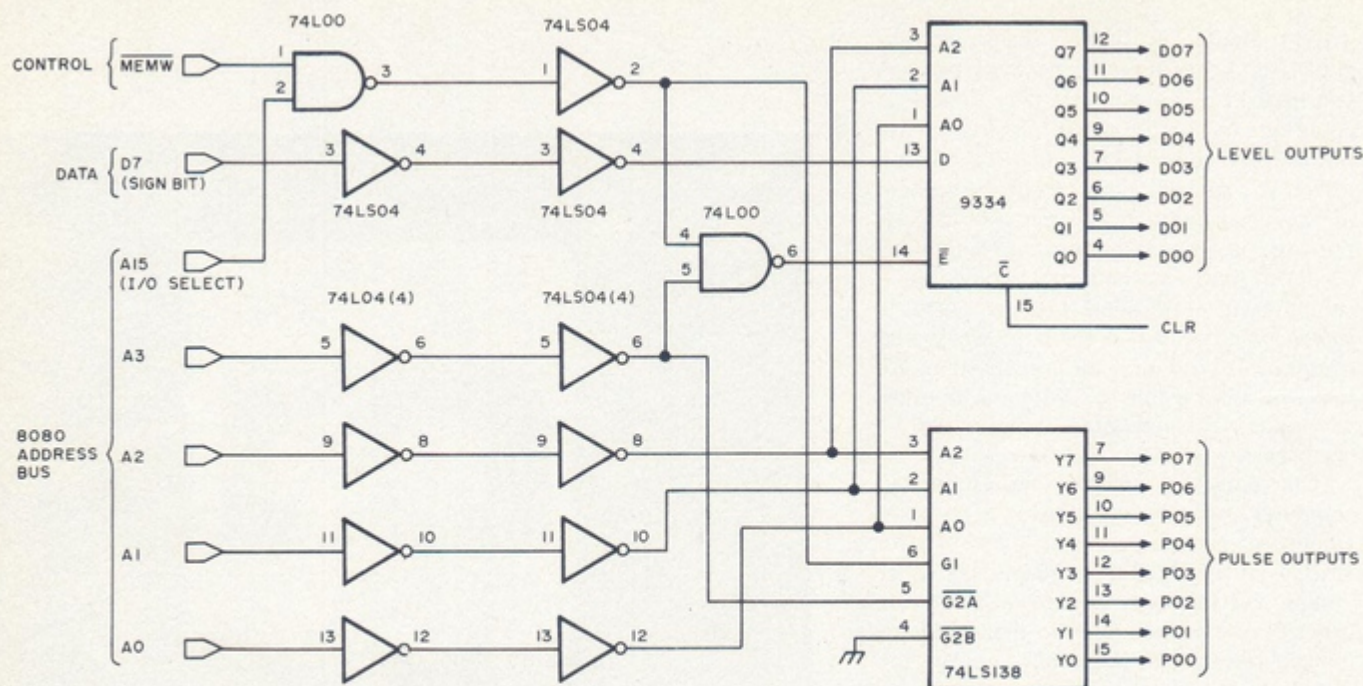


Figure 2: Output decoding logic used to create discrete (level) outputs and pulse outputs. This circuitry connects to an 8080 data bus through memory address space. The outputs are used to drive the various lines of stepper motor interfaces such as those shown in figure 3. Each stepper motor individually uses one discrete level output (for forward or reverse direction control) and one pulse output (to initiate one step sequence). Two discrete level outputs are also used to control the running versus standby, or off, status of the motors.

drives to the approximate position of the spot, and with LED proximity sensors and the tactile sensors on the manipulator, it attempts to sense and recognize the existence of a wall, which is hopefully associated with the spot. If this is the case, the robot moves out perpendicular to the wall for a fixed distance, sets the image sensor focus accordingly, and rescans for the spot which hopefully should now begin to have the shape of a wall outlet from this short range. By a series of small maneuvers, the robot positions itself exactly in front of the outlet at a precise distance away. At this time, the manipulator begins a vertical motion until the outlet is centered in the field of view of the auxiliary image sensor. The computer now has a straight on, in focus view such as that seen in photo 2.

At this point, fairly simple pattern recognizing algorithms are used to analyze the shape, size and topology of the image (feature extraction). These features are compared with known properties of wall outlets. If enough features match within certain error bounds, the recognition is successful. (In this respect, it is much easier to attempt to recognize a specific object than to try to analyze a general scene.) If the recognition is not successful, several more attempts are made at slightly different distances to over-

come the granularity effects of the sensor. If a recognition is still not successful, the robot heads off to explore other spots. Assuming success at this point, the robot moves in for a high resolution image of a single receptacle (see photo 3). Again, the image is analyzed, and further recognition tests are made. If everything is okay, the exact height and distance to the electrical contacts are determined, and the manipulator jaws close to the separation appropriate to the space between the two contacts. An attempt is now made to plug in, with the computer monitoring the voltage between the two plug prongs on the manipulator. Small searching motions are made until contact is established, whereupon the jaws open slightly to make good electrical connections. The computer then directs the recharging operation until the correct charge is reached. It is expected that the robot will normally be able to go about six hours between recharging operations.

Stepping Motor Drives

In order to execute all of the motions required of it, the robot must have precise control of the motors. This is accomplished by means of the stepping motor drive modules which constitute the interface between the on board computer and the stepping motors. There is a total of seven

identical drives: two for the motive system (one for each wheel), three for the manipulator (grasp, lift and rotate), and two for the sensory turret (pan and tilt).

Figure 2 illustrates the logic used to provide both pulse and discrete (level) outputs from the computer. For simplicity, only eight pulse outputs and eight discrete outputs are shown. The scheme can be expanded to many more outputs, with the addition of suitable logic. Low power TTL provides the connection to the address, data and control buses of the microprocessor. Lines A0 to A3 select one of 16 possible outputs, with A15 acting as an input/output select line. The sign bit, D7, is used to signify whether a one or a zero will be written into the selected discrete output line. Generation of the negative going output pulses and strobing of the data are done by the MEMW pulse. A 9334 addressable latch and 74LS138 binary to octal decoder provide the outputs to the system. If desired, one pulse output can be wired to the CLR line of the 9334 latch for simultaneous resetting of all the discrete outputs.

Figure 3 shows an individual stepping motor drive circuit which is sufficiently simple and general to permit many other non-robot applications where precise computer controlled motion is required. The circuit generates the 4 phase pulse sequence of high currents necessary to operate a Slo-Syn™ bifilar type stepping motor manufactured by the Superior Electric Company. Inputs to the stepping motor circuit are connected as desired to the discrete and pulse outputs shown in figure 2. The forward or reverse input (\overline{FR}) determines whether the motor shaft turns clockwise ($\overline{FR} = 1$) or counter-clockwise ($\overline{FR} = 0$) as viewed from the shaft end of the motor. The forward limit input (\overline{FL}) prevents the motor from moving forwards (clockwise) if $\overline{FL} = 0$. The reverse limit input (\overline{RL}) prevents the motor from moving backwards if $\overline{RL} = 0$. These inputs are generally not controlled by the computer, but are wired to simple limit switches to prevent excessive motion of the motor in a particular direction. In the case of the robot, these "reflex inputs" are independent to allow the computer to change the direction of motion after a mechanical limit has been reached, and to establish "zero point" settings for the various possible motions. The pulse input (\overline{P}) triggers on a negative going transition and causes the motor to advance one step (1.8 degrees) in the direction specified by the \overline{FR} input. Thus, 200 pulses on the \overline{P} input cause the motor shaft to rotate exactly one revolution. The maximum pulse rate is of the order of several hundred pulses per second without error, but this

depends strongly on the motor size and the driven load. The clear input (\overline{CLR}), when momentarily brought low, resets the counting flip flops to zero in case it is necessary to establish a standard shaft position when power is turned on. The off input (\overline{OFF}), when brought low, removes all current from the motor windings, and frees the shaft except for a small holding torque caused by the permanent magnets. If $\overline{OFF} = 0$ and standby (SBY) is brought high (SBY = 1), the shaft position is maintained, but at greatly reduced current. The off and standby inputs are critical for application in the robot, since they permit substantial power savings. In figure 3, series 7400 logic can be substituted for the 74L and 74LS series if desired, but at the expense of extra power. Resistance values in the drive circuits are typical, and might need to be optimized for a particular motor. The diodes in the emitter circuit of the 2N3055s should have a high surge rating, and moderate heat sinking for 2N3055s is advised.

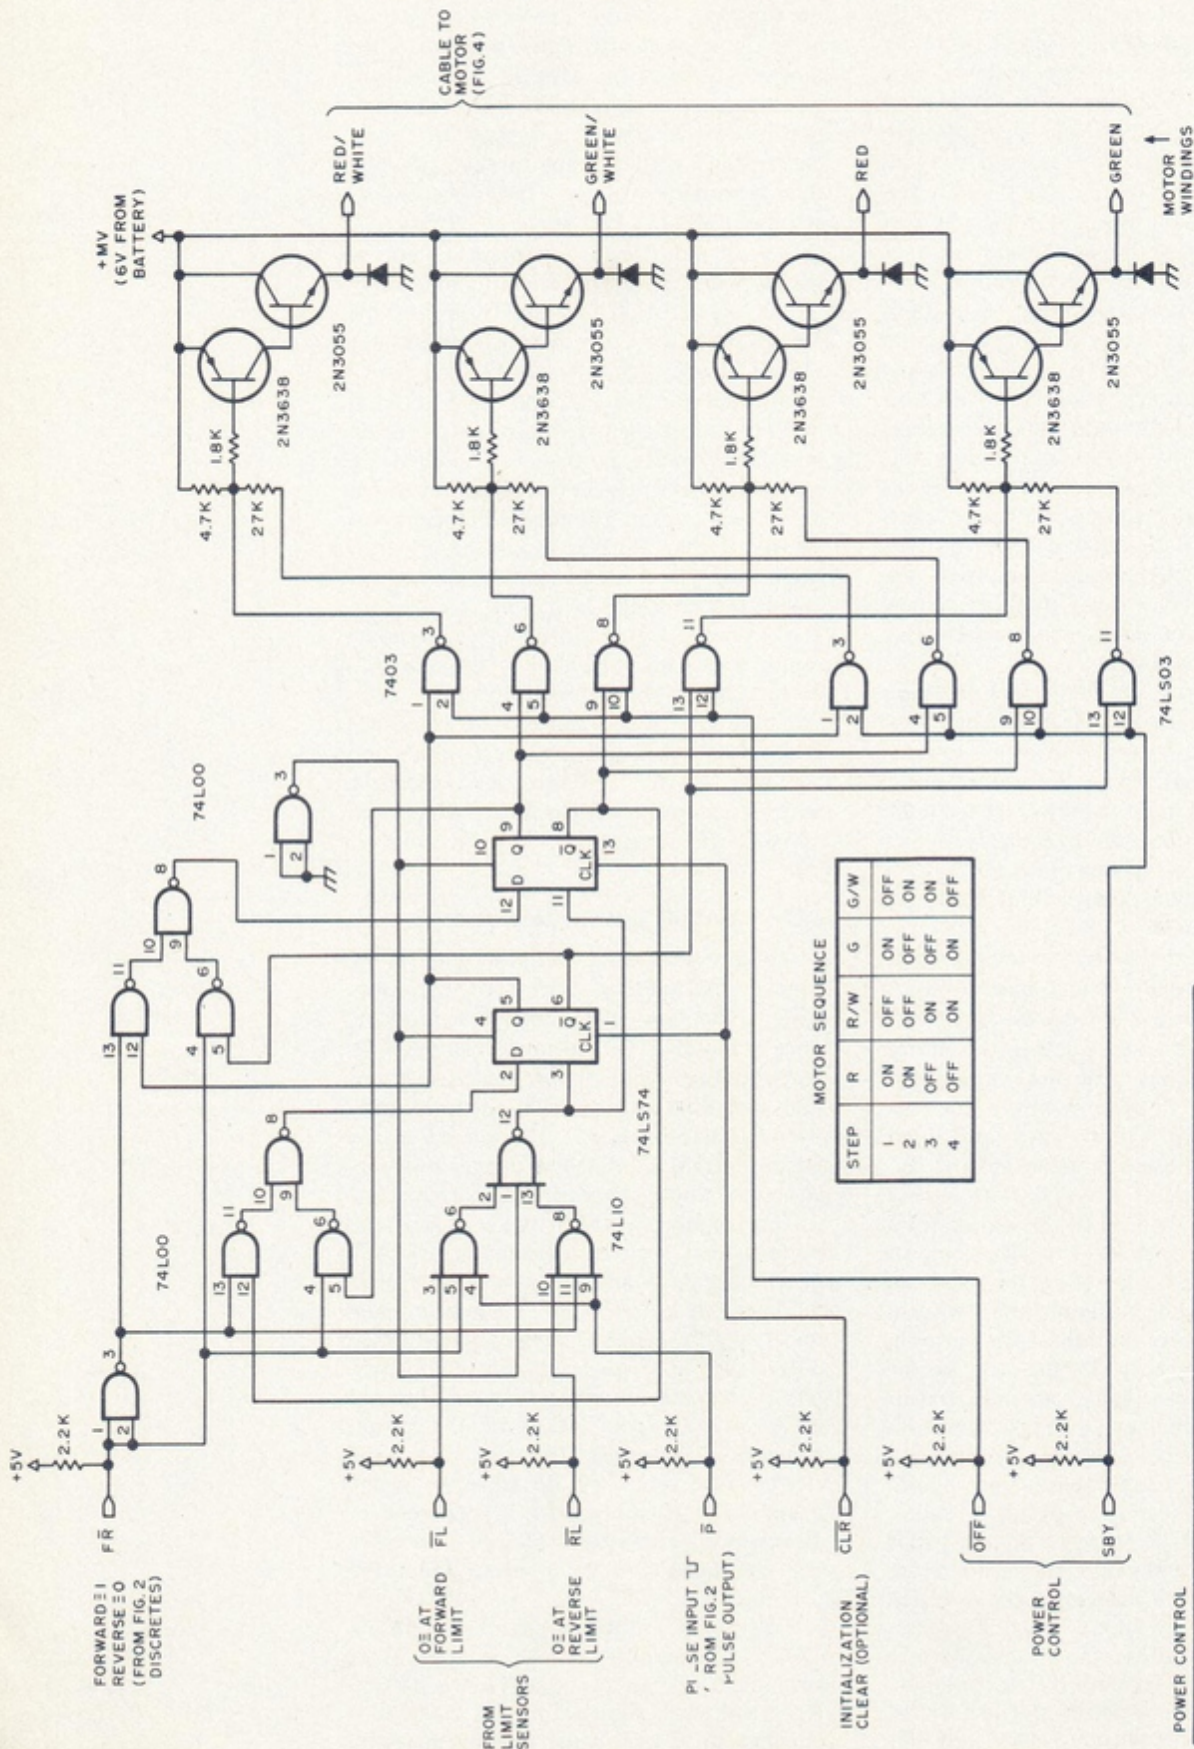
Connections to a Slo-Syn™ stepping motor are shown in figure 4. In the robot, all motors are powered directly from the main 6 V battery (+MV). Values for the resistance R can vary from zero to several ohms and are selected to be compatible with the motor's maximum current rating and dynamic characteristics.

Robot Control: Software Aspects

The previous discussion dealt with several general and specific aspects of the robot system and how it is intended to operate, once completed. Of the three major subsystems, motive, manipulator and sensory turret, only the motive subsystem is fully operational at present. The others are in various stages of detailed design and construction. Photo 4 reflects the current state of construction, which includes several temporary items such as the hand wired computer backplane and stepping motor drive modules, the crude "hand" made from sheet metal, and the cables leading to conventional power supplies. The 16 kg (35 pound) battery is on board to maintain balance, but not used, awaiting completion of the power conversion electronics. The open area at the front of the robot will be taken up by the manipulator assembly when it is completed. The power conversion modules and the sensory turret will go in place above the computer.

Software for the robot is created on the CDC-6400 computer using an 8080 cross assembler written in FORTRAN IV by Robert Mitchell. Object code for the 8080 is punched on 8 level paper tape by the 6400

Figure 3: Each stepper motor (there are seven presently incorporated into Newt's design) requires a drive circuit as shown here. This circuit contains counting logic and high current drivers needed to create the 4 step sequence of signals required by a



Slow-Syn™ stepping motor. The inputs to this interface circuit include one level (\overline{FR}) signal and one pulse signal (\overline{P}) derived from an output port such as the one shown in figure 2. The forward and reverse limit inputs are used to override commands when a sensor detects that rotation of the motor has reached some mechanical limit of the motor's activity. The power control logic (summarized in the chart) is used to conserve battery power by allowing a "standby" mode for the motor when it is not being actively driven. This standby mode holds the motor shaft position actively, but uses less current than the full torque state used to move the shaft from step to step.

system, and is subsequently read into the robot's memory through an ASR33 Teletype. A modest resident monitor program is stored in the first 666 bytes of EROM, which allows dumping and modifying of memory, punching and reading of paper tape, and branching to any memory location. The monitor is essential, because the computer has no conventional "front panel" with switches and lights.

The current robot control program occupies about 1000 bytes of programmable memory, not including table areas of variable size. It provides a method for exercising direct control of the robot for testing purposes. The coding, written by Dennis Toms, is very general, modular and compact. In order to start and stop the stepping motors without error, it is necessary to provide a profile of acceleration and deceleration. This is most conveniently done by using a time-delay table giving the appropriate time intervals between stepping motor pulses. The table is precomputed to yield a uniform acceleration over some time period. It is sufficient to have a single table for all motors. Each motor has associated with it a 14 byte motor status word (MSW) whose format is given in figure 5. The motor number and motor flag each occupy one byte; the other entries are two bytes (one word) each. The motor number byte is a fixed number which associates the MSW with a particular motor. The motor flag byte is a code which gives the current state of the motor (off, on, standby, accelerating, decelerating, emergency stop, etc). The speed pointer word is an index which points to the current entry in the time delay table. The top speed word specifies the index to the delay table entry, giving the shortest permissible delay between stepping motor pulses (maximum motor speed). The total steps word specifies the number of steps the motor is to execute. The acceleration, constant speed and deceleration counters specify the number of steps to be taken in the acceleration, constant speed and deceleration phases of the motion, respectively. These counters are decremented appropriately as the motion takes place, so that if necessary, at any time during the motion the state of the motor can be determined by

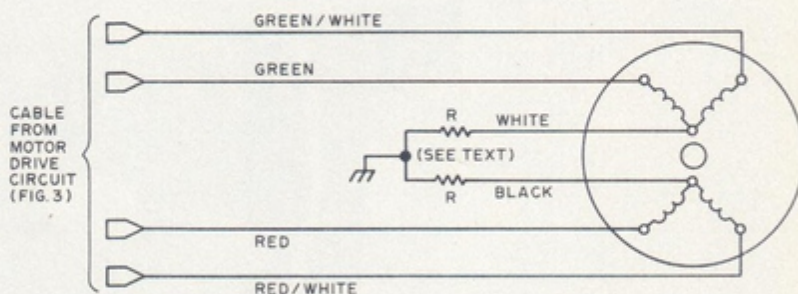


Figure 4: Wiring and color code for the Superior Electric Slow-Syn™ bifilar stepping motors. The resistance R should be set (see text) to reflect typical operating characteristics of the motor's use. The wiring is shown by color designations, and is connected to the drive circuit of figure 3. The ground return to figure 3 can be via the robot's chassis or through an additional circuit in the cabling.

interrupt handling routines. After all bytes of the appropriate MSW have been loaded with proper values, low level software takes over to carry out the motion. A pulse is issued to the selected motor causing it to advance 1.8 degrees. After a delay specified by the first entry in the delay table, a second pulse is issued, and the speed pointer is incremented to point to the second value in the delay table, and the acceleration counter is decremented. This process continues until the acceleration counter reaches zero. When this occurs, the acceleration phase is complete and the constant speed phase is entered. For this phase, the speed pointer is

Figure 5: Format of the motor status word. There is one motor status word allocated to each stepper motor; an interrupt driven process updates and times the operations of the motors. The information includes motor identification, status flags and parameters which specify the details of a cycle of operation consisting of acceleration, constant running speed and deceleration.

Motor Status Word (MSW)	
Motor Number	Motor Flag
Speed Pointer	
Top Speed	
Total Steps	
Acceleration Counter	
Constant Speed Counter	
Deceleration Counter	




Photo 4: The robot Newt as it appears in the current state of construction. The hand wired computer backplane and stepping motor drive modules and sheet metal "hand" and power cables to an external source are all temporary.

held fixed, pointing to the time delay value corresponding to the top speed index, while the constant speed counter is decremented for each step. After the constant speed phase is completed, the speed pointer is decremented for each pulse, causing the motor to decelerate. The motor finally stops after a number of pulses have been issued equaling total steps. The motor is then placed in

standby condition and this fact is stored in the motor flag byte. If unforeseen conditions arise which make it inadvisable or impossible to complete the motion as planned, the interrupt software stores the appropriate emergency stop code in the motor flag and the motor is halted. In all but extreme emergency cases requiring instant stopping of a motor running at high speed,

the software can recover the total number of steps actually taken, saving it for navigational purposes.

To exercise direct control of the robot, a simple interpretive command system was developed. The currently implemented commands are listed in table 1. Each command consists of the 1 byte ASCII character F, B, L, R, W, V or Q, followed by a 2 byte argument. In the case of the F, B, L and R commands, the argument is simply the number of motor steps to be taken (less than or equal to $2^{15} - 1$). For the W command, the argument is the waiting time in units of 10ths of a second. For the V command, the argument specifies the new maximum speed of the robot in steps per second. For the Q command, the argument gives the 16 bit address of a sequence of commands stored in memory. Of course, many other commands, such as those appropriate for sensory turret and manipulator motion, will be added to the list of table 1. A program of robot activity consists of a simple sequence of these 3 byte commands, one after the other. For example, the following 15 byte sequence (arguments shown in hexadecimal notation):

```
F 0A10
L 000B
B 0010
W 0066
B 0002
```

would cause the robot to move forward 2576 steps (about 1.4 meters), turn left 11 steps (a few degrees), back up for 16 steps, stop and wait for 10.2 seconds, and then back up two steps. All accelerations and decelerations are taken care of automatically by the lower level software.

Commands can be given to the robot in several different ways (command modes) which are specified in table 2. For the D mode, the robot simply executes the commands directly as they are entered on the Teletype. After each command is given, it is necessary to wait until it is carried out before entering another command. In the R mode nothing is executed, but each command is recorded in sequence in the memory. The N mode is a combination of the D and R modes. It permits a sequence of commands to be recorded as they are being executed. The P mode allows playback of any previously recorded sequence. In this mode, the Teletype cable can be disconnected to allow the robot to roam freely. If the C mode is specified, the robot creeps along at a very slow pace under an F, B, L or R command for an indefinite number of steps. The motion is terminated by de-

Table 1: Currently implemented commands. The demonstrations of operation seen in photo 5 were created using an interpretive sequence including these seven basic commands. The command list is open ended in that many additions to the possible operations are expected as the software is developed further.


ASCII Command Code	Command Definition
F	Move robot forward
B	Move robot backward
L	Rotate robot left
R	Rotate robot right
W	Make robot wait
V	Set robot maximum speed
Q	Execute command sequence

pressing any key on the Teletype, after which the number of steps taken is printed out. The T mode is a combination of the C and R modes. The E, L and O modes are "bookkeeping" in nature, and permit erasing, listing, and setting the origin of a command sequence, respectively. During the creation of a program of action for the robot (command sequence), any combination of modes can be used as desired. It is easy to think of many other useful modes which could be added to the list of table 2.


Using the commands listed in table 1, it has been possible to do a number of preliminary experiments with the robot which serve to test both the hardware and software, and to prove some of the fundamental ideas. For example, one can place wooden alphabet blocks on the floor at random, and then program the robot (using the N and T modes) to "pick them up" by trapping them in its temporary "hand." After putting the blocks and the robot back in their initial positions, the playback (P) mode can be used to repeat the motions automatically. One such test, rather in the form of a demonstration, is shown in photos 5a to 5g. In the first frame, photo 5a, the robot is moving towards the nearest block from its

Table 2: Command modes for the robot. The software presently implemented is used to try out various exercises of the robot's machinery. The software is structured into several command modes described by this list.

Command Mode	Type of Operation
D	Direct (simply execute the commands)
R	Remember (record commands)
N	Normal (combine D and R)
P	Playback (run through a command sequence)
C	Creep (directly execute motions at a turtle's pace)
T	Teach (combine C and R)
E	Erase
L	List
O	Set origin




(e)



(f)

initial position in the left corner of the playpen. In 5b, it is maneuvering into position for "picking up" the block. In 5c, the block has been picked up, moved and released in a new position. In 5d, a second block has been fetched and placed next to the first block, and the robot is turning around and heading for a third block. In 5e, it is about ready to pick up the third block. In 5f, the third block has been placed, and the fourth block is in the "hand." In the final frame 5g, the fourth block has been placed, spelling out N E W T, which is, of course, the robot's name. Approximately 100 3 byte commands were required for this sequence. The program can be executed dozens of times without error, and the final positions of the blocks and the robot after the sequence is finished has a scatter of a few millimeters ($\pm 1/8$ inch) in each direction. Of course, when the robot is complete with all its senses, it will seldom have to execute such a lengthy sequence in a completely open loop fashion. Perhaps it can eventually do tasks such as searching for and finding specific alphabet blocks in a random pile of blocks. This would require frequent closing of the feedback loops through the visual senses.

As the robot goes through its motions, such as depicted in photos 5a to 5g, it seems to possess an almost uncanny grace and precision. Small children, when watching it, are frightened at first, but this soon gives way to playful interest and warm curiosity. Even hardened computer experts are amazed to see a computer driving itself around on wheels!



(g)

Some Personal Remarks on Building Robots

It will not be possible for one person alone to write the software for the robot. This is far too large a job. It is hoped, however, that as the hardware nears completion, it will act as a focal point for many persons wanting to experiment with the robot by writing their own software. A

fascinating project would be to create a general "Robot Control Language" which would free each programmer from the details of the hardware. What a rich experience it would be to work together, exchanging ideas in a highly interactive way. There are many tutorial possibilities for computer

science classes and beginning students learning about computers. There are problems in psychology, procedural languages, human-to-machine communication, functioning of parts of the brain . . . (Working with robots is certainly one way to gain a much greater appreciation for the complex-

SOME TERMINOLOGY

Amp hour: A unit of energy for rating batteries. The battery of NEWT, with its 84 Amp hour capacity, can store enough charge to drive a steady 1 A load for 84 hours, or a 14 A load for six hours. The finite capacity of any practical battery means that any mobile robot must incorporate some programs for seeking electrical outlets and recharging batteries periodically.

Azimuth: As used here, an angle relative to a fixed direction in the horizontal plane.

Cognition: As Webster has it, this is "the act or process of knowing, including both awareness and judgement." [*Webster's New Collegiate Dictionary*, 1976 edition.] In the context of robots and artificial intelligence, this term refers to programmed models which approximate the behavior of natural cognition.

Degrees of freedom: The state of a robot mechanism (or any other system) can be described by specifying the current value of each variable parameter. Thus, if a robot arm has seven joints, the position of its "hand" might be determined by the angular setting of each joint. Each such independently variable parameter of a system is called a "degree of freedom," so the seven jointed arm would have seven degrees of freedom.

Heuristic: A heuristic computer program is one which starts out with an approximate method of solving of a problem within the context of some goal, and uses feedback from the effects of the solution to improve upon its own performance. Heuristic programming is one of the major contemporary artificial intelligence techniques, and is a key to developing a cognitive robot.

Manipulator systems: A generic term for any mechanical device which a robot uses to directly manipulate its environment. In the NEWT robot, this is currently (see photos) a simple sheet metal frame which can catch a block and slide it across the floor as the robot moves, with no active grasping; NEWT is intended to eventually have a much more flexible system of manipulation as described in the text. Most industrial robots currently in use consist of manipulators alone, without much in the way of sensory feedback or motive systems.

Motive systems: A generic term for the mechanisms used to convey the robot around its environment. In the NEWT robot, this refers to the two drive wheels, balancing caster and stepper motors which propel the robot.

Open loop, closed loop: A closed loop system is one which operates with feedback from errors. The feedback is intended to correct for the errors and thus approach the truth; an open loop system ignores error signals and operates on the sometimes

naive assumption that no errors occur. The terms must be qualified by a reference to the time intervals involved in the system: NEWT, for example, is a closed loop system over long time periods, since it is intended to navigate using feedback from its sensors; however, due to the processing loads associated with sensors, NEWT operates open loop between navigation sightings in a manner analogous to the dead reckoning method of navigation used occasionally by airplane pilots or captains of ships.

Round off errors: In operations such as addition, multiplication or calculating transcendental functions, there is often some uncertainty in the least significant part of the result. In an extended calculation in which these operations are repeated over and over, appreciable round off errors can accumulate. In a digitally controlled vehicle guidance system such as that used for a robot, these numerical errors are a major source of uncertainty in the vehicle position, and are just as important as more obvious sources of error such as step quantization or slippage in the drive mechanism.

Senses: In a robotics context, senses are specialized peripherals which convert information about the environment into signals which can be analyzed by a computer or used directly by the electronics, as in a reflex. Sensory information may be obtained from devices as simple as a microswitch with a "feeler" arm, or as complicated as photoelectric imaging arrays with zoom lenses and pointing mechanisms.

Stepping motors: An ordinary electric motor is characterized by continuous motion when energized. A stepping motor uses a different design philosophy to achieve a motor which will move its shaft in small incremental angular steps on command, and will actively maintain its position in between each command. This type of motor is very well adapted to digital control of mechanical systems, and is used by NEWT for all mechanical motions in the robot.

Step quantization: The stepper motors have a finite angular resolution built into their design. This means that any mechanical motion derived from the motor will have a certain minimum step size, so that any attempt to position to a finer tolerance must be approximated.

Trajectory: The path of a moving object is its trajectory. In the case of the mobile robot, a trajectory is planned before motion takes place, given a desired goal position and a world model which covers its course and objects which may be in the way.

World model: A world model is the result of cognition as implemented in robots. Formally, it is an information structure built up in the memory of the robot, based on both initialization and heuristic interaction with the environment.

ity and capability of the brain.) In addition, experiments can continue on sensor development and interfaces from sensor to computer. There are a great many practical spin-offs from this kind of work.

Many people believe that as more and more advances are made in microelectronics, the prospects of mass producing robots will become attractive, and the prices of these hypothetical machines will plummet. (Let us hope we will have learned something from *RUR*.) If this occurs, many applications will open up. Besides such things as planetary surface exploration, such as already demonstrated by the Viking robots, one can envision undersea robots working on oil pipelines and well heads, coal mining robots, fire fighting robots, agricultural robots, robots on assembly lines producing customized articles, robot-like prosthetic devices, and many other types of robots for specialized and general service, doing jobs which are too difficult, too dangerous, or which are otherwise undesirable for humans.

I have been asked many times the questions, "Why are you building a robot?" and "What will it do when it is finished?" The answer to the second of these questions is easy: I simply don't know what the robot will be able to do. This is the whole point of building the robot. Given a modest amount of hardware and a greater amount of software, thoroughly integrated to form a system, the idea is to find out just what such a system is capable of doing. The whole is likely to be far greater than the sum of the parts. The system is pushed as far as the available time, money and energy will allow in order to learn what can be done and what cannot be done; in other words, to explore the frontier of robot research, and to know and understand the problems involved. The necessary knowledge can neither be obtained by theoretical studies, nor by simulations using large computers.

In regard to the first of the questions, I have been fascinated with robots since the mid-1950s and have constructed several robot devices prior to the one described in this article. The construction of such machines presents many interesting challenges. A functioning robot is a most curious blend of electronics, mechanics, computer design, computer programming and artificial intelligence. All these fields come together in the design and construction of a robot, and each must be explored in depth. Added to this are the challenges and excitement of locating obscure components in surplus store parts bins around the country, planning, building, and then replanning, rebuilding, and, of course, experimenting and learning. To me,

About NEWT's Name and Family Tree

The origin of NEWT's name is buried in an often quoted verse from *Macbeth* by William Shakespeare . . .

*In the cauldron boil and bake;
Eye of newt and toe of frog,
Wool of bat and tongue of dog,
Adder's fork and blind-worm's sting,
Lizard's leg and howlet's wing,
For a charm of powerful trouble,
Like a hell-broth boil and bubble.*

Newt I, the present robot's predecessor, was a light-seeking robot consisting of a large eye on a stalk rising above a motor driven platform.

all these are fascinating aspects of endeavors which are perhaps best left to amateurs. I say this, because I strongly believe that the amateur computer enthusiast has a golden opportunity to participate in advances in the field of robotics. In fact, the amateur has several advantages over the professional. The research can be as abstract as the amateur wishes it to be and can be conducted without regard to immediate payoff potential in the marketplace. There is no need to spend time writing elaborate proposals, no need to continually justify the direction of the work, and no need to get hard results every few months to write up and stick into quarterly reports. History has shown that precisely this atmosphere of freedom which surrounds the amateur is the atmosphere in which brilliant innovations and discoveries are sometimes made. ■

Acknowledgment: The author wishes to express gratitude to Dennis Toms for his enthusiastic help and interest in this project.