Reactivity in a Logic-Based Robot Programming Framework*

Yves Lespérance Kenneth Tam Michael Jenkin

Dept. of Computer Science, York University,
4700 Keele St., Toronto, ON Canada, M3J 1P3
{lesperan,kenneth,jenkin}@cs.yorku.ca

Abstract. A robot must often react to events in its environment and exceptional conditions by suspending or abandoning its current plan and selecting a new plan that is an appropriate response to the event. This paper describes how high-level controllers for robots that are reactive in this sense can conveniently be implemented in ConGolog, a new logic-based agent/robot programming language. Reactivity is achieved by exploiting ConGolog's prioritized concurrent processes and interrupts facilities. The language also provides nondeterministic constructs that support a form of planning. Program execution relies on a declarative domain theory to model the state of the robot and its environment. The approach is illustrated with a mail delivery application.

1 Introduction

Reactivity is usually understood as having mainly to do with strict constraints on reaction time. As such, much work on the design of reactive agents has involved non-deliberative approaches where behavior is hardwired [3] or produced from compiled universal plans [18, 17]. However, there is more to reacting to environmental events or exceptional conditions than reaction time. While some events/conditions can be handled at a low level, e.g., a robot going down a hallway can avoid collision with an oncoming person by slowing down and making local adjustments in its trajectory, others require changes in high-level plans. For example, an obstacle blocking the path of a robot attempting a delivery may mean that the delivery must be rescheduled. Here as in many other cases, the issue is not real-time response. What is required is *reconsideration of the robot's plans in relation to its goals and the changed environmental conditions.* Current plans may need to be suspended or terminated and new plans devised to deal with the exceptional event or condition.

* The research described received financial support from Communications and Information Technology Ontario (and its earlier incarnation ITRC) and the Natural Science and Engineering Research Council of Canada. Hector Levesque came up with the idea of handling sensing through exogenous events; he also helped with the iterative deepening route planning procedure. Ho Kong Ng did much of the work involved in porting the controller to the Nomad Super Scout and testing the resulting system. We thank them as well as Ray Reiter, Jeff Lloyd, Mikhail Soutchanski, Giuseppe De Giacomo, and Daniele Nardi for helpful discussions related to this work. Many of our papers are available at http://www.cs.yorku.ca/~lesperan/.

To provide the range of responses required by environmental events and exceptional conditions, i.e. reactivity in the wide sense, the best framework seems to be a hierarchical architecture. Then, urgent conditions can be handled in real-time by a low-level control module, while conditions requiring replanning are handled by a high-level control module that models the environment and task, and manages the generation, selection, and scheduling of plans.

Synthesizing plans at run-time provides great flexibility, but it is often computationally infeasible in complex domains, especially when the agent does not have complete knowledge and there are exogenous events (i.e. actions by other agents or natural events). In [10], it was argued that *high-level program execution* was a more practical alternative. The idea, roughly, is that instead of searching for a sequence of actions that takes the robot from an initial state to some goal state, the task is to find a sequence of actions that constitutes a legal execution of some high-level program. By high-level program, we mean one whose primitive instructions are domain-dependent actions of the robot, whose tests involve domain-dependent predicates that are affected by the actions, and whose code may contain nondeterministic choice points where lookahead is necessary to make a choice that leads to successful termination. As in planning, to find a sequence that constitutes a legal execution of a high-level program, one must reason about the preconditions and effects of the actions within the program. However, if the program happens to be almost deterministic, very little searching is required; as more and more nondeterminism is included, the search task begins to resemble traditional planning. Thus, in formulating a high-level program, the user gets to control the search effort required.

In [10], *Golog* was proposed as a suitable language for expressing high-level programs for robots and autonomous agents. Golog was used to design a high-level robot control module for a mail delivery application [20]. This module was interfaced to systems providing path planning and low-level motion control, and successfully tested on several different robot platforms, including a Nomad 200, a RWI B21, and a RWI B12.

A limitation of Golog for this kind of applications is that it provides limited support for writing reactive programs. In [4], *GonGolog*, an extension of Golog that provides concurrent processes with possibly different priorities as well as interrupts was introduced. In this paper, we try to show that ConGolog is an effective tool for the design of high-level reactive control modules for robotics applications. We provide an example of such a module for a mail delivery application. The approach proposed should also be applicable in contexts other than robotics, where an agent is embedded in some environment and needs to react to failures and changes in its environment.

2 ConGolog

As mentioned, our high-level programs contain primitive actions and tests of predicates that are domain-dependent. Moreover, an interpreter for such programs must reason about the preconditions and effects of the actions in the program to find a legal terminating execution. We specify the required domain theories in the situation calculus [11], a language of predicate logic for representing dynamically changing worlds. In this language, a possible world history, which is simply a sequence of actions, is represented by a first order term called a *situation*. The constant S_0 is used to denote the initial situation —

that in which no actions (of interest) have yet occurred. There is a distinguished binary function symbol do and the term $do(\alpha, s)$ denotes the situation resulting from action α being performed in situation s. Relations whose truth values vary from situation to situation, called *predicate fluents*, are denoted by predicate symbols taking a situation term as the last argument. For example, $Holding(o, s)$ might mean that the robot is holding object o in situation s. Similarly, functions whose value varies with the situation, *functional fluents*, are represented by function symbols that take a situation argument. The special predicate $Poss(\alpha, s)$ is used to represent the fact that primitive action α is executable in situation s. A domain of application will be specified by theory that includes the following types of axioms:

- Axioms describing the initial situation, S_0.
- Action precondition axioms, one for each primitive action α, which characterizes $Poss(\alpha, s)$.
- Successor state axioms, one for each fluent F, which characterize the conditions under which $F(\boldsymbol{x}, do(a, s))$ holds in terms of what holds in situation s; these axioms may be compiled from effects axioms, but provide a solution to the frame problem [14].
- Unique names axioms for the primitive actions.
- Some foundational, domain independent axioms.

Thus, the declarative part of a ConGolog program implementing a high-level controller for a robot will be such a theory.

A ConGolog program also includes a procedural part which specifies the behavior of the robot. This is specified using the following constructs:

α,	primitive action
$\phi?$,	wait for a condition[1]
$(\sigma_1; \sigma_2)$,	sequence
$(\sigma_1 \mid \sigma_2)$,	nondeterministic choice between actions
$\pi\, \boldsymbol{x}\, [\sigma]$,	nondeterministic choice of arguments
σ^*,	nondeterministic iteration
if ϕ **then** σ_1 **else** σ_2 **endIf**,	conditional
while ϕ **do** σ **endWhile**,	loop
$(\sigma_1 \parallel \sigma_2)$,	concurrent execution
$(\sigma_1 \rangle\!\rangle\, \sigma_2)$,	concurrency with different priorities
$\sigma^{\parallel}$,	concurrent iteration
$<\boldsymbol{x} : \phi \rightarrow \sigma>$,	interrupt
proc $\beta(\boldsymbol{x})\ \sigma$ **endProc**,	procedure definition
$\beta(\boldsymbol{t})$,	procedure call
noOp	do nothing

The nondeterministic constructs include $(\sigma_1 \mid \sigma_2)$, which nondeterministically choses between programs σ_1 and σ_2, $\pi \boldsymbol{x}[\sigma]$, which nondeterministically picks a binding for the variables $\boldsymbol{x}$ and performs the program σ for this binding of $\boldsymbol{x}$, and σ^*, which means

[1] Here, ϕ stands for a situation calculus formula with all situation arguments suppressed; $\phi(s)$ will denote the formula obtained by restoring situation variable s to all fluents appearing in ϕ.

performing σ zero or more times. Concurrent processes are modeled as interleavings of the primitive actions involved. A process may become blocked when it reaches a primitive action whose preconditions are false or a wait action ϕ? whose condition ϕ is false. Then, execution of the program may continue provided another process executes next. In $(\sigma_1 \rangle\rangle \sigma_2)$, σ_1 has higher priority than σ_2, and σ_2 may only execute when σ_1 is done or blocked. $\sigma^{\|}$ is like nondeterministic iteration σ^*, but the instances of σ are executed concurrently rather than in sequence. Finally, an interrupt $< \boldsymbol{x} : \phi \rightarrow \sigma >$ has variables $\boldsymbol{x}$, a trigger condition ϕ, and a body σ. If the interrupt gets control from higher priority processes and the condition ϕ is true for some binding of the variables, the interrupt triggers and the body is executed with the variables taking these values. Once the body completes execution, the interrupt may trigger again. With interrupts, it is easy to write programs that are reactive in that they will suspend whatever task they are doing to handle given conditions as they arise. A more detailed description of ConGolog and a formal semantics appear in [4]. We give an example ConGolog program in section 4.

A prototype ConGolog interpreter has been implemented in Prolog. This implementation requires that the axioms in the program's domain theory be expressible as Prolog clauses; note that this is a limitation of this particular implementation, not the framework.

In applications areas such as robotics, we want to use ConGolog to program an embedded system. The system must sense conditions in its environment and update its theory appropriately *as it is executing* the ConGolog control program.[2] This requires adapting the high-level program execution model presented earlier: the interpreter cannot simply search all the way to a final situation of the program. An adapted model involving incremental high-level program execution is developed in [5]. However in this paper, we sidestep these issues by making two simplifying assumptions:

1. that the interpreter can make an arbitrary choice among the primitive actions that are allowed by a nondeterministic program at each step and immediately commit and execute this action, and
2. that there is a set of exogenous events detectable by the system's sensors (e.g. a mail pick up request is received or the robot has arrived at the current destination) and that the environment is continuously monitored for these; whenever such an exogenous event is detected to have occurred, it is immediately inserted in the execution.

We can get away with this because our application program performs essentially no search and the exogenous events involved are easy to detect. In the incremental execution model of [5], sensing actions can be included at specific points in the program, and while nondeterminism is interpreted as arbitrary choice as per assumption (1), a new "search block" construct is also added to the language, so that within such a block the interpreter must search to ensure that the primitive action selected is on a path to a terminating configuration for the block.

---

[2] Here, the environment is anything outside the ConGolog control module about which information must be maintained; so the sensing might only involve reading messages from another module through a communication socket.

3 Interfacing the High-Level Control Module

As mentioned earlier, we use a hierarchical architecture to provide both real-time response as well as high-level plan reconsideration when appropriate. At the lowest level, we have a reactive control system that performs time-critical tasks such as collision avoidance and straight line path execution. In a middle layer, we have a set of components that support navigation through path planning, map building and/or maintenance, keeping track of the robot's position, etc. and support path following by interacting with the low-level control module. On top of this, there is the ConGolog-based control module that supports high-level plan execution to accomplish the robot's tasks; this level treats navigation somewhat like a black box.

In this section, we describe how the ConGolog-based high-level control module is interfaced to rest of the architecture. The high-level control module needs to run asynchronously with the rest of the architecture so that other tasks can be attended to while the robot is navigating towards a destination. It also needs to interact with the navigation module to get tasks accomplished. To support this, we need to give the high-level control module a *model of the navigation module*. We have defined a simple version of such a model. With respect to navigation, the robot is viewed by the high-level control module as always being in one of the following set of states:

$$RS = \{Idle, Moving, Reached, Stuck, Frozen\}.$$

The current robot state is represented by the functional fluent $robotState(s)$. The robot's default state is $Idle$; when in this state, the robot is not moving towards a destination, but collision avoidance is turned on and the robot may move locally to avoid oncoming bodies. With the robot in $Idle$ state, the high-level control module may execute the primitive action $startGoTo(place)$; this changes the robot's state to $Moving$ and causes the navigation module to attempt to move the robot to $place$. If and when the robot reaches the destination, the navigation module generates the exogenous event $reachDest$, which changes the robot's state to $Reached$. If on the other hand the navigation module encounters obstacles it cannot get around and finds the destination unreachable, then it generates the exogenous event $getStuck$, which changes the robot's state to $Stuck$. In any state, the high-level control module may execute the primitive action $resetRobot$, which aborts any navigation that may be under way and returns the robot to $Idle$ state. Finally, there is the $Frozen$ state where collision avoidance is disabled and the robot will not move even if something approaches it; this is useful when the robot is picking up or dropping off things; humans may reach into the robot's carrying bins without it moving away. All other actions leave the robot's state unchanged. This is specified in the following successor state axiom for the $robotState$ fluent:

$$\begin{aligned}
&robotState(do(a,s)) = i \equiv \\
&\quad \exists p\, a = startGoTo(p) \wedge i = Moving \vee \\
&\quad a = reachDest \wedge i = Reached \vee \\
&\quad a = getStuck \wedge i = Stuck \vee \\
&\quad a = resetRobot \wedge i = Idle \vee \\
&\quad a = frezeRobot \wedge i = Frozen \vee \\
&\quad i = robotState(s) \wedge \forall p\, a \neq startGoTo(p) \wedge a \neq reachDest \wedge \\
&\qquad a \neq getStuck \wedge a \neq resetRobot \wedge a \neq frezeRobot
\end{aligned}$$

We also have precondition axioms that specify when these primitive actions and exogenous events are possible. For example, the following says that the action of directing the robot to start moving toward a destination p is possible whenever the robot is in $Idle$ state:

$$Poss(startGoTo(p), s) \equiv robotState(s) = Idle$$

We omit the other precondition axioms as they are obvious from the model description.

We also use two additional functional fluents: $robotDestination(s)$ refers to the last destination the robot was set in motion towards, and $robotPlace(s)$ refers to the current location of robot as determined from the model. Their successor state axioms are:

$$\begin{aligned}
&robotDestination(do(a, s)) = p \equiv \\
&\quad a = startGoTo(p) \lor \\
&\quad p = robotDestination(s) \land \forall p\, a \neq startGoTo(p)
\end{aligned}$$

$$\begin{aligned}
&robotPlace(do(a, s)) = p \equiv \\
&\quad \exists p'\, a = startGoTo(p') \land p = Unknown \lor \\
&\quad a = reachDest \land p = robotDestination(s) \lor \\
&\quad p = robotPlace(s) \land \\
&\qquad \forall p\, a \neq startGoTo(p) \land a \neq reachDest
\end{aligned}$$

Note that in this model, we treat navigation to a location as an "activity" in the sense of Gat [8], i.e. the primitive actions are essentially instantaneous changes that do not constitute the activity, but only initiate or terminate it. In contrast to [8], we do specify how the world model is to be updated. In extending our model, we should be able to use the same approach to include other activities, e.g. perceptual search routines.

4 A Mail Delivery Example

To test our approach, we have implemented a simple mail delivery application. The high-level control module for the application must react to two kinds of exogenous events:

- new shipment orders, which are represented by the event $orderShipment(sender, recipient, priority)$, and
- signals from the navigation module, namely the $reachDest$ event announcing that the destination has been reached and the $getStuck$ event announcing that the robot has failed to reach its destination.

The first kind is typical of the communication interactions a robot may have with its environment, while the second kind is typical of the control interactions a task-level module may have with the rest of the robot's architecture. To require more reactivity from the robot, we assume that shipment orders come with different priority levels and that the system must interrupt service of a lower priority order when a higher priority one comes in. Also, we want the robot to make a certain number of attempts to get to a customer's mailbox as some of the obstacles it runs into may be temporary. This is handled by assigning a certain amount of credit to customers initially and reducing their credit when an attempt to go to their mailbox fails. When customers run out of credit, they are suspended and shipments sent to them are returned to the sender when possible.

In addition to the navigation primitive actions and exogenous events already described, the application uses the following primitive actions:

$ackOrder(n)$	acknowledge reception of servable order
$declineOrder(n)$	decline an unservable order
$pickUpShipment(n)$	pick up shipment n
$dropOffShipment(n)$	drop off shipment n
$cancelOrder(n)$	cancel an unservable order
$reduceCredit(c)$	reduce customer c's credit
$notifyStopServing(c)$	notify unreachable customer

Note that shipment orders are identified by a number n that is assigned from a counter when the $orderShipment$ event occurs. We have precondition axioms for these primitive actions, for example:

$$
\begin{array}{l}
Poss(pickUpShipment(n), s) \equiv \\
\quad orderState(n, s) = ToPickUp \wedge \\
\quad robotPlace(s) = mailbox(sender(n, s))
\end{array}
$$

The primitive fluents for the application are:

$orderState(n, s) = i$	order n is in state i
$sender(n, s) = c$	sender of order n is c
$recipient(n, s) = c$	recipient of order n is c
$orderPrio(n, s) = p$	priority of order n is p
$orderCtr(s) = n$	counter for orders arriving
$credit(c, s) = k$	customer c has credit k
$Suspended(c, s)$	service to customer c is suspended

We have successor state axioms for these fluents. For example, the state of an order starts out as $NonExistent$, then changes to $JustIn$ when the $orderShipment$ event occurs, etc.; the following successor state axiom specifies this:

$$
\begin{array}{l}
orderState(n, do(a, s)) = i \equiv \\
\quad \exists c, r, p\, a = orderShipment(c, r, p) \wedge orderCtr = n \wedge i = JustIn \vee \\
\quad a = ackOrder(n) \wedge i = ToPickUp \vee \\
\quad a = pickUpShipment(n) \wedge i = OnBoard \vee \\
\quad a = dropOffShipment(n) \wedge \\
\quad\quad robotPlace(s) = mailbox(recipient(n, s)) \wedge i = Delivered \vee \\
\quad a = dropOffShipment(n) \wedge \\
\quad\quad robotPlace(s) = mailbox(sender(n, s)) \wedge i = Returned \vee \\
\quad a = dropOffShipment(n) \wedge \\
\quad\quad robotPlace(s) = CentralOffice \wedge i = AtCentralOffice \vee \\
\quad a = cancelOrder(n) \wedge i = Cancelled \vee \\
\quad a = declineOrder(n) \wedge i = Declined \vee \\
\quad i = orderState(n, s) \wedge \neg(\exists c, r, p\, a = orderShipment(c, r, p) \wedge orderCtr = n) \\
\quad\quad \wedge\, a \neq ackOrder(n) \wedge a \neq pickUpShipment(n) \wedge \\
\quad\quad a \neq dropOffShipment(n) \wedge a \neq cancelOrder(n) \wedge a \neq declineOrder(n)
\end{array}
$$

We omit the rest of the successor state and action precondition axioms for space reasons.

The initial state of the domain might be specified by the following axioms:

$$
\begin{array}{l}
Customer(Yves) \quad Customer(Ken) \\
Customer(Hector) \quad Customer(Michael) \\
Customer(c) \supset credit(c, S_0) = 3 \\
orderCtr(S_0) = 0 \\
orderState(n, S_0) = NonExistent \\
robotState(S_0) = Idle \\
robotPlace(S_0) = CentralOffice
\end{array}
$$

Let us now specify the behavior of our robot using a ConGolog program. Exogenous events are handled using prioritized interrupts. The main control procedure concurrently executes four interrupts at different priorities:

proc $mainControl$
  $<\ n : orderState(n) = JustIn \rightarrow handleNewOrder(n)\ >$
  $\rangle\!\rangle$
  $<\ n : (orderState(n) = ToPickUp \wedge Suspended(sender(n)))$
    $\rightarrow cancelOrder(n)\ >$
  $\|$
  $<\ n : (orderState(n) = ToPickUp \vee orderState(n) = OnBoard$
      $\vee\ robotPlace \neq CentralOffice)$
    $\rightarrow robotMotionControl\ >$
  $\rangle\!\rangle$
  $<\ robotState = Moving \rightarrow$ **noOp** $>$
endProc

The top priority interrupt takes care of acknowledging or declining new shipment orders. This ensures that customers get fast feedback when they make an order. At the next level of priority, we have two other interrupts, one that takes care of cancelling orders whose senders have been suspended service, and another that controls the robot's motion. At the lowest priority level, we have an interrupt with an empty body that prevents the program from terminating when the robot is in motion and all other threads are blocked.

The top priority interrupt deals with a new shipment order n by executing the following procedure:

proc $handleNewOrder(n)$
  if $Suspended(sender(n)) \vee Suspended(recipient(n))$ **then**
    $declineOrder(n)$
  else
    $ackOrder(n);$
    if $robotState = Moving \wedge orderPrio(n) > curOrderPrio$ **then**
      $resetRobot$ % abort current service
    endIf
  endIf
endProc

This sends a rejection notice to customers making an order whose sender or recipient is suspended; otherwise an acknowledgement is sent. In addition, when the new shipment order has higher priority than the order currently being served, the robot's motion is aborted, causing a reevaluation of which order to serve ($curOrderPrio$ is a defined fluent whose definition appears below).

The second interrupt in $mainControl$ handles the cancellation of orders when the sender gets suspended; its body executes the primitive action $cancelOrder(n)$.

The third interrupt in $mainControl$ handles the robot's navigation, pick ups, and deliveries. When the interrupt's condition is satisfied, the following procedure is called:

proc $robotMotionControl$
  if $\exists c\, CustToServe(c)$ **then** $tryServeCustomer$
  else $tryToWrapUp$;
  endIf
endProc

This tries to serve a customer when there is one to be served and tries to return to the central office and wrap up otherwise. $CustToServe(c, s)$ is a defined fluent:

$$\begin{array}{l} CustToServe(c,s) \stackrel{\text{def}}{=} \exists n[\\ \quad (orderState(n,s) = ToPickUp \wedge sender(n,s) = c \\ \qquad \wedge \neg Suspended(recipient(n,s),s)) \vee \\ \quad (orderState(n,s) = OnBoard \wedge (recipient(n,s) = c \\ \quad\; \vee\, sender(n,s) = c \wedge Suspended(recipient(n,s),s)))] \\ \;\wedge \neg Suspended(c,s) \end{array}$$

To try to serve a customer, we execute the following:

proc $tryServeCustomer$
  $\pi\, c\, [BestCustToServe(c)?;$
    $startGoTo(mailbox(c));$
    $(robotState \neq Moving)?;$
    if $robotState = Reached$ **then**
      $freezeRobot;$
      $dropOffShipmentsTo(c);$
      $pickUpShipmentsFrom(c);$
      $resetRobot$
    else if $robotState = Stuck$ **then**
        $resetRobot;$ % abandon attempt
        $handleServiceFailure(c)$
    % else when service aborted nothing more to do
    endIf]
endProc

This first picks one of the best customers to serve, directs the robot to start navigating towards the customer's mailbox, and waits until the robot halts. If the robot reaches the customer's mailbox, then shipments for the customer are dropped off and shipments from him/her are picked up. If on the other hand, the robot halts because it got stuck, the $handleServiceFailure$ procedure is executed. Finally, if the robot halts because a higher priority order came in and the top priority interrupt executed a $resetRobot$, then there is nothing more to be done. $BestCustToServe(c, s)$ is a defined fluent that captures all of the robot's order scheduling criteria:

$$
\begin{array}{l}
BestCustToServe(c, s) \stackrel{\text{def}}{=} CustToServe(c, s) \wedge \\
\quad custPriority(c, s) = maxCustPriority(s) \wedge \\
\quad credit(c, s) = maxCreditFor(maxCustPriority(s), s)
\end{array}
$$

$$
\begin{array}{l}
custPriority(c, s) = p \stackrel{\text{def}}{=} \\
\quad \exists n\, OrderForCustAtPrio(n, c, p, s) \wedge \\
\quad \forall n', p'(OrderForCustAtPrio(n', c, p', s) \supset p' \leq p)
\end{array}
$$

$$
\begin{array}{l}
OrderForCustAtPrio(n, c, p, s) \stackrel{\text{def}}{=} \neg Suspended(c, s) \wedge \\
\quad [orderState(n, s) = ToPickUp \wedge sender(n, s) = c \wedge \\
\qquad orderPrio(n, s) = p \vee \\
\quad\; orderState(n, s) = OnBoard \wedge orderPrio(n, s) = p \wedge \\
\qquad (recipient(n, s) = c \vee \\
\qquad\; sender(n, s) = c \wedge Suspended(recipient(n, s), s))]
\end{array}
$$

$$
\begin{array}{l}
maxCustPriority(s) = p \stackrel{\text{def}}{=} \exists c\, custPriority(c, s) = p \\
\qquad\qquad \wedge\, \forall c'\, custPriority(c', s) \leq p
\end{array}
$$

$$
\begin{array}{l}
maxCreditFor(p, s) = k \stackrel{\text{def}}{=} \\
\quad \exists c[custPriority(c, s) = p \wedge credit(c, s) = k \wedge \\
\qquad \forall c'(custPriority(c', s) = p \supset credit(c', s) \leq k)]
\end{array}
$$

This essentially says that the best customers to serve are those that have the highest credit among those having the highest priority orders. We can now also define the priority of the order currently being served as follows:

$$
\begin{array}{l}
curOrderPrio(s) = p \stackrel{\text{def}}{=} \\
\quad \forall c[robotState(s) = Moving \wedge \\
\qquad robotDestination(s) = mailbox(c) \\
\qquad \supset p = custPriority(c, s)] \wedge \\
\quad [\neg(robotState(s) = Moving \wedge \\
\quad\; \exists c\, robotDestination(s) = mailbox(c)) \supset p = -1]
\end{array}
$$

The $handleServiceFailure$ procedure goes as follows:

proc $handleServiceFailure(c)$
  $reduceCredit(c);$
  if $credit(c) = 0$ **then**
    $notifyStopServing(p);$
  endIf;
endProc

When the robot gets stuck on the way to customer c's mailbox, it first reduces c's credit, and then checks whether it has reached zero and c has just become $Suspended$; if so, c is notified that he/she will no longer be served.

The $tryToWrapUp$ procedure is similar to $tryServeCustomer$, except that it attempts to drive the robot to the $CentralOffice$ and then drop off undeliverable shipments there.

Procedure $dropOffShipmentsTo(c)$ delivers to customer c all shipments on board such that c is the shipment's recipient or c is the shipment's sender and the recipient has been suspended:

proc $dropOffShipmentsTo(c)$
  while $\exists n\,(orderState(n) = OnBoard \wedge$
    $(recipient(n) = c \vee$
     $sender(n) = c \wedge Suspended(recipient(n))))$ **do**
   $\pi\, n\,[(orderState(n) = OnBoard \wedge$
    $(recipient(n) = c \vee$
      $sender(n) = c \wedge Suspended(recipient(n))))?;$
     $dropOffShipment(n)]$
  endWhile
endProc

Procedure $pickUpShipmentsFrom(c)$ simply picks up all outgoing shipments from customer c's mailbox.

Note that by handling the cancellation of pick ups in a separate thread from that dealing with navigation and order serving, we allow the robot to be productive while it is in motion and waiting to reach its destination. This makes a better use of resources.

To run the system, we execute $mainControl$ after placing a few initial orders:

$$
\begin{array}{l}
orderShipment(Yves, Ken, 0) \parallel \\
orderShipment(Ken, Hector, 1) \rangle\!\rangle \\
mainControl
\end{array}
$$

5 Implementation

The high-level control module for the mail delivery application has been ported to an RWI B12 and a Nomad Super Scout mobile robot and tested on some sample delivery tasks (see figure 1). The other software components for this were based on a system developed during an earlier project concerned with building an experimental vehicle to conduct survey/inspection tasks in an industrial environment [12]. The system supports point to point navigation in a previously mapped environment and can use prepositioned visual landmarks to correct odometry errors. It relies on sonar sensors to sense unmodeled obstacles.

The system's architecture conforms to the general scheme described earlier. It provides two levels of control. An onboard low-level controller [16] performs all time-critical tasks such as collision avoidance and straight line path execution. The low-level controller assumes that the robot is always in motion and communicates with an offboard

Fig. 1. Our robot facing an obstacle.

global path planner and user interface module known as the *Navigator*. The Navigator takes as inputs a metric/topological map of the environment in which the robot is located and the coordinates (as defined in the map) of the two end points, i.e., the source and the destination of a path. By making use of some predefined path-finding algorithms such as breadth-first search or A^* the Navigator identifies a feasible path between the source and the destination. To follow the path, the Navigator decomposes it into segments (a segment is a straight line between two adjacent way-points) and then forwards the segments to the low-level controller for execution. The Navigator supervises the low-level controller and identifies failures in the low-level controller's ability to execute a path segment.

The ConGolog-based high-level control module interacts with the rest of the architecture by communicating with the Navigator through a socket interface. The high-level controller, Navigator, and low-level controller all run asynchronously. The primitive actions in the ConGolog interface model are implemented using operations provided by the Navigator (currently, the mail pickup and drop off actions are only simulated). For example, the ConGolog primitive action $startGoTo(p)$ is implemented as `[planPath(coordinatesOf(`p`)); followPath]`, where `planPath` and `follow_path` are operations supplied by the Navigator.

Our test scenarios confirmed the system's ability to deal with navigation failures and to interrupt the current task when an urgent shipment order is made. For more details on the implemented system, see [21].

6 Discussion

In this work, we have shown how ConGolog can be used to implement high-level robot controllers that can cope with dynamic and unpredictable environments — controllers that are reactive and support high-level plan reconsideration in response to exogenous events and exceptional conditions. Our work demonstrates that a logic-based approach can be used to build effective systems. In the ConGolog framework, application domains and their dynamics are specified declaratively and the axioms are used by the interpreter to automatically update its world model. This is less error-prone than having the user program his own ad-hoc world representation and update procedures. ConGolog controllers have clear formal specifications and are easier to extend and adapt to different environments or tasks. Existing robot control frameworks with similar objectives include RAP [7] and ATLANTIS [8]. In ConGolog, the control framework is less constrained and how the world changes is specified declaratively in a very expressive language.

The main limitation of the work accomplished so far is the lack of search/planning in the current high-level control program. ConGolog is designed to support run-time planning through search over a nondeterministic program. This is the main reason for using a declarative representation of domain dynamics. So the next step in our work will be to extend the controller developed so far to perform run-time planning for route optimization and dealing with failures (e.g., in navigation).

It is not difficult to write search/planning code in ConGolog. For example, here's how one might implement an iterative deepening search to find the shortest route through customers needing service — one starts the search by invoking $planRoute(0)$:

proc $planRoute(n)$
  $serveAllCust(n) \mid planRoute(n+1)$
endProc

proc $serveAllCust(n)$
  $\neg\exists c CustToServe(c)? \mid$
  $\pi\, c, d[(CustToServe(c) \wedge d = distanceTo(c) \wedge d \leq n)?;$
    $goServe(c); serveAllCust(n-d)]$
endProc

But adding such planning to a reactive control program raises a lot of complex issues. The planning/search must be interleaved with action execution and sensing the environment — as mentioned earlier, the standard ConGolog execution model does not support this. The generated plan must also be reevaluated when conditions change; it may no longer be executable, or achieve the goal, or be appropriate in the new conditions (e.g. when a more urgent order has just arrived). As well, when generating a plan, we may want to ignore the possibility that some actions may fail to achieve their objectives (e.g. navigating to a customer), and just deal with such failures when they occur.

Recently, De Giacomo and Levesque [5] have developed a new execution model for ConGolog that supports incremental high-level program execution in the presence of sensing. This should allow us to incorporate controlled search/planning in our programs while retaining a clean semantics. We are examining ways of dealing with the other issues mentioned. A Golog-based approach to execution monitoring and plan repair has been proposed in [6]. The use of Golog for planning is discussed in [15].

Another limitation of our work so far is that the system developed is rather small. We need to experiment with more complex tasks to see whether our approach and the use of prioritized interrupts to provide reactivity scales up. As well, interrupts support the suspension of the current plan but not its termination; we may add a conventional exception throwing/catching mechanism that terminates the current plan.

Another area under investigation is information acquisition tasks. [21] tackles an application where packages must be delivered to the recipient "in person" and where the robot must use smart search strategies to locate the recipient, for example, asking someone whether he has seen the recipient.

The high-level program execution model of robot/agent control that underlies our approach is related to work on resource-bounded deliberative architectures [2, 13] and agent programming languages [19, 9, 22]. One difference is that in our approach, plan selection is coded in the program; there is no goal-directed invocation of plans. This makes for a less declarative and perhaps more complex specification, but eliminates some overhead. On the other hand, the robot's world is modeled using a domain action theory and the world model is updated automatically using the successor state axioms; there is no need to perform asserts and retracts. Moreover, the evaluation of a test may involve arbitrary amounts of inference, although following logic programming philosophy, we take the programmer to be responsible for its efficiency/termination. Perhaps a more central difference is that our robots/agents can be understood as executing programs, albeit in a rather smart way — they have a simple operational semantics. Modeling the operation of an agent implemented using a resource-bounded deliberative architecture requires a much more complex account. In [1], Baral and Son extend ConGolog to handle plans specified as hierarchical task networks; this somewhat bridges the gap between ConGolog and "rule-based" agent languages.

References

1. Chitta Baral and Tran Cao Son. Extending ConGolog to allow partial ordering. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin, 2000. In this volume.
2. M.E. Bratman, D.J. Israel, and M.E. Pollack. Plans and ressource-bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.
3. R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2(1):14–23, 1986.
4. Giuseppe De Giacomo, Yves Lespérance, and Hector J. Levesque. Reasoning about concurrent execution, prioritized interrupts, and exogenous actions in the situation calculus. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 1221–1226, Nagoya, Japan, August 1997.
5. Giuseppe De Giacomo and Hector J. Levesque. An incremental interpreter for high-level programs with sensing. In *Cognitive Robotics — Papers from the 1998 AAAI Fall Symposium*, pages 28–34, Orlando, FL, October 1998. AAAI Press. AAAI Tech. Report FS-98-02.
6. Giuseppe De Giacomo, Raymond Reiter, and Mikhail E. Soutchanski. Execution monitoring of high-level robot programs. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixth International Conference (KR'98)*, pages 453–464, Trento, Italy, June 1998.

7. R. James Firby. An investigation into reactive planning in complex domains. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 202–206, Seattle, WA, July 1987.
8. Erann Gat. Integrating planning and reacting in a heterogenous asynchronous architecture for controlling real-world mobile robots. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 809–815, San Jose, CA, July 1992.
9. K. V. Hindriks, F. S. de Boer, W. van der Hoek, and J.-J. Ch. Meyer. A formal semantics for an abstract agent programming language. In M.P. Singh, A. Rao, and M.J. Wooldridridge, editors, *Intelligent Agents IV — Proceedings of ATAL'97*, pages 215–229. Springer-Verlag, June 1998.
10. Hector J. Levesque, Raymond Reiter, Yves Lespérance, Fangzhen Lin, and Richard B. Scherl. GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31(59–84), 1997.
11. John McCarthy and Patrick Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, Edinburgh, UK, 1979.
12. B. Nickerson, P. Jasiobedzki, D. Wilkes, M. Jenkin, E. Milios, J. Tsotsos, A. Jepson, and O.N. Bains. The ARK project: Autonomous mobile robots for known industrial environments. *Robotics and Autonomous Systems*, 25:83–104, 1998.
13. A.S. Rao and M.P. Georgeff. An abstract architecture for rational agents. In Bernhard Nebel, Charles Rich, and William Swartout, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*, pages 439–449, Cambridge, MA, 1992. Morgan Kaufmann Publishing.
14. Raymond Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In Vladimir Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, San Diego, CA, 1991.
15. Raymond Reiter. Knowledge in action: Logical foundations for describing and implementing dynamical systems. Draft Monograph, available at `http://www.cs.toronto.edu/~cogrobo`, 1999.
16. Matt Robinson and Michael Jenkin. Reactive low level control of the ARK. In *Proceedings, Vision Interface '94*, pages 41–47, Banff, AB, May 1994.
17. Stanley J. Rosenschein and Leslie P. Kaelbling. A situated view of representation and control. *Artificial Intelligence*, 73:149–173, 1995.
18. M. J. Schoppers. Universal plans for reactive robots in unpredictable environments. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 1039–1046, 1987.
19. Yoav Shoham. Agent-oriented programming. *Artificial Intelligence*, 60(1):51–92, 1993.
20. K. Tam, J. Lloyd, Y. Lespérance, H. Levesque, F. Lin, D. Marcu, R. Reiter, and M. Jenkin. Controlling autonomous robots with GOLOG. In *Proceedings of the Tenth Australian Joint Conference on Artificial Intelligence (AI-97)*, pages 1–12, Perth, Australia, November 1997.
21. Kenneth Tam. Experiments in high-level robot control using ConGolog — reactivity, failure handling, and knowledge-based search. Master's thesis, Dept. of Computer Science, York University, 1998.
22. Rogier M. van Eijk, Frank S. de Boer, Wiebe van der Hoek, and John-Jules Ch. Meyer. Open multi-agent systems: Agent communication and integration. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin, 2000. In this volume.