MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A. I. Memo 864 **September, 1985**

A ROBUST LAYERED CONTROL SYSTEM FOR A MOBILE ROBOT

Rodney A. Brooks

Abstract. We describe a new architecture for controlling mobile robots. Layers of control system are built to let the robot operate at increasing levels of competence. Layers are made up of asynchronous modules which communicate over low bandwidth channels. Each module is an instance of a fairly simple computational machine. Higher level layers can subsume the roles of lower levels by suppressing their outputs. However, lower levels continue to function as higher levels are added. The result is a robust and flexible robot control system. The system is intended to control a robot that wanders the office areas of our laboratory, building maps of its surroundings. In this paper we demonstrate the system controlling a detailed simulation of the robot.

Acknowledgments. This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the research is provided in part by an IBM Faculty Development Award, in part by a grant from the Systems Development Foundation, in part by an equipment grant from Motorola, and in part by the Advanced Research Projects Agency under Office of Naval Research contracts N00014–80–C–0505 and N00014–82–K–0334.

© Massachusetts Institute of Technology 1985

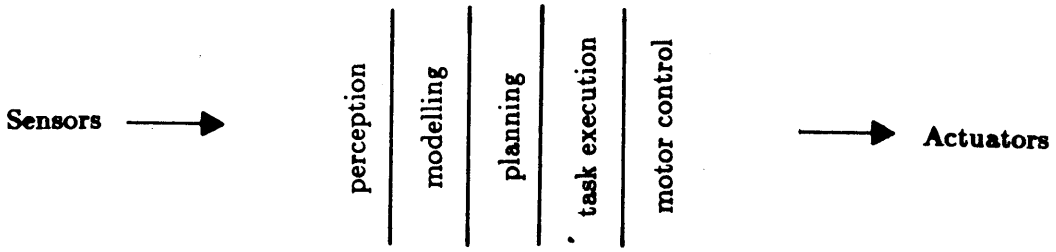

Figure 1. A traditional decomposition of a mobile robot control system into functional modules.

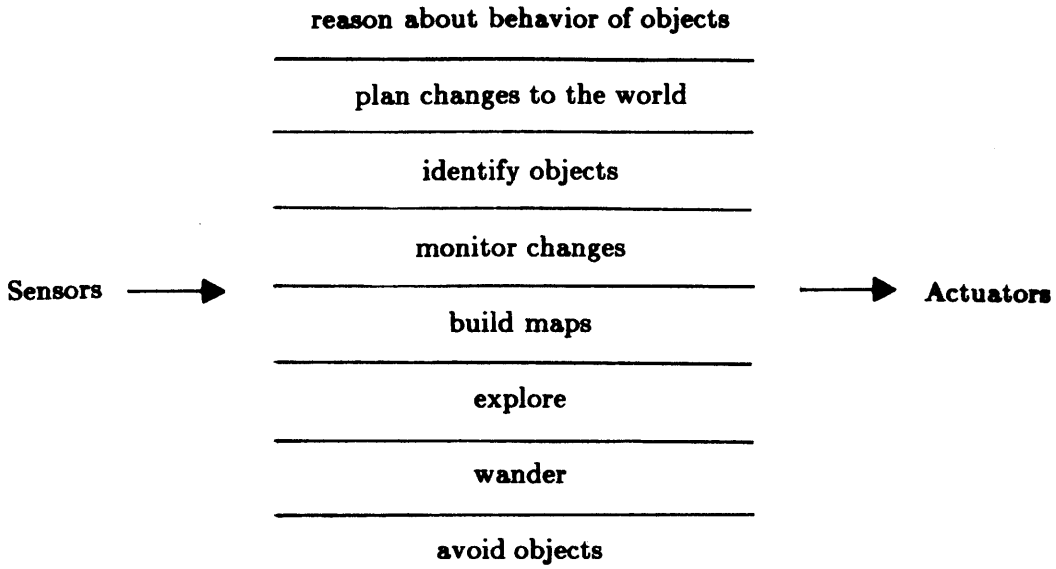

Figure 2. A decomposition of a mobile robot control system based on task achieving behaviors.

1. Introduction

A control system for a completely autonomous mobile robot must perform many complex information processing tasks in real time. It operates in an environment where the boundary conditions (viewing the instantaneous control problem in a classical control theory formulation) are changing rapidly. In fact the determination of those boundary conditions is done over very noisy channels since there is no straightforward mapping between sensors (e.g. TV cameras) and the form required of the boundary conditions.

The usual approach to building control systems for such robots is to decompose the problem into a series (roughly) of *functional units* as illustrated by a series of vertical slices in figure 1. After analysing the computational requirements for a mobile robot we have decided to use *task achieving behaviors* as our primary decomposition of the problem. This is illustrated by a series of horizontal slices in figure 2. As with a functional decomposition we implement each slice explicitly then tie them all together to form a robot control system. Our new decomposition leads to a radically different architecture for mobile robot control systems, with radically different implementation strategies plausible at the hardware level, and with a large number of advantages concerning robustness, buildability and testability.

1.1 Requirements

We can identify a number of requirements of a control system for an intelligent autonomous mobile robot. They each put constraints on possible control systems we might build and employ.

- **Multiple Goals.** Often the robot will have multiple goals, some conflicting, which it is trying to achieve. It may be trying to reach a certain point ahead of it while avoiding local obstacles. It may be trying to reach a certain place in minimal time while conserving power reserves. Often the relative importance of goals will be context dependent. Getting off the railroad tracks when a train is heard becomes much more important than inspecting the last 10 track ties of the current track section. The control system must be responsive to high priority goals, while still servicing necessary "low level" goals (e.g. in getting off the railroad tracks it is still important that the robot maintains its balance so it doesn't fall down).
- **Multiple Sensors.** The robot will most likely have multiple sensors (e.g. TV cameras, encoders on steering and drive mechanisms, and perhaps infrared beacon detectors, an inertial navigation system, acoustic rangefinders, infrared rangefinders, access to a global positioning satellite system, etc.). All sensors have an error component in their readings. Furthermore, often there is no direct analytic mapping from sensor values to desired physical quantities. Some of the sensors will overlap in the physical quantities they measure. They will often give inconsistent readings—sometimes due to normal sensor error and sometimes due to the measurement conditions being such that the sensor (and subsequent processing) is being used outside its domain of applicability. Often there will be no analytic characterization of the domain of applicability (e.g. under what precise conditions does the Sobel operator return valid edges?). The robot must make decisions under these conditions.
- **Robustness.** The robot ought to be robust. When some sensors fail it should be able to adapt and cope by relying on those still functional. When the environment changes drastically it should be able to still achieve some modicum of sensible behavior, rather then sit in shock, or wander aimlessly or irrationally around. Ideally it should also continue to function well when there are faults in parts of its processor(s).
- **Additivity.** As more sensors and capabilities are added to a robot it needs more processing power; otherwise the original capabilities of the robot will be impaired relative to the flow of time.

1.2 Other Approaches

- **Multiple Goals.** [Elfes and Talukdar 83] designed a control language for [Moravec 83]'s robot which tried to accommodate multiple goals. It mainly achieved this by letting the user explicitly code for parallelism and to code an exception path to a special handler for each plausible case of unexpected conditions.
- **Multiple Sensors.** [Flynn 85] explicitly investigated the use of multiple sensors, with complementary characteristics (sonar is wide angle but reasonably accurate in depth, while infrared is very accurate in angular resolution but terrible in depth measurement). Her system has the virtue that if one sensor fails the other still delivers readings that are useful to the higher level processing. [Giralt et al 83] use a laser range finder for map making, sonar sensors for local obstacle detection, and infrared beacons for map calibration. The robot operates in a mode where one particular sensor type is used at a time and the others are completely ignored, even though they may be functional. In the natural world multiple redundant sensors are abundant. For instance [Kreithen 83] reports that pigeons have more than four independent orientation sensing systems (e.g. sun position compared to internal biological clock). It is interesting that the sensors do not seem to be combined but rather, depending on the environmental conditions and operational level of sensor subsystems, the data from one sensor tends to dominate.
- **Robustness.** The above work tries to make systems robust in terms of sensor availability, but little has been done with making either the behavior or the processor of a robot robust.
- **Additivity.** There are three ways this can be achieved without completely rebuilding the physical control system. (1) Excess processor power which was previously being wasted can be utilized. Clearly this is a bounded resource. (2) The processor(s) can be upgraded to an architecturally compatible but faster system. The original software can continue to run, but now excess capacity will be available and we can proceed as in the first case. (3) More processors can be added to carry the new load. Typically systems builders then get enmeshed in details of how to make all memory uniformly accessible to all processors. Usually the cost of the memory to processor routing system soon comes to dominate the cost (the measure of cost is not important—it can be monetary, silicon area, access time delays, or something else) of the system. As a result there is usually a fairly small upper bound (on the order of hundreds for traditional style processing units; on the order to tens to hundreds of thousands for extremely simple processors) on the number of processors which can be added.

1.3 Starting Assumptions

Our design decisions for our mobile robot are based on nine dogmatic principles (six of these principles were presented more fully in [Brooks 85]):

- Complex (and useful) behavior need not necessarily be a product of an extremely complex control system. Rather, complex behavior may simply be the reflection of a complex environment [Simon 69]. It may be an observer who ascribes complexity to an organism—not necessarily its designer.
- Things should be simple. This has two applications. (1) When building a system of many parts one must pay attention to the interfaces. If you notice that a particular

interface is starting to rival in complexity the components it connects, then either the interface needs to be rethought or the decomposition of the system needs redoing. (2) If a particular component or collection of components solves an unstable or ill-conditioned problem, or, more radically, if its design involved the solution of an unstable or ill-conditioned problem, then it is probably not a good solution from the standpoint of robustness of the system.

- We want to build cheap robots which can wander around human inhabited space with no human intervention, advice or control and at the same time do useful work. Map making is therefore of crucial importance even when idealized blue prints of an environment are available.
- The human world is three dimensional; it is not just a two dimensional surface map. The robot must model the world as three dimensional if it is to be allowed to continue cohabitation with humans.
- Absolute coordinate systems for a robot are the source of large cumulative errors. Relational maps are more useful to a mobile robot. This alters the design space for perception systems.
- The worlds where mobile robots will do useful work are not constructed of exact simple polyhedra. While polyhedra may be useful models of a realistic world, it is a mistake to build a special world such that the models can be exact. For this reason we will build no artificial environment for our robot.
- Sonar data, while easy to collect, does not by itself lead to rich descriptions of the world useful for truly intelligent interactions. Visual data is much better for that purpose. Sonar data may be useful for low level interactions such as real time obstacle avoidance.
- For robustness sake the robot must be able to perform when one or more of its sensors fails or starts giving erroneous readings. Recovery should be quick. This implies that built-in self calibration must be occurring at all times. If it is good enough to achieve our goals then it will necessarily be good enough to eliminate the need for external calibration steps. To force the issue we do not incorporate any explicit calibration steps for our robot. Rather we try to make all processing steps self calibrating.
- We are interested in building *artificial beings*—robots which can survive for days, weeks and months, without human assistance, in a dynamic complex environment. Such robots must be self sustaining.

1.4 The Physical Robot

In the rest of the paper we describe a layered robot control system. It is intended to drive an actual robot. To give some concreteness to the cosntraints on the control system we describe that robot here.

We have constructed a mobile robot shown in photo 1. It is about 17 inches in diameter and about 30 inches from the ground to the top platform. Most of the processing occurs offboard on Lisp machines.

The drive mechanism was purchased from Real World Interface of Sudbury, Massachusetts. Three parallel drive wheels are steered together. The two motors are servoed by a single microprocessor. The robot body is attached to the steering mechanism and always

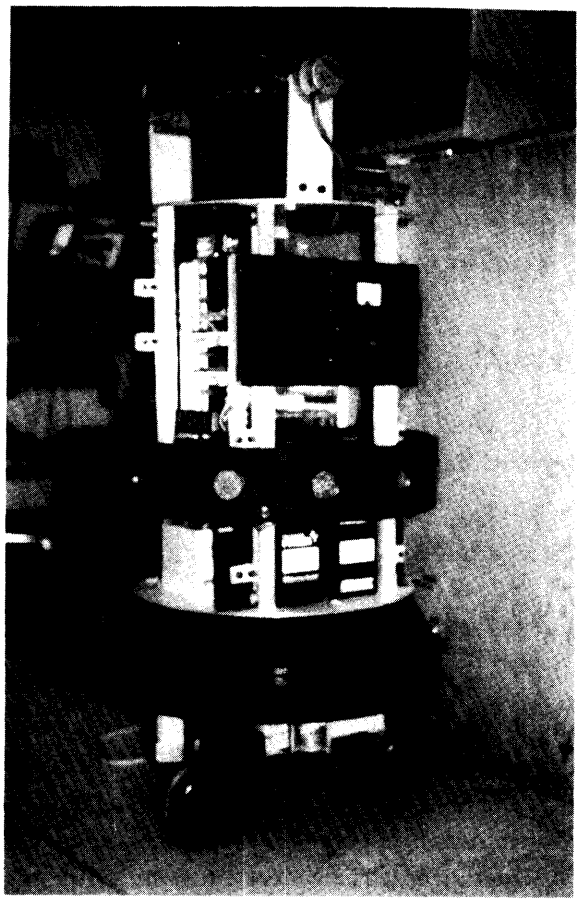

Photo 1. The MIT AI Lab mobile robot.

points in the same direction as the wheels. It can turn in place (actually it inscribes a circle about 1 cm in diameter).

Currently installed sensors are a ring of twelve Polaroid sonar time of flight range sensors and two Sony CCD cameras. The sonars are arranged symmetrically around the rotating body of the robot. The cameras are on a tilt head (pan is provided by the steering motors). We plan to install feelers which can sense objects at ground level about six inches from the base extremities.

A central cardcage contains the main onboard processor, an Intel 8031. It communicates with offboard processors via a 12Kbit/sec duplex radio link. The radios are modified Motorola digital voice encryption units. Error correction cuts the effective bit rate to less than half the nominal rating. The 8031 passes commands down to the motor controller processor and returns encoder readings. It controls the sonars, and the tilt head and switches the cameras through a single channel video transmitter mounted on top of the robot. The latter transmits a standard TV signal to a Lisp machine equipped with a demodulator and frame grabber. Multiple Lisp machines work together to control the robot, communicating

rather infrequently over a Chaos net.

2. Levels and Layers

There are many possible approaches to building an autonomous intelligent mobile robot. As with most engineering problems they all start by decomposing the problem into pieces, solving the subproblems for each piece, and then composing the solutions. We think we have done the first of these three steps differently to other groups. The second and third steps also differ as a consequence.

2.1 Levels of Competence

Typically mobile robot builders (e.g. [Nilsson 84], [Moravec 83], [Giralt et al 83], [Kanayama 83], [Tsuji 84], [Crowley 85]) have sliced the problem into some subset of:

- sensing,
- mapping sensor data into a world representation,
- planning,
- task execution, and
- motor control.

This decomposition can be regarded as a horizontal decomposition of the problem into vertical slices. The slices form a chain through which information flows from the robot's environment, via sensing, through the robot and back to the environment, via action, closing the feedback loop (of course most implementations of the above subproblems include internal feedback loops also). An instance of each piece must be built in order to run the robot at all. Later changes to a particular piece (to improve it or extend its functionality) must either be done in such a way that the interfaces to adjacent pieces do not change, or the effects of the change must be propagated to neighboring pieces, changing their functionality too.

We have chosen instead to decompose the problem vertically as our primary way of slicing up the problem. Rather than slice the problem on the basis of internal workings of the solution we slice the problem on the basis of desired external manifestations of the robot control system.

To this end we have defined a number of *levels of competence* for an autonomous mobile robot. A level of competence is an informal specification of a desired class of behaviors for a robot over all environments it will encounter. A higher level of competence implies a more specific desired class of behaviors.

We have used the following levels of competence (an earlier version of these was reported in [Brooks 84]) as a guide in our work:

0. Avoid contact with objects (whether the objects move or are stationary).

1. Wander aimlessly around without hitting things.

2. "Explore" the world by seeing places in the distance which look reachable and heading for them.

3. Build a map of the environment and plan routes from one place to another.

4. Notice changes in the "static" environment.

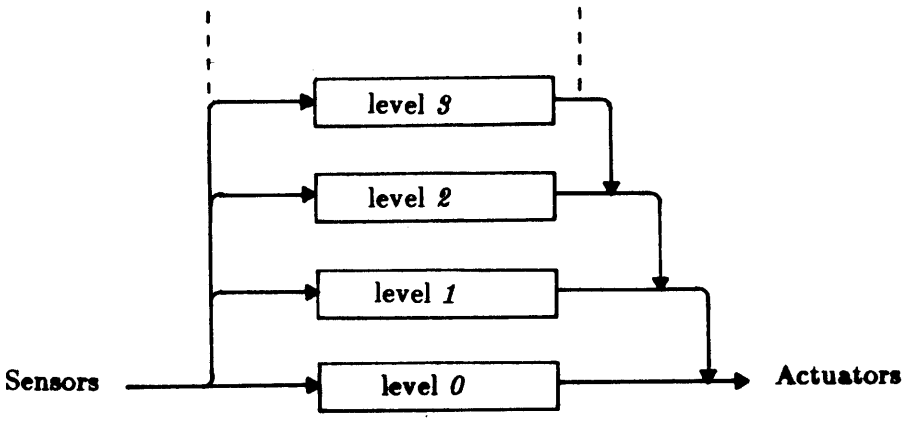

Figure 3. Control is layered with higher level layers subsuming the roles of lower level layers when they wish to take control. The system can be partioned at any level, and the layers below form a complete operational control system.

5. Reason about the world in terms of identifiable objects and perform tasks related to certain objects.
6. Formulate and execute plans which involve changing the state of the world in some desirable way.
7. Reason about the behavior of objects in the world and modify plans accordingly.

Notice that each level of competence includes as a subset each earlier level of competence. Since a level of competence defines a class of valid behaviors it can be seen that higher levels of competence provide additional constraints on that class.

2.2 Layers of Control

The key idea of levels of competence is that we can build layers of a control system corresponding to each level of competence and simply add a new layer to an existing set to move to the next higher level of overall competence.

We start by building a complete robot control system which achieves level *0* competence. It is debugged thoroughly. We never alter that system. We call it the zeroth level control system, Next we build a another control layer, which we call the first level control system. It is able to examine data from the level *0* system and is also permitted to inject data into the internal interfaces of level *0* suppressing the normal data flow. This layer, with the aid of the zeroth, achieves level *1* competence. The zeroth layer continues to run unaware of the layer above it which sometimes interferes with its data paths.

The same process is repeated to achieve higher levels of competence. See figure 3.

We call this architecture a *subsumption architecture.*

In such a scheme we have a working control system for the robot very early in the piece—as soon as we have built the first layer. Additional layers can be added later, and the initial working system need never be changed.

We claim that this architecture naturally lends itself to solving the problems for mobile robots delineated in section 1.1.

- **Multiple Goals.** Individual layers can be working on individual goals concurrently. The suppression mechanism then mediates the actions that are taken. The advantage here is that there is no need to make an early decision on which goal should be pursued. The results of pursuing all of them to some level of conclusion can be used for the ultimate decision.
- **Multiple Sensors.** In part we can ignore the sensor fusion problem as stated earlier using a subsumption architecture. Not all sensors need to feed into a central representation. Indeed certain readings of all sensors need not feed into central representations—only those which perception processing identifies as extremely reliable might be eligible to enter such a central representation. At the same time however the sensor values may still be being used by the robot. Other layers may be processing them in some fashion and using the results to achieve their own goals, independent of how other layers may be scrutinizing them.
- **Robustness.** Multiple sensors clearly add to the robustness of a system when their results can be used intelligently. There is another source of robustness in a subsumption architecture. Lower levels which have been well debugged continue to run when higher levels are added. Since a higher level can only suppress the outputs of lower levels by actively interfering with replacement data, in the cases that it can not produce results in a timely fashion the lower levels will still produce results which are sensible, albeit at a lower level of competence.
- **Additivity.** An obvious way to handle additivity is to make each new layer run on its own processor. We will see below that this is practical as there are in general fairly low bandwidth requirements on communication channels between layers. In addition we will see that the individual layers can easily be spread over many loosely coupled processors.

2.3 The Structure of Layers

But what about building each individual layer? Don't we need to decompose a single layer in the traditional manner? This is true to some extent, but the key difference is that we don't need to account for all desired perceptions and processing and generated behaviors in a single decomposition. We are free to use different decompositions for different sensor-set task-set pairs.

We have chosen to build layers from a set of small processors which send messages to each other.

Each processor is a finite state machine with the ability to hold some data structures. Processors send messages over connecting "wires". There is no handshaking or acknowledgement of messages. The processors run completely asynchronously, monitoring their input wires, and sending messages on their output wires. It is possible for messages to get lost—it actually happens quite often. There is no other form of communication between processors, in particular there is no shared global memory.

All processors (which we refer to as modules) are created equal in the sense that within a layer there is no central control. Each module merely does its thing as best it can.

Inputs to modules can be suppressed and outputs can be inhibited by wires terminating from other modules. This is the mechanism by which higher level layers subsume the role of lower levels.

3. A Robot Control System Specification Language

There are two aspects to the components of our layered control architecture. One is the internal structure of the modules, and the second is the way in which they communicate. In this section we flesh out the details of the semantics of our modules and explain a description language for them.

3.1 Finite State Machines

Each module, or processor, is a finite state machine, augmented with some instance variables which can actually hold Lisp data structures.

Each module has a number of input lines and a number of output lines. Input lines have single element buffers. The most recently arrived message is always available for inspection. Messages can be lost if a new one arrives on an input line before the last was inspected.

There is a distinguished input to each module called **reset**.

Each state is named. When the system first starts up all modules start in the distinguished state named **NIL**. When a signal is received on the reset line the module switches to state **NIL**. A state can be specified as one of four types.

- **Output.** An output message, computed as a function of the module's input buffers and instance variables, is sent to an output line. A new specified state is then entered.

- **Side effect.** One of the module's instance variables is set to a new value computed as a function of its input buffers and variables. A new specified state is then entered.

- **Conditional dispatch.** A predicate on the module's instance variables and input buffers is computed and depending on the outcome one of two subsequent states is entered.

- **Event dispatch.** An sequence of pairs of conditions and states to branch to are monitored until one of the events is true. The events are in combinations of arrivals of messages on input lines and the expiration of time delays.*

*The exact semantics are as follows. After an event dispatch is executed all input lines are monitored for message arrivals. When the next event dispatch is executed it has access to latches which indicate whether new messages arrived on each input line. Each condition is evaluated in turn. If it is true then the dispatch to the new state happens. Each condition is an and/or expression on the input line latches. In addition, condition expressions can include delay terms which become true a specified amount of time after the beginning of the execution of the event dispatch. An event dispatch waits until one of its condition expressions is true.

An example of a module defined in our specification language is the **avoid** module:

```
(defmodule avoid
  :inputs (force heading)
  :outputs (command)
  :instance-vars (resultforce)
  :states
    ((nil (event-dispatch (and force heading) plan))
     (plan (setf resultforce (select-direction force heading))
           go)
     (go (conditional-dispatch (significant-force-p resultforce 1.0)
                               start
                               nil))
     (start (output command (follow-force resultforce))
            nil)))
```

Here, `select-direction`, `significant-force-p` and `follow-force` are all lisp functions, while `setf` is the modern lisp assignment special form.

The force input line inputs a force with magnitude and direction found by treating each point found by the sonars as the site of a repulsive force decaying as the fifth power of distance. The fifth power was chosen as the force drops off reasonably quickly as objects are further away than the manouevering resolution of the robot, but increase dramatically as objects become very close to the robot. Function `select-direction` takes this and combines it with the input on the heading line considered as a motive force. It selects the instantaneous direction of travel by summing the forces acting on the robot. (This simple technique computes the tangent to the minimum energy path computed by [Khatib 83].)

Function `significant-force-p` checks whether the resulting force is above some threshold—in this case it determines whether the resulting motion would take less than a second. The dispatch logic then ignores such motions.

Function `follow-force` converts the desired direction and force magnitude into motor velocity commands.

This particular module is part of the level *1* control system described below. It essentially does local navigation, making sure obstacles are avoided by diverting a desired heading away from obstacles. It does not deliver the robot to a desired location—that is the task of level *2* competence.

3.2 Communication

Figure 4 shows the best way to think about these finite state modules for the purposes of communications. They have some input lines and some output lines. An output line from one module is connected to input lines of one or more other modules. One can think of these lines as wires, each with sources and a destination.

Additionally outputs may be inhibited, and inputs may be suppressed.

An extra wire can terminate (i.e. have its destination) at an output site of a module. If *any* signal travels along this wire it *inhibits* any output message from the module along

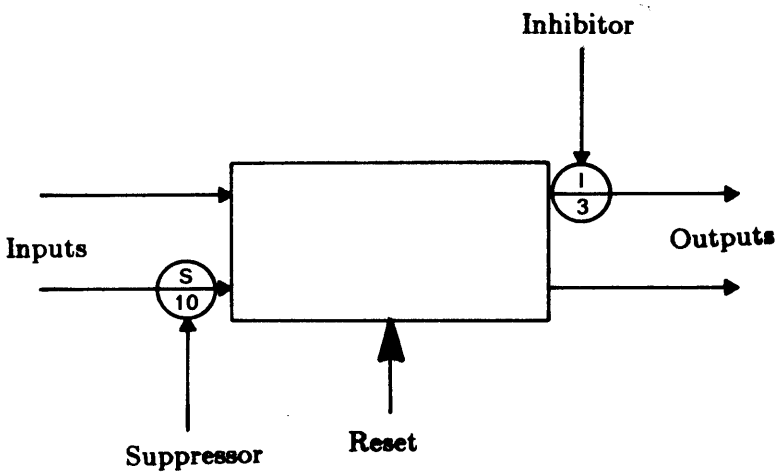

Figure 4. A module has input and output lines. Input signals can be suppressed and replaced with the suppressing signal. Output signals can be inhibited. A module can also be reset to state NIL.

that line for some pre-determined time. Any messages sent by the module to that output during that time period is lost.

Similarly an extra wire can terminate at an input site of a module. Its action is very similar to that of inhibition, but additionally, the signal on this wire, besides inhibiting signals along the usual path, actually gets fed through as the input to the module. Thus it *suppresses* the usual input and provides a replacement. If more than one suppressing wire is present they are essentially 'or'-ed together.

For both suppression and inhibition we write the time constants inside the circle.

In our specification language we write wires as a source (i.e. an output line) followed by a number of destinations (i.e. input lines). For instance the connection to the force input of the **avoid** module defined above might be the wire defined as:

```
(defwire (feelforce force) (runaway force) (avoid force))
```

This links the force output of the **feelforce** module to two inputs.

Suppression and inhibition can also be described with a small extension to the syntax above. Below we see the suppression of the command input of the **motor** module, a level *0* module by a signal from the level *1* module **avoid**.

```
(defwire (avoid command) ((suppress (motor command) 1.5)))
```

In a similar manner a signal can be connected to the reset input of a module.

4. A Robot Control System Instance

We have implemented a mobile robot control system to achieve levels *0* and *1* competence as defined above, and have started implementation of level *2* bringing it to a stage which exercises the fundamental subsumption idea effectively. We need more work on an early vision algorithm to complete level *2*.

The complete code for the modules and interconnections of levels *0*, *1*, and *2* is given in the Appendix.

4.1 Zeroth Level

The lowest level layer of control makes sure that the robot does not come into contact with other objects. It thus achieves level *0* competence. If something approaches the robot it will move away. If in the course of moving itself it is about to collide with an object it will halt. Together these two tactics are sufficient for the robot to flee from moving obstacles, perhaps requiring many motions, without colliding with stationary obstacles. The robot is not invincible of course, and a sufficently fast moving object, or a very cluttered environment might result in a collision.

- The **motor** module communicates with the actual robot. The module has an extra communication mechanism, allowing it to send and receive commands to and from the physical robot directly. It accepts a command specifying angle and speed of turn, magnitude of forward travel and velocity. It sends a high, or busy, signal on a status line it maintains then waits in a delay-induced loop for it to be finished. If at any time a halt message is received, it commands the robot to halt. (Any additional motion commands received during transit are lost.) At the termination of a motion, whether by completion or halting, the module sets the status output line to low.
- The **sonar** module takes a vector of sonar readings, filters them for invalid readings, and effectively produces a robot centered map of obstacles in polar coordinates.
- The **collide** module monitors the sonar map and if it detects objects dead ahead it sends a signal on the halt line to the **motor** module. The **collide** module does not know or care whether the robot is moving. Halt messages sent while the robot is stationary are essentially lost.
- The **feelforce** module sums the results of considering each detected object as a repulsive force, generating a single resultant force.
- The **runaway** module monitors the 'force' produced by the sonar detected obstacles and sends command to the **motor** module if it ever becomes significant.

Figure 5 gives a complete description of how the modules are connected together. Notice that the status line of the **motor** module is not used at this level.

4.2 First Level

The first level layer of control, when combined with the zeroth, imbues the robot with the ability to wander around aimlessly without hitting obstacles. This was defined earlier as level *1* competence. This control level relies in a large degree on the zeroth level's aversion to hitting obstacles. In addition it uses a simple heuristic to plan ahead a little in order to avoid potential collisions which would need to be handled by the zeroth level.

- The **wander** module generates a new heading for the robot every 10 seconds or so.
- The **avoid** module, described in more detail in section 3, takes the result of the force computation from the zeroth level, and combines it with the desired heading to produce a modfied heading which usually points in roughly the right direction, but is perturbed

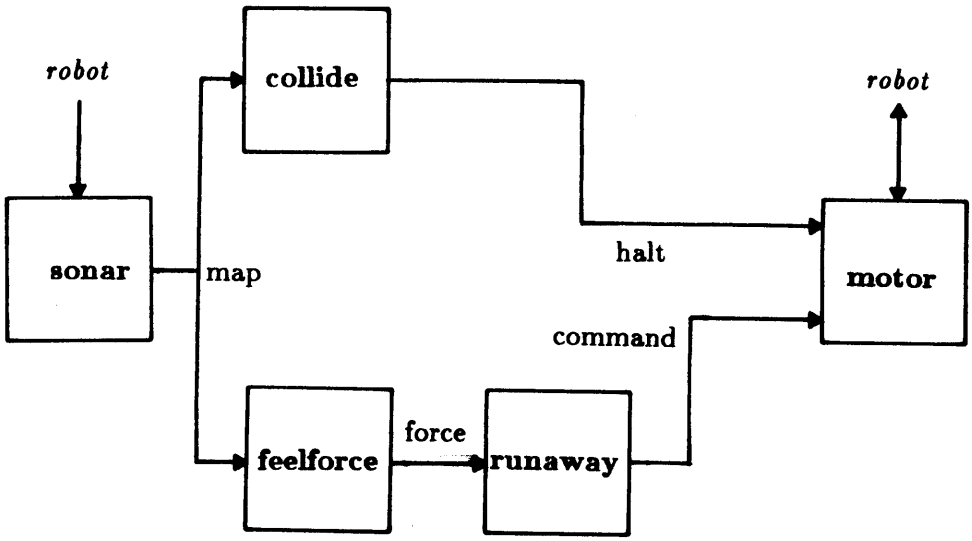

Figure 5. The level *0* control system.

to avoid any obvious obstacles. This computation implicitly subsumes the computations of the **runaway** module, in the case that there is also a heading to consider. In fact the output of the **avoid** module suppresses the output from the **runaway** module as it enters the **motor** module.

Figure 6 gives a complete description of how the modules are connected together. Note that it is simply figure 5 with some more modules and wires added.

4.3 Second Level

Level two is meant to add an exploratory mode of behavior to the robot, using visual observations to select interesting places to visit. At the time of writing only the non-vision aspects of this level have been fully implemented. They provide a means of position servoing the robot to a desired relative position despite the presence of local obstacles on its path (as detected with the sonar sensing system). The wiring diagram is shown in figure 7. Note that it is simply figure 6 with some more modules and wires added.

- The **grabber** module ensures that level *2* has control of the motors by sending a halt signal to the **motor** module, then temporarily inhibiting a number of communications paths in the lower levels so that no new actions can be initiated (thus for a brief two seconds the robot is unable to avoid approaching objects), and waiting for the **motor** module to indicate that it is no longer controlling a robot motion. At this point the sensors will be giving stable readings sufficient to plan a detailed motion, so a goal can be sent to the **pathplan** module.
- The **monitor** module continually monitors the status of the **motor** module. As soon as that module becomes inactive the **monitor** module queries the robot via a direct connection to get a reading from its shaft encoders on how far it has travelled. Thus

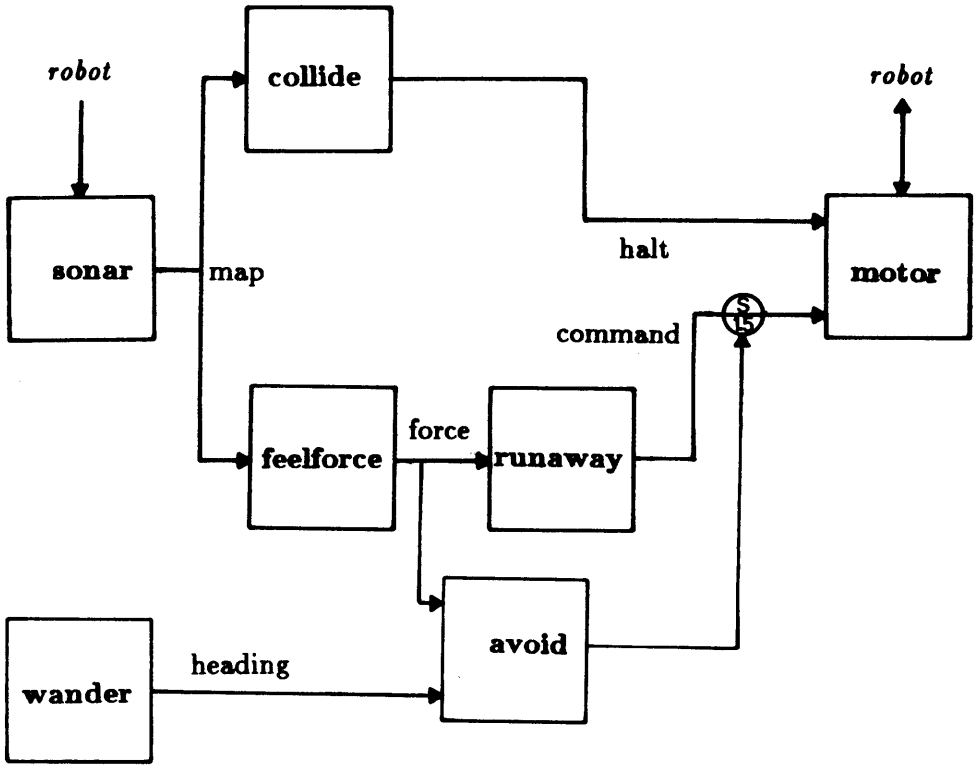

Figure 6. The level *0* control system augmented with the level *1* system.

it is able to track each motion, whether it terminated as intended or if there was an early halt due to looming obstacles.

- The **integrate** module accumulates reports of motions from the **monitor** module and always sends its most recent result out on its integral line. It gets restarted by application of a signal to its reset input.
- The **pathplan** module takes a goal specification (in terms of an angle to turn, a distance to travel, and a final orientation) and attempts to reach that goal. To do this it sends headings to the **avoid** module, which may perturb them to avoid local obstacles, and monitors its integral input which is an integration of actual motions. The messages to the **avoid** module suppress random wanderings of the robot, so long as the higher level planner remains active. When the position of the robot is close to the desired position (the robot is unaware of control errors due to wheel slippage etc., so this is a dead reckoning decision) it outputs the goal to the **straighten** module.
- The **straighten** module is responsible for modifying the final orientation of the robot. Any command with just an angular heading will not get through the avoid module as it filters out small motions, since the pressure of forces from remote obstacles would otherwise make it "buzz" with large turns and small forward motions. Therefore the final orienting step of a planned path needs to use another mechanism. Since the

robot only has to turn in place there is no danger of a collision, although there is always the danger of an approaching object. The **straighten** module thus sends its commands directly to the **motor** module, and monitors the integral line to make sure it is successful. For good measure it also inhibits the collision reflex, since there is no chance of a collision from forward motion, but the **collide** module might indicate there is a collision imminent as the robot rotates past an obstacle. This extra control signal is not necessary, as the **straighten** module would reinitiate another turn if the first were unsuccessful, but it does make for a smoother final reorientation.

The current wiring of the second level of control is shown in figure 7, augmenting the two lower level control systems. Note that four of the inhibition and suppression nodes are used only to gain control of the lower levels of the system, not to supplant it during operation of the second layer. The zeroth and first layers still play an active role during normal operation of the second layer.

The remainder of control for level *2* competence is not implemented at the time of writing. It will take stereo depth data [Grimson 85], produce descriptions of freeways, or corridors of space, and send commands to the **goal** line of the **pathplan** module. In addition extra monitoring of the **integral** ouput of the **integrate** module will ensure that if the robot wanders outside of the perceived limits of the corridor of free space then it will stop and take more visual observations. This last capability will make additional use of the reset capability to halt the path planning behavior.

5. Performance

At the time of this writing (August 1985) the physical robot and communications links are not yet complete enough to support tests of the layered control system. We have, however, tested all of layers *0* and *1* and part of layer *2* on a simulated robot.

5.1 A Simulated Robot

The simulation tries to simulate all the errors and uncertainties that exist in the world of the real robot. When commanded to turn through angle α and travel distance d the simulated robot actually turns through angle $\alpha+\delta\alpha$ and travels distance $d+\delta d$. Its sonars can bounce off walls multiple times, and even when they do return they have a noise component in the readings modeling thermal and humidity effects. We feel it is important to have such a realistic simulation. Anything less leads to incorrect control algorithms.

The simulator runs off a clock and runs at the same rate as would the actual robot. It actually runs on the same processor that is simulating the subsumption architecture. Together they are nevertheless able to perform a realtime simulation of the robot and its control and also drive graphics displays of robot state and module performance monitors. Figure 8 shows the robot (which itself is not drawn) receiving sonar reflections at some of its 12 sensors. Other beams did not return within the time allocated for data collection. The beams are being reflected by various walls. There is a small bar in front of the robot perpendicular to the direction the robot is pointing.

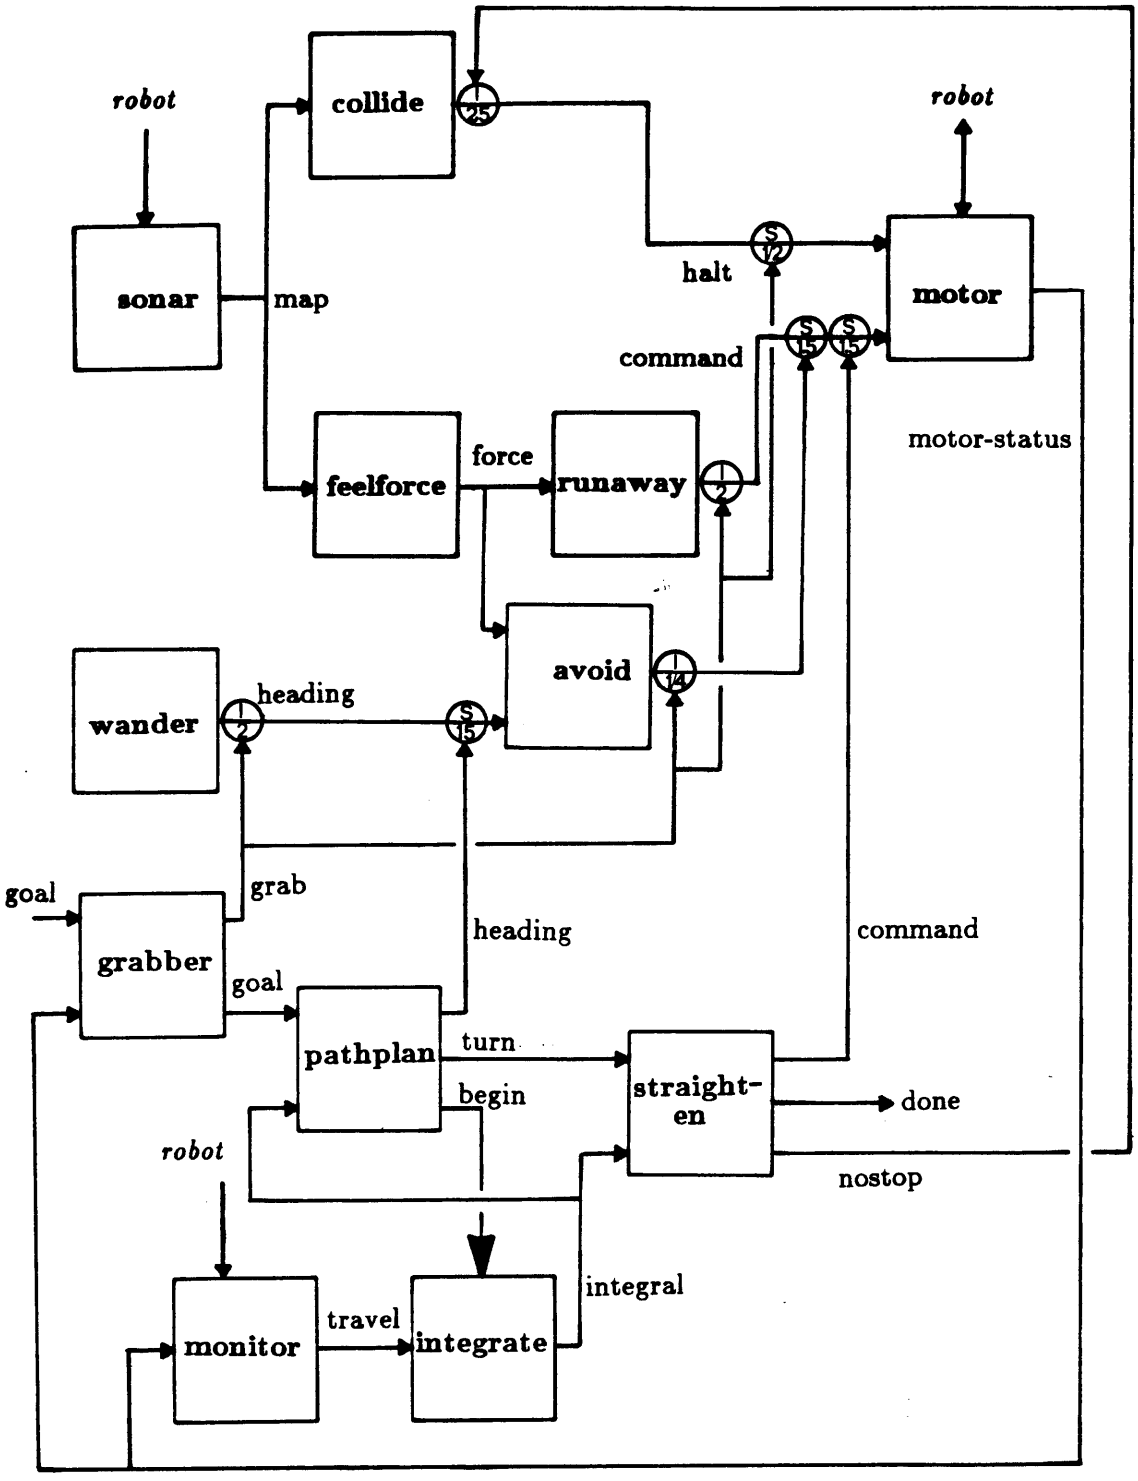

Figure 7. The level *0* and *1* control systems augmented with the level *2* system.

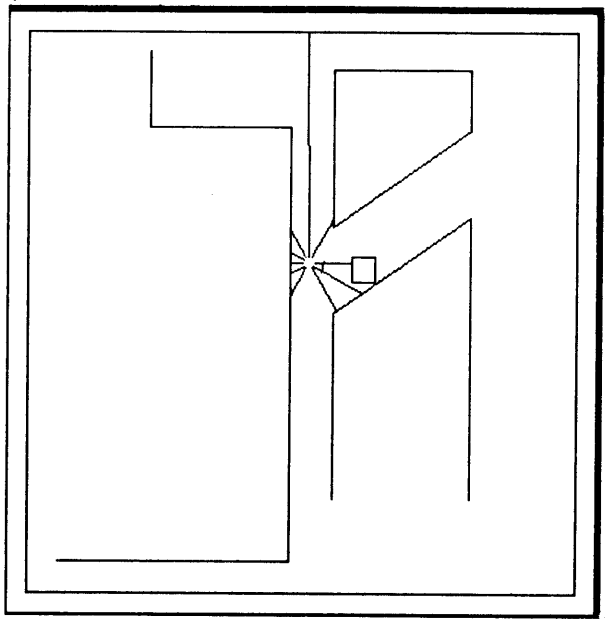

Figure 8. The simulated robot receives 12 sonar readings. Some sonar beams glance off walls and do not return within a certain time.

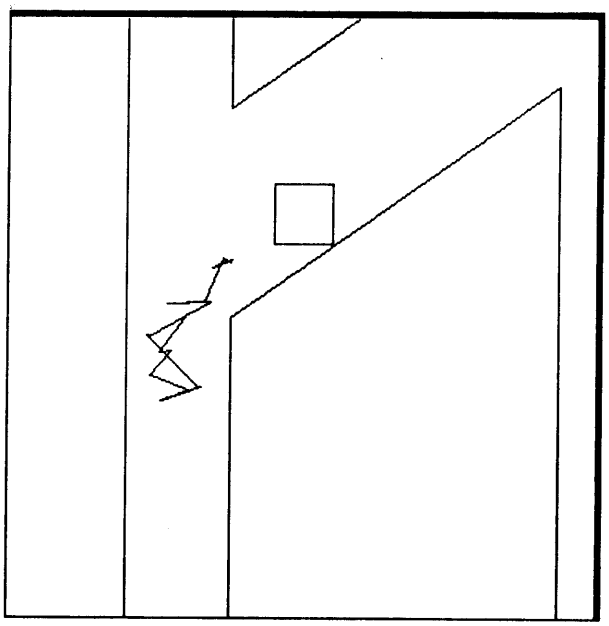

Figure 9. Under levels *0* and *1* control the robot wanders around aimlessly. It does not hit obstacles.

5.2 Zeroth and First Level

Figure 9 shows an example world in two dimensional projection. The simulated robot with a first level control system connected was allowed to wander from an initial position. The squiggly line traces out its path. Note that it was wandering aimlessly and that it hit no obstacles.

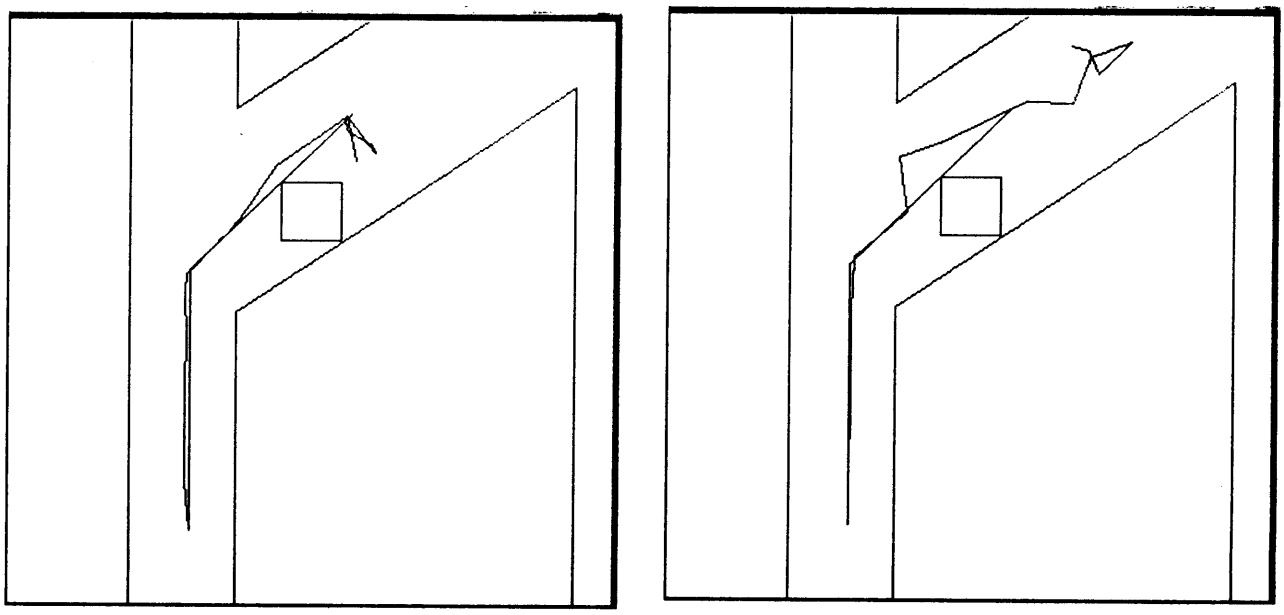

Figure 10. With level *2* control the robot tries to achieve commanded goals. The nominal goals are the two straight lines. After reaching the second goal, since there are no new goals forthcoming, the robot reverts to aimless level *1* behavior.

5.3 Second Level

Figure 10 shows two examples of the same scene and the motion of the robot with the second level control system connected. In this case two goals were fed to the goal line of the **grabber** module. The two straight lines show the nominal commanded path. While achieving these goals the lower level wandering behavior was suppressed. The goals were not reached exactly. The simulator models a uniformly distributed error of $\pm 5\%$ in both turn and forward motion. As soon as the goals had been achieved satisfactorily the robot reverted to its wandering behavior.

6. Implementation Issues

One of the motivations for developing the layered control system was additivity of processing power. The fact that it is decomposed into asynchronous processors with low bandwidth communication and no shared memory should certainly assist in achieving that goal. New processors can simply be added to the network by connecting their inputs and outputs at appropriate places—there are no bandwith or synchronization considerations in such connections.

The bug in this idea is that our modules, while themselves simple finite state machines, have access to Lisp programs (e.g. `select-direction`) which require additional computational power.

6.1 A Spatial Processor

In building the control system we noted that all non-trivial uses of Lisp (e.g., some pred-

icates used are so trivial that they can be implemented as a single gate) relate to spatial reasoning. Approximately 60 trivial processors, each with a one bit alu, arranged in a circular network seem sufficient to carry out the hard Lisp computations of any given module. The communication network topology would model the space around the robot. Such a network could easily be fitted a single chip. Connection machine processors seem suitable for such arrays [Hillis 85].

6.2 Sizing The Processors

The finite state processors need not be large. Sixteen states is more than sufficient for all modules we have so far written. (Actually 8 states are sufficient under the model of the processors we have presented here and used in our simulations. However we have refined the design somewhat towards gate level implementation and there we use simpler more numerous states.) Such processors could easily be put in the corner of an array processor chip, or many could be packed on a single chip.

6.3 Shorter Term Approaches

While we have been able to simulate sufficient processors on a single Lisp Machine up until now, that capability will soon pass as we bring on line our vision work (the algorithms have been debugged as traditional serial algorithms but we plan on re-implementing them within the subsumption architecture). Building the architecture in custom chips is a long term goal. In the shorter term the connection machine [Hillis 85] seems bendable to support a large scale simulation. The communication within modules can use the NEWS network, and the between module communication can use the cross omega network. The low bandwidth intermodule communication requirements suggest that the network will not soon be saturated.

7. Conclusion

The key ideas in this paper are:

- The mobile robot control problem can be decomposed in terms of behaviors rather than in terms of functional modules.
- It provides a way to incrementally build and test a complex mobile robot control system.
- Useful parallel computation can be performed on a low bandwidth loosely coupled network of asynchronous simple processors. The topology of that network is relatively fixed.
- There is no need for a central control module of a mobile robot. The control system can be viewed as a system of agents each busy with their own solipsist world.

Besides leading to a different implementation strategy it is also interesting to note the way the decomposition affected the capabilities of the robot control system we have built. In particular our control system deals with moving objects in the environment at the very lowest level, and has a specific module (**runaway**) for that purpose. Traditionally mobile robot projects have delayed handling moving objects in the environment beyond the scientific life of the project.

Most lacking in this paper are:

- A demonstration of the control system with a real robot.
- Details of how the full level *2* control system will be implemented.
- A clear separation of the algorithms for control of the robot from the implementation medium. We felt this was necessary to convince the reader of the utility of both. It is unlikely that the subsumption architecture would appear to be useful without a clear demonstration of how a respectable and useful algorithm can run on it. Mixing the two descriptions as we have done demonstrates the proposition.
- We are not completely happy with the fact that a level *0* module has an output (the motor-status output of the **motor** module) which is not used within that level but only in level *2*.

Nevertheless we are confident that the methodology presented in this paper will continue to bear fruit as we bring it up on our physical robot.

Acknowledgement

Tomás Lozano-Pérez, Eric Grimson, Jon Connell and Anita Flynn have all provided helpful comments on earlier drafts of this paper.

References

[**Brooks 84**] "Aspects of Mobile Robot Visual Map Making", Rodney A. Brooks, *Robotics Research 2, Hanafusa and Inoue eds, MIT Press*, 1984, 369–375.

[**Brooks 85**] "Visual Map Making for a Mobile Robot", Rodney A. Brooks, *IEEE Conference on Robotics and Automation, St Louis*, March, 1985, 824–829.

[**Crowley 85**] "Navigation for an Intelligent Mobile Robot", James L. Crowley, *IEEE Journal of Robotics and Automation*, RA-1, March 1985, 31–41.

[**Elfes and Talukdar 83**] "A Distributed Control System for the CMU Rover", Alberto Elfes and Sarosh N. Talukdar, *Proceedings IJCAI, Karlsruhe, West Germany*, August 1983, 830–833.

[**Flynn 85**] "Redundant Sensors for Mobile Robot Navigation", Anita Flynn, *MS Thesis, Department of Electrical Engineering and Computer Science, MIT*, July 1985.

[**Giralt et al 83**] "An Integrated Navigation and Motion Control SYstem for Autonomous Multisensory Mobile Robots", Georges Giralt, Raja Chatila, and Marc Vaisset, *Robotics Research 1, Brady and Paul eds, MIT Press*, 1983, 191–214.

[**Grimson 85**] "Computational Experiments with a Feature Based Stereo Algorithm", W. Eric L. Grimson, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-7, January 1985, 17–34.

[**Hillis 85**] "The Connection Machine", Danny Hillis, *PhD Thesis, Department of Electrical Engineering and Computer Science, MIT*, May 1985.

[**Kanayama 83**] "Concurrent Programming of Intelligent Robots", Yutaka Kanayama, *Proceedings IJCAI, Karlsruhe, West Germany*, August 1983, 834–838.

[**Khatib 83**] "Dynamic Control of Manipulators in Operational Space", Oussama Khatib, *Sixth IFTOMM Congress on Theory of Machines and Mechanisms, New Delhi*, December 1983.

[**Kreithen 83**] "Orientational Strategies in Birds: A Tribute to W. T. Keeton", Melvin L. Kreithen, in *Behavioral Energetics: The Cost of Survival in Vertebrates, Ohio State University Press*, 1983, 3–28.

[**Moravec 83**] "The Stanford Cart and the CMU Rover", Hans P. Moravec, *Proceedings of the IEEE*, 71, July 1983, 872–884.

[**Nilsson 84**] "Shakey the Robot", Nils J. Nilsson, *SRI AI Center Technical Note 323*, April 1984.

[**Simon 69**] "Sciences of the Artificial", Herbert A. Simon, *MIT Press*, 1969.

[**Tsuji 84**] "Monitoring of a Building Environment by a Mobile Robot", Saburo Tsuji, *Robotics Research 2, Hanafusa and Inoue eds, MIT Press*, 1984, 349–356.

Appendix

On the following pages is the complete control code for the mobile robot control system instance described in this paper. The lisp functions which compute forces etc. are omitted. All modules and wires are conditionalized with reader macros that identify which layer they belong in. A read macro conditionally tells the Lisp reader whether to process or ignore the foloowing expression. Thus expressions prefixed by `#+level1` are only processed if some global statement has said that level *1* modules and wires should be included. In this way we can use a single source file, but still control which layers get loaded.

```
#+level0
(defmodule motor
  :inputs (command halt)
  :outputs (motor-status)
  :states
    ((nil (event-dispatch command drive))
     (drive (output motor-status hi)
            send-motion)
     (send-motion (send *robot* :move command)
                  transit)
     (transit (event-dispatch halt halt-robot
                              (delay 0.5) idle-p))
     (halt-robot (send *robot* :halt)
                 transit)
     (idle-p (conditional-dispatch (lo? *bus-status*) done transit))
     (done (output motor-status lo)
           nil)))

#+level0
(defmodule sonar
  :inputs ()
  :outputs (map)
  :states
    ((nil (event-dispatch (delay 1.0) read))
     (read (output map (convert-sonar-array-to-map *bus-readings*))
       nil)))

#+level0
(defmodule collide
  :inputs (map)
  :outputs (halt)
  :states
    ((nil (event-dispatch map monitor))
     (monitor (conditional-dispatch (dangerous-motion-p map 5.0)
                                    quit
                                    nil))
     (quit (output halt hi)
           nil)))

#+level0
(defmodule feelforce
  :inputs (map)
  :outputs (force)
  :states
    ((nil (event-dispatch map compute))
     (compute (output force (compute-force map))
              nil)))
```

```
#+level0
(defmodule runaway
  :inputs (force)
  :outputs (command)
  :states
    ((nil (event-dispatch force decide))
     (decide (conditional-dispatch (significant-force-p force 3.0)
                                   runaway
                                   nil))
     (runaway (output command (follow-force force))
              nil)))

#+level1
(defmodule wander
  :inputs ()
  :outputs (heading)
  :states
    ((nil (event-dispatch (delay 10.0) generate))
     (generate (output heading (make-random-heading 200.0))
               nil)))

#+level1
(defmodule avoid
  :inputs (force heading)
  :outputs (command)
  :instance-vars (resultforce)
  :states
    ((nil (event-dispatch (and force heading) plan))
     (plan (setf resultforce (select-direction force heading))
           go)
     (go (conditional-dispatch (significant-force-p resultforce 1.0)
                               start
                               nil))
     (start (output command (follow-force resultforce))
            nil)))

#+level2
(defmodule grabber
  :inputs (goal motor-status)
  :outputs (grab outgoal)
  :states
    ((nil (event-dispatch goal reset))
     (reset (output grab hi)
            wait)
     (wait (event-dispatch (delay 0.25) test))
     (test (conditional-dispatch (lo? motor-status) gotit reset))
     (gotit (output outgoal goal)
            nil)))
```

```
#+level2
(defmodule monitor
  :inputs (motor-status)
  :outputs (travel)
  :states
    ((nil (event-dispatch motor-status check))
     (check (conditional-dispatch (lo? motor-status) out nil))
     (out (output travel (send *robot* :travel?))
          nil)))


#+level2
(defmodule integrate
  :inputs (travel)
  :outputs (integral)
  :instance-vars (accum-motion)
  :states
    ((nil (event-dispatch travel initialize))
     (initialize (setf accum-motion (make-zero-relative-motion))
                 add)
     (add (setf accum-motion (increment-motion accum-motion travel))
          add2)
     (add2 (output integral accum-motion)
           accumulate)
     (accumulate (event-dispatch travel add))))


#+level2
(defmodule pathplan
  :inputs (integral goal)
  :outputs (begin heading turn)
  :states
    ((nil (event-dispatch goal init))
     (init (output begin hi)
           firstmove)
     (firstmove (output heading (decide-local-goal goal '()))
                wait)
     (move (output heading (decide-local-goal goal integral))
           wait)
     (wait (event-dispatch integral test))
     (test (conditional-dispatch (close-enuf-p goal integral)
                                 turn
                                 move))
     (turn (output turn goal)
           nil)))
```

```
#+level2
(defmodule straighten
  :inputs (goal integral)
  :outputs (nostop command done)
  :states
    ((nil (event-dispatch goal decide))
     (decide (conditional-dispatch (close-enuf-orientation-p goal
                                                             integral)
                                   over
                                   send))
     (send (output nostop hi)
           sendturn)
     (sendturn (output command (select-turn goal integral))
               wait)
     (wait (event-dispatch integral over))
     (over (output done hi)
           nil)))
;;; level 0 wiring diagram
#+level0 (defwire (sonar map) (collide map) (feelforce map))
#+level0 (defwire (collide halt) (motor halt))
#+level0 (defwire (feelforce force) (runaway force)
                                    #+level1 (avoid force))
#+level0 (defwire (runaway command) (motor command))
;;; level 1 wiring diagram
#+level1 (defwire (wander heading) (avoid heading))
#+level1 (defwire (avoid command) ((suppress (motor command) 1.5)))
;;; level 2 wiring diagram
#+level2 (defwire (motor motor-status) (monitor motor-status)
                                       (grabber motor-status))
#+level2 (defwire (monitor travel) (integrate travel))
#+level2 (defwire (grabber outgoal) (pathplan goal))
#+level2 (defwire (grabber grab) ((inhibit (wander heading) 2.0))
                                 ((inhibit (avoid command) 0.25))
                                 ((inhibit (runaway command) 2.0))
                                 ((suppress (motor halt) 0.5)))
#+level2 (defwire (pathplan begin) ((reset integrate)))
#+level2 (defwire (pathplan heading) ((suppress (avoid heading) 15.0)))
#+level2 (defwire (integrate integral) (pathplan integral)
                                       (straighten integral))
#+level2 (defwire (pathplan turn) (straighten goal))
#+level2 (defwire (straighten nostop) ((inhibit (collide halt) 2.5)))
#+level2 (defwire (straighten command)
                  ((suppress (motor command) 1.5)))
```