

Understanding the Design Re

1988

Trial by Fire: Requirements for Agents in Complex Environments

Paul R. Cohen, Michael L. Greenberg, David M. Hart, and Adele E. Howe

The Phoenix Research Agenda

The Phoenix project is directed by three complementary goals. First, there are immediate technical aims: a real-time, adaptive planner for controlling simulated forest fires;




Figure 1. A Portion of Yellowstone National Park as Viewed within the Phoenix Simulator.

This figure shows a small fraction (about 7 km by 8 km) of the entire 75-km-square map. The northern tip of Yellowstone Lake is at the bottom of the screen. Thin black lines represent elevation contours, slightly wider lines represent roads, and the widest lines represent the fireline cut by bulldozers. Blue lines represent rivers and streams. The blue B in the bottom left corner marks the location of the fireboss, the agent that directs all the others. The white icons represent buildings. Shades of green represent different vegetation, and shades of red represent different fire intensities.

toward goals; and distributed planning algorithms. Second, there are motivating issues, of which the foremost is to understand how complex environments constrain on the design of intelligent agents. We seek general rules that justify and explain why an agent should be designed one way rather than another. The terms in these rules describe characteristics of environments, tasks and behaviors, and the architectures of agents. Last, because AI is still inventing itself, Phoenix is a commentary on the aims and methods of the field. Our position is that most AI systems have been built for trivial environments which offer few constraints on their design and, thus, few opportunities to learn how environments constrain and inform system design (Cohen 1989). To afford ourselves this opportunity, we began the Phoenix project by designing a real-time, spatially distributed, multiagent, dynamic, ongoing, unpredictable environment.

In the following pages, we describe Phoenix from the perspective of our technical aims and motives. The second section describes the Phoenix task—controlling simulated forest fires—and explains why we use a simulated environment instead of a real, physical one. Environmental Constraints on Agent Design discusses the characteristics of the forest fire environment and the constraints they place on the design of agents. The two lowest

layers of Phoenix, described in The Phoenix Environment, Layers 1 and 2, implement the simulated environment and maintain the illusion that the forest fire and agents are acting simultaneously. Above these layers are two others: a specific agent design (see Agent Design, Layer 3) and our organization of multiple fire-fighting agents (The Organi-

zation of Fire-Fighting Agents in Phoenix). These sections describe how Phoenix agents plan in real time but do not provide the minute detail that is offered elsewhere (Cohen et al. forthcoming). The next section illustrates Phoenix agents controlling a forest fire. The last section describes the current status of the project and our immediate goals.

The Problem

The Phoenix task is to control simulated forest fires by deploying simulated bulldozers, crews, airplanes, and other objects. We discuss how the simulation works in The Phoenix Environment, Layers 1 and 2 but concentrate here on how it appears to the viewer and the problems it poses for planners.

The Phoenix environment simulates fires in Yellowstone National Park, for which we have constructed a representation from Defense Mapping Agency data. Figure 1 shows a view of an area of the park.

Fires spread in irregular shapes and at variable rates, determined by ground cover, elevation, moisture content, wind speed and




Figure 3. Fire at 3:00 PM: Bulldozers Are Beginning to Build a Line around the Fire.




Figure 2. Fire at 12:30 PM: The Fireboss Is Alerted to New Fire and Dispatches Two Bulldozers to Fight It.

The left pane displays the actual state of the world; the right pane displays the current state of the world as the fireboss sees it.




Figure 4. Fire at 8:00 PM: The Fire is Nearly Encircled. Bulldozers are Close to Meeting at the Fire Front.

The left pane displays the real world; the right pane displays Bulldozer 11's view of the world.

direction, and natural boundaries. For example, fires spread more quickly in brush than in mature forest, are pushed in the direction of the wind and uphill, burn dry fuel more readily, and so on. These conditions also determine the probability that the fire will jump the fire-line and any natural boundaries. Except

Phoenix is a real-time, adaptive planner that manages forest fires in a simulated environment. Alternatively, Phoenix is a search for functional relationships between the designs of agents, their behaviors, and the environments in which they work. In fact, both characterizations are appropriate and together exemplify a research methodology that emphasizes complex, dynamic environments and complete, autonomous agents. Within the Phoenix system, we empirically explore the constraints the environment places on the design of intelligent agents. This article describes the underlying methodology and illustrates the architecture and behavior of Phoenix agents.

for two particular conditions (convective and crown fires), Phoenix is an accurate simulator of forest fires. Fire-fighting objects are also accurately simulated; for example, bulldozers move at a maximum speed of 40 kph in transit, 5 kph traveling cross-country, and 0.5 kph when cutting a fireline. To give a sense of scale, the fire in figure 1 is about 1.5 km in diameter and has burned for about eight simulated hours. The fire's history, which can be read in figures 2, 3, and 4, is as follows: At noon in simulation time (figure 2) a fire is ignited; later, it is detected by a watchtower (not visible in the figures). A little later, two bulldozers start a journey from the fire station, marked by a B in the southwest corner, to the rear of the fire. Because the wind is from the north, the rear is the northern end of the fire. At 3 PM (figure 3), the bulldozers arrive at the rear and start cutting a fireline. Figure 4 illustrates conditions at 8 PM simulation time. The fire is contained a little later. This entire simulation takes about 1 minute on a Texas Instruments Explorer.

Fires are fought by removing one or more of the elements that keep them burning: fuel, heat, and air. Cutting a fireline removes fuel. Dropping water and flame retardant removes heat and air, respectively. In major forest fires, controlled backfires are set to burn areas in the path of wildfires, thus denying them fuel. Huge project fires, such as those in Yellowstone last summer, are managed by many geographically dispersed firebosses and hundreds of firefighters.

The current Phoenix planner is a bit more modest. One fireboss directs a few bulldozers to cut a line near the fire boundary. Currently, we are implementing indirect attacks, which exploit natural boundaries as firebreaks (such as the river in figure 1), and parallel attacks, which involve backfires. The Phoenix planner does use common fire-fighting plans, such as

the "two bulldozer surround" illustrated in figures 2, 3, and 4. In this plan, two bulldozers begin at the rear of the fire and work their way around to the front, pinching it off.

The Phoenix fireboss directs bulldozers but does not completely control them. In fact, the fireboss gives fairly crude directions, such as "go to location

x,y," and individual agents decide how to interpret and implement them. Thus, bulldozers and other agents are semiautonomous. Other organizational structures are enabled by increasing or decreasing the degree of autonomy; for example, an earlier fire planner, designed by David Day, had a single fireboss that controlled every action of all its agents. At the other extreme, we are working with Victor Lesser on a completely distributed version of Phoenix, in which agents negotiate plans in the absence of a single fireboss. We can experiment with different organizational structures because all agents have exactly the same architecture, and so each can assume an autonomous, semi-autonomous, or completely subservient role.

Although Phoenix agents and their environment are all part of a large software system, we designed them to give the impression of independent agents playing against simulated forest fires, much as we play a video game. In fact, early in the project, we built an interface to allow us, instead of an automated planner, to direct fire-fighting agents. It required us to simultaneously control several agents, and demanded considerable foresight and planning. We found it impossible to control more than a couple of bulldozers in real time in the vicinity of the fire, so we gave bulldozers simple reflexes, enabling them to scurry away from encroaching fire. Since this time, the basic style of interaction between the Phoenix environment and the Phoenix planners has not changed: One or more planners, AI or human, direct semiautonomous agents to move around a map, building a line around continuously burning fires.

To some, the decision to develop and test Phoenix agents in a simulated environment is profoundly wrong. One argument is that by building the environment and the interface

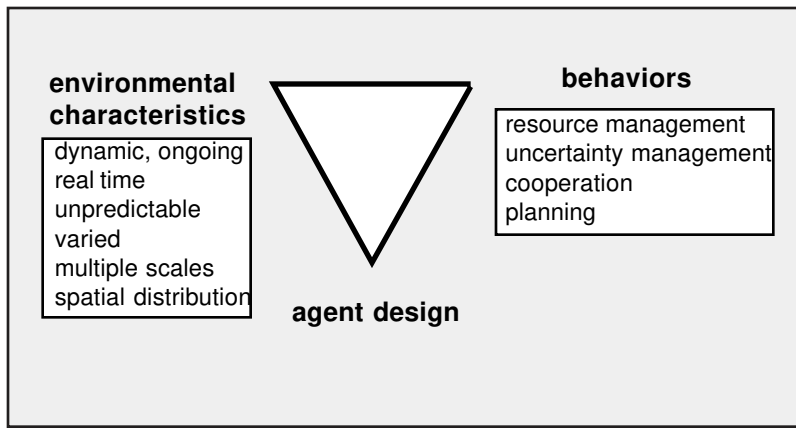


Figure 5. The Behavioral Ecology Triangle.

to agents, we risk deferring or ignoring difficult problems. For example, if we build a simulated agent that has a completely accurate internal map of its simulated environment, and when it moves, its “wheels” don’t slip, then all its planning and acting can be dead reckoning. Obviously, we can create trivial environments and develop techniques that won’t work in real environments, but why would we? The point of using simulators is to create realistic and challenging worlds, not to avoid these challenges. In response to the criticism that simulators can never provide faithful models of the real, physical world, we argue that the fire environment is a real-time, spatially distributed, ongoing, multiactor, dynamic, unpredictable world, irrespective of whether it is an accurate model of how forest fires spread. As it happens, the fire environment is an accurate model of forest fires, but this accuracy isn’t necessary for the environment to challenge our current planning technology. Moreover, we want to leave open the possibility of working in simulated worlds which are unlike any physical world that we have encountered.

The advantages of simulated environments are that they can be instrumented and controlled and provide variety, all essential characteristics for experimental research. Specifically, simulators offer control, repeatability, replication, variety, and interfaces.

Control

Simulators are highly parameterized, so we can experiment with many environments. For example, to test the robustness of real-time planning mechanisms, we can change

the rate at which wind direction shifts or speed up the rate at which fire burns. Most important, from the standpoint of our work on real-time planning, is the fact that we can manipulate the amount of time an agent is allowed to think relative to the rate at which the environment changes, thus exerting (or decreasing) the time pressure on the agent (see The Phoenix Environment, Layers 1 and 2).

Repeatability

We can guarantee identical initial conditions from one run to the next; we can exactly play back some histories of environmental conditions and selectively change others.

Replication

Simulators are portable; so they enable replications and extensions of experiments at different laboratories. They enable direct comparisons of results, which would otherwise depend on uncertain parallels between the environments in which the results were collected.

Variety

Simulators allow us to create environments which don’t occur naturally or which aren’t accessible or observable.

Interfaces

We can construct interfaces to the simulator that allow us to defer questions we would have to address if our agents interacted with the physical world, such as the vision problem. We can also construct interfaces to show elements that aren’t easily observed in the physical world; for example, we can show the different views which agents have of the fire, their radius of view, their destinations, the paths they are trying to follow, and so on. The Phoenix environment graphics make it easy to see what agents are doing and why.

Environmental Constraints on Agent Design

From the preceding descriptions of the Phoenix environment and tasks, one can begin to see the challenges they present to Phoenix agents. As researchers, our challenge is to design these agents for the Phoenix environment. The relationships between agent design, desired agent behaviors, and environment characteristics are clarified by what we call the behavioral ecology triangle, shown in figure 5.

The vertices of the triangle are the agent’s design (that is, internal structures and processes), its environment, and its behavior

(that is, the problems it solves and the ways it solves them). In this context, our five tasks (and, indeed, the tasks of all AI research on intelligent agents) are:

1. Environment Analysis: What characteristics of an environment most significantly constrain agent design?
2. Design: What architecture produces the desired behaviors under the expected range of environmental conditions?
3. Prediction: How does a particular agent behave in particular environmental conditions?
4. Explanation: Why does an agent behave as it does in particular environmental conditions?
5. Generalization: Over what range of environmental conditions can we expect particular behaviors from the agent? Over what range of problems? Over what range of designs?

To date, the Phoenix project has concentrated on environment analysis (later in this section), the design task (see Agent Design, Layer 3 and The Organization of Fire-Fighting Agents in Phoenix), and the building of an environment in which the other tasks can be empirically pursued. Figure 5 implicitly captures many hypotheses and explanatory tasks. We can think of anchoring two corners and solving for a third. For example, we can anchor an environment and a set of behaviors and solve for an agent design, or we can anchor a design and an environment and test predictions about behavior. Another more exploratory research strategy is to anchor just one corner, such as the environment, and look for trade-offs in the other corners. For example, given the Phoenix environment, how is adaptability to changing time pressures affected by the design decision to search for plans in memory rather than generate them from scratch?

Let us survey the characteristics of the Phoenix environment that constrain the design of Phoenix agents and the behaviors the agents must display to succeed at their tasks. The fire environment is dynamic because everything changes: wind speed and direction, humidity, fuel type, the size and intensity of the fire, the availability and position of fire-fighting objects, the quantity and quality of information about the fire, and so on. The environment is ongoing in the sense that there isn't a single, well-defined problem to be solved, after which the system quits, but rather a continuous flow of problems, most of which are unanticipated. The environment is

How do the characteristics of the Phoenix environment, in concert with the desired behaviors of Phoenix agents, constrain the design of the agents?

real time in the sense that the fire sets the pace to which the agent must adapt. The agent's actions, including thinking, take time, and during this time, the environment is changing. These characteristics require an agent to have some concept of relative or passing time. The agent must reason about the potential effects of its actions and, particularly, about how much time these actions might require. Additionally, it must be able to perceive changes in its environment, either directly through its own senses or indirectly through communication with other agents.

The environment is unpredictable because fires can erupt at any time and any place; because weather conditions can change abruptly; and because agents can encounter unexpected terrain, fire, or other agents as they carry out plans. An agent must respond to the unexpected outcomes of its own actions (including the actions taking more or less time than expected) and to changes in the state of the world. This responsiveness requires interleaving planning, execution, and monitoring and suggests that detailed plans of long duration are likely to fail before successful completion. The unpredictability of the environment requires agents to be flexible, particularly in the way they handle temporal resources. In fact, all resources, including time, fire-fighting agents, money, fuel, and equipment, are limited and nonrenewable. Because the environment is ongoing, decisions about resources have long-term effects that constrain later actions and require agents to manage their resources intelligently, with a global perspective. For this reason, among others, Phoenix agents cannot be exclusively reactive.

Whereas unpredictability is a characteristic of the Phoenix environment, uncertainty arises in agents. Uncertainty is the result of the fire continuously moving, unpredictable changes in wind speed and direction, communication delays between agents, and the limited world views of the individual agents. For example, to the northwest of Bulldozer

11, in the right-hand pane of figure 4 is a small black patch of the fireline. This information is all Bulldozer 11 knows about the location and progress of the other bulldozer (whose actual location is shown in the left-hand pane of figure 4) and illustrates how far Bulldozer 11 can see. It follows that Bulldozer 11's fire map, as shown in the right-hand pane, must merge what it currently sees with what it recalls. As one would expect, the recollection is inaccurate; Bulldozer 11 thinks the fire at its northern point is a few hundred meters from the fireline because that's where it was when Bulldozer 11 cut the fireline. In fact, the fire has spread all the way to the fireline, as shown in the left-hand pane. As a consequence of these types of uncertainty, agents must allot resources for information gathering. Agents must be able to integrate and disseminate local information, and because of their own localized views, they must be able to communicate and coordinate with each other.

The fact that events happen at different scales in the Phoenix environment has profound consequences for agent design. Temporal scales range from seconds to days, spatial scales from meters to kilometers. Agents' planning activities also take place at disparate scales; for example, a bulldozer agent must react quickly enough to follow a road without straying as a result of momentary inattention and must plan several hours of fire-fighting activity, both within the time constraints imposed by the environment.

Given the size and variation in the world map, the degree to which the environment can change, and the possible actions of agents, the environment can produce a large variety of states. Consequently, an agent must know how to act in many different situations. The ramifications for agent design depend on whether small differences in environmental conditions can produce large differences in plan utilities. For example, if every fire scenario is truly different in the sense that each requires a unique, scenario-specific plan, then it might be pointless to provide agents with memories of previous plans. In fact, we believe that although the fire environment presents a wide variety of states, these differences do not require radically different plans.

The Phoenix environment is spatially distributed, and individual agents have only limited, local knowledge of the environment. Moreover, most fires are too big for a single agent to control; their perimeters grow much faster than a single agent can cut a fireline. These constraints dictate multiagent, dis-

tributed solutions to planning problems. They also expand the scope of our research from a study of agent design to a study of organizational design. We have drawn a line, temporarily, and excluded the latter.

In sum, to perform their designated tasks under the constraints of the Phoenix environment, Phoenix agents must engage in particular behaviors. In gross terms, these behaviors are resource management, uncertainty management, planning, and cooperative problem solving; we just discussed the more specific behaviors. The question we address in Agent Design, Layer 3 and The Organization of Fire-Fighting Agents in Phoenix is how the characteristics of the Phoenix environment, in concert with the desired behaviors of Phoenix agents, constrain the design of the agents. Specifically, what architecture is capable of planning in real time, responding to events at different time scales, coordinating the efforts of several agents, collecting and integrating data about a changing environment, and so on.

The Phoenix Environment, Layers 1 and 2

To facilitate experiments, Phoenix is built in four layers. The lowest is a task coordinator that on a serial machine, maintains the illusion of simultaneity among many cognitive, perceptual, reflexive, and environmental processes. The next layer implements the Phoenix environment itself—the maps of Yellowstone National Park and the fire simulations. The third layer contains the definitions of the agent components—our specific agent design. The fourth layer describes the agents organization, their communication and authority relationships. Layers 3 and 4 are described in later sections.

The two lowest layers in Phoenix, called the task coordinator layer and the map layer, respectively, make up the Phoenix discrete event simulator. We discuss the task coordinator first. It is responsible for the illusion of simultaneity among the three following events and actions.

1. Fires. Multiple fires can burn simultaneously in Phoenix. Fires are essentially cellular automata that spread according to local environmental conditions, including wind speed and direction, fuel type, humidity, and terrain gradient.
2. Agents' physical actions. Agents move from one place to another, report what they perceive, and cut a fireline.

3. Agents' internal actions. Internal actions include sensing, planning, and reflexive reactions to immediate environmental conditions.

These tasks are not generated at the task coordinator level of Phoenix, just scheduled on the central processing unit (CPU) there. Fire tasks are generated at the map layer, and agent tasks are generated as described in later sections.

Typically, the task coordinator manages the physical and internal actions of several agents (for example, one fireboss, four bulldozers, and a couple of watchtowers) and one or more fires. The illusion of continuous, parallel activity on a serial machine is maintained by segregating each process and agent activity into a separate task and executing them in small, discrete time quanta, ensuring that no task ever gets too far ahead or behind the others. The default setting of the synchronization quantum is five minutes, so all tasks are synchronized within five minutes of each other. The quantum can be increased, which improves the CPU utilization of tasks and makes the testbed run faster but increases the simulation-time disparity between tasks, magnifying coordination problems such as communicating and knowing the exact state of the world at a particular time. Conversely, decreasing the quantum reduces how out of sync processes can be but increases the running time of the simulation.

The task coordinator manages two types of time: CPU time and simulation time. CPU time refers to the length of time that processes run on a processor. Simulation time refers to the time of day in the simulated environment. Within the predefined time quantum, all simulated parallel processes begin or end at roughly the same simulation time. To exert real-time pressure on the Phoenix planner, every CPU second of thinking is followed by K simulation-time minutes of activity in the Phoenix environment. Currently $K = 5$, but this parameter can be modified to experiment with how the Phoenix planner copes with different degrees of time pressure.

The fire simulator resides in Phoenix's map layer; that is, the map layer generates tasks

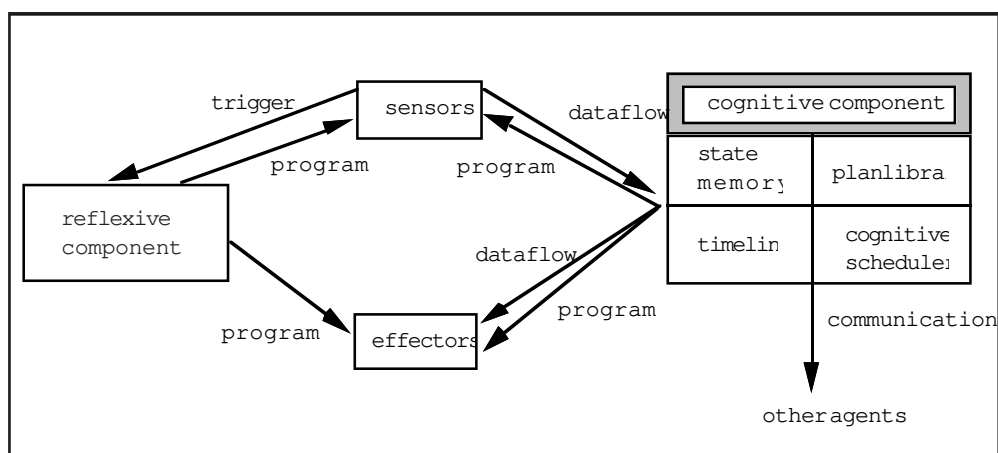


Figure 6. Phoenix Agent Design.

that produce dynamic forest fires when executed by the task coordinator. Phoenix's map, which represents Yellowstone National Park, is a composite of several two-dimensional structures and stores information for each coordinate about ground cover, elevation, features (roads, rivers, houses, and so on), and fire state. The fire itself is implemented as a cellular automaton in which each cell at the boundary decides whether to spread to its neighbors, depending on the local conditions just mentioned and global conditions such as wind speed and direction (currently, we do not model local variations in weather conditions). These conditions also determine the probability that the fire will jump the fireline and any natural boundaries.

The Phoenix discrete event simulation is generic. It can manage any simulations that involve maps and processes. For example, we could replace the forest fire environment with an oil spill environment. We could replace our map of Yellowstone with oceanographic maps of, say, Prince William Sound in Alaska. Fire processes have spatial extent and spread according to wind speed, direction, fuel type, terrain, and so on. They could easily be replaced with oil slick processes, which also have spatial extent and spread according to other rules. Similarly, we could replace the definitions of bulldozers and airplanes with definitions of boats and booms.

Agent Design, Layer 3

The third layer of Phoenix is our specific agent design, which is constrained by the forest fire environment as described earlier. For example, because events happen on two dramatically different time scales, we




Figure 7. Contents of Fireboss's Timeline after Being Notified of a New Fire: Action to Search for a Plan to Deal with the Fire.

The combination of a reflexive and cognitive component is designed to handle time scale mismatches inherent in the fire-fighting environment.

designed an agent with two parallel and nearly independent mechanisms for generating actions (figure 6). One mechanism quickly generates reflexive actions—on the order of a few seconds of simulated time—and the other generates plans that can take hours of simulated time to execute. This longer-term planning can be computationally intensive because it incurs a heavy time penalty for switching contexts when interrupted. For this reason, the cognitive component is designed to perform only one action at a time (unlike sensors, effectors, or reflexes where multiple activities execute in parallel). Both the cognitive and reflexive components have access to sensors, and both control effectors, as shown in figure 6.

The agent interacts with its environment through its sensors and effectors, and action is mediated by the reflexive and cognitive components. Sensory information can be provided autonomously or can be requested, and sensors' sensitivity can be adjusted by the cognitive component. Effectors produce actions in the world, such as gathering information, building a fireline, and moving.

Reflexes are triggered by sensor output. They change the programming of effectors to prevent catastrophes or fine tune the operation of effectors. For example, a bulldozer is stopped by a reflex if it is about to move into the fire, and reflexes handle the fine tuning necessary for the bulldozer to follow a road. Reflexes are allotted almost no CPU time and have no memory of events, so they cannot produce coordinated sequences of actions. They are designed for rapid, unthinking action. Although some researchers have suggested that longer-term plans can emerge from compositions of reflexes (Agre and Chapman 1987; Brooks 1986), we do not believe that compositions of reflexes can handle temporally extensive planning tasks

such as resource management or spatially extensive tasks such as path planning with rendezvous points for several agents. Thus, we have adopted a design in which reflexes handle immediate tasks, and a cognitive component handles everything else.

The cognitive component of an agent is responsible for generating and executing plans. Instead of generating plans de novo, as classical hierarchical planners did, the Phoenix cognitive component instantiates and executes stored skeletal plans. We believe this design is a good one for the forest fire environment because, first, a relatively small number of skeletal plans is probably sufficient to cope with a wide range of fires and, second, the store-recompute trade-off suggests relying on stored plans, rather than computing them, in real-time situations. In addition to controlling sensors and effectors, the cognitive component handles communications with other agents (including integrating sensor reports), and it responds to flags set when reflexes execute. It also engages in a wide range of internal actions, including projection (for example, where the fire will be in 20 minutes), plan selection and scheduling, plan monitoring, error recovery, and replanning. Our implementations of some of these capabilities are quite rudimentary and leave much room for improvement, as we discuss in Current Status and Future Work.

As an overview, this is how the cognitive component works: In response to a situation such as a new fire, an appropriate plan is retrieved from the plan library and placed on the timeline (figure 6). State memory stores information, such as weather, resource conditions, and sensory input, that helps the cognitive agent select appropriate plans and instantiate the variables of the chosen plan for the current situation. For example, if the fire is small and nearby, and the weather is calm, then a one-bulldozer plan is retrieved and instantiated with situation-specific information such as the wind speed and the current location of the fire. The actions in a plan are eventually selected for execution by the cognitive scheduler, which is described shortly. At any time during this process, sensory data can trigger reflexive actions; for example, if the cognitive component is executing a command to move to a destination, and a sensor reports fire ahead, then the reflexive component sends a command to reverse direction. This reaction happens very fast relative to the cycle time of the cognitive component, so the reflexive component sets a flag to tell the cognitive component what it did. When the cognitive component notices

the flag, it might modify its plan. The analogy here is to our own reflexes, which yank us away from hot surfaces long before our cognitive apparatus becomes aware of the problem.

With this overview in mind, let us consider the operation of the cognitive component in detail. We focus on the operation of the fireboss agent, which plans the activities of other agents such as bulldozers and crews. Each of these, in turn, plans how to carry out the directives of the fireboss. Because bulldozers and crews have the same architecture as the fireboss (figure 6), they can reason in exactly the same way. In the following discussion, we first describe planning when things go according to plan and then describe error handling, interruptions, and other unexpected events.

When a fire is reported, an action called “deal with fire” is retrieved from the plan library and used to create a timeline entry, in this case called “deal with fire 27,” which is added to the timeline (see figure 7). *Actions* are general representations of the cognitive activities the agent can perform, such as path planning or communication, and describe applicability conditions, resource constraints, and uninstantiated variables. Creating a timeline entry instantiates an action, binding its variables and adding the temporal constraints that relate it to other actions the agent has chosen to execute. Although timeline entries represent actions, it is not quite accurate to say they are executed (although we use this terminology where the accurate description is too awkward). In fact, when a timeline entry is created, it inherits a set of execution methods from the action it instantiates. Each of these methods executes the desired action; they differ along dimensions, such as the time they require and the quality of their output. For example, a single action “plan a path” points to several path-planning algorithms, some which run quickly and return adequate paths and some which run longer but produce shorter paths. When a timeline entry is selected for execution, the execution method most appropriate to the current circumstances is chosen. By delaying the choice of methods, the cognitive scheduler can reason about its own use of time and select execution methods that are suited to emerging time constraints.

If entries exist on the timeline (for example, “deal with fire 27”), then the cognitive scheduler of the Phoenix cognitive component makes three decisions: (1) which action to execute next, (2) how much time is available for its execution, and (3) what execution method should implement the action.

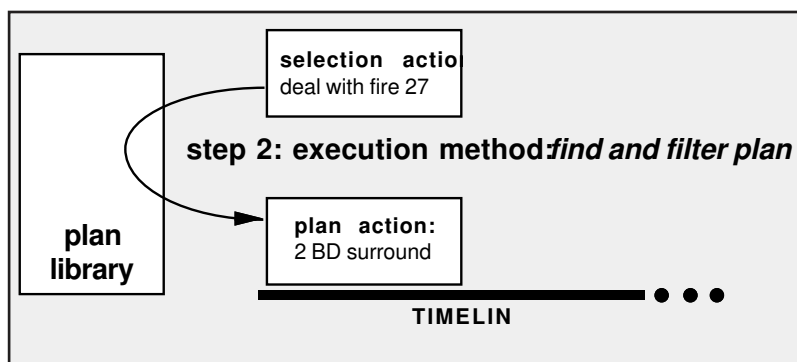


Figure 8. The Fireboss Executes Timeline Action “deal with fire 27.”

This action searches the plan library, selects the 2 BD surround plan as appropriate for dealing with a new fire, and places the new plan on the timeline.

The cognitive scheduler always selects the next action on the timeline to execute, but often, several actions have this distinction, and a choice must be made. Actions on the timeline can be unordered (and, thus, equally entitled to go first) for several reasons: Skeletal plans often leave actions unordered so that the cognitive scheduler has flexibility at execution time to select the best order. Frequently, the agent is executing several plans simultaneously, for example, when several fires are reported. The planner formulates plans for each but doesn’t specify temporal constraints among actions from different plans. In the current example, however, the only action on the timeline is “deal with fire 27,” so the cognitive scheduler determines how much time is available to execute it and selects an execution method. In this case, it selects a method called the find-and-filter plan (step 2, figure 8). When executed, its effect is to search the plan library for a plan to “deal with fire 27.” First, it finds all plans for dealing with fires of this type; then it filters the infeasible ones; it selects from the candidates to find the most appropriate one; and last, it adds a new action to the timeline called “2 BD surround.” (This plan involves sending two bulldozers to a rendezvous point, then to the fire, after which they cut a fireline in opposite directions around the fire.)

Once again, the cognitive scheduler selects an action (the only one is “2 BD surround”), assesses how much time is available, and selects an execution method. In this case, the method is to expand the plan. The result is to add a network of actions, partially ordered over time, to the timeline (step 3, figure 9). The network starts with a placeholder action, s, followed by two unordered actions that allocate bulldozers 1 and 2, respectively. The next action determines the rendezvous point

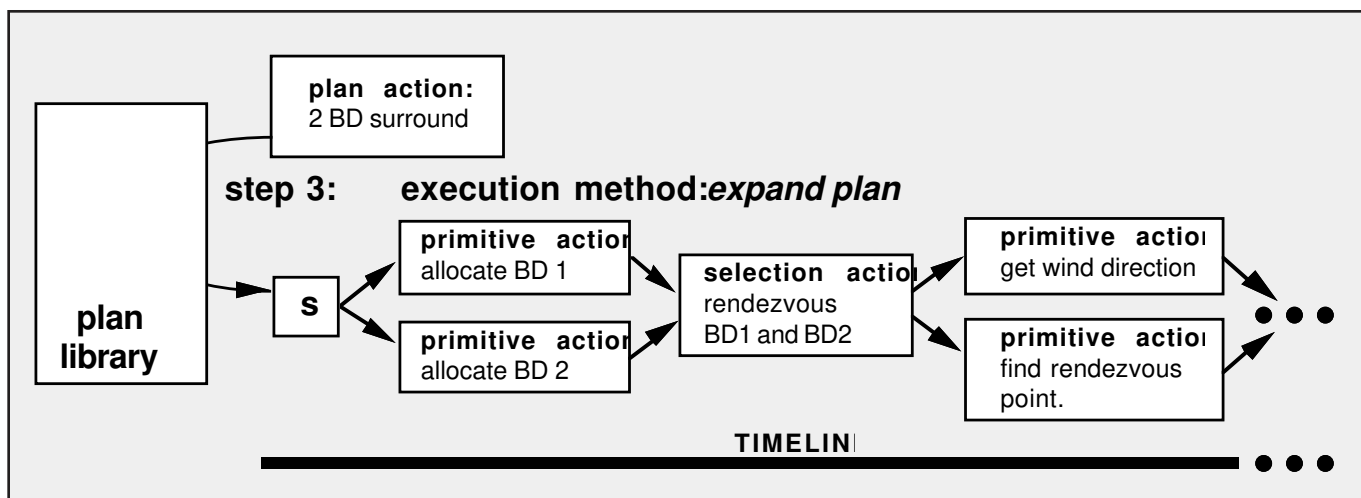


Figure 9. The Fireboss Executes the Timeline Action, “2 BD surround,” which Expands into a Network of Plan Steps.

for the bulldozers. Then two unordered actions bind the variables in the plan with the current wind direction and the previously determined rendezvous point. (Space precludes showing the rest of the plan in figure 9.)

The cognitive scheduler again looks at the timeline and now must make a decision about which action to select. The “allocate bulldozer” actions are unordered, so one must be selected to go first. Then, as before, the cognitive scheduler assesses the available time and selects an execution method. We discuss this example further, later.

Three kinds of actions can be differentiated by their effects on the timeline when they are executed: First, *selection actions*, such as “deal with fire 27,” result in a search of the plan library. Next, a *plan action*, such as “2 BD surround,” is posted on the timeline. Plan actions are placeholders for plans; executing them results in plan expansions being posted on the timeline. Many of the actions in a plan are of the third type, *primitive actions*, that result in a computation (for example, calculating a route) or a command to a sensor or effector. However, a plan can contain any of the three types of actions; for example, the expansion of “2 BD surround” contains a selection action. When executed, it results in a search of the plan library for a plan to have the two bulldozers rendezvous. Plans can also contain plan actions, which, when executed, add subplans to the network. This nesting of plans within plans is our mechanism for representing hierarchical plans. Last, plans can contain just a single, primitive action, such as finding the rendezvous point for two bulldozers.

We discussed how actions are scheduled and executed when everything goes accord-

ing to plan, but in the Phoenix environment it rarely does. Phoenix agents have three abilities, all rudimentary, for handling unexpected events. Reflexes, operating on a short time scale, can halt or modify potentially injurious actions, such as straying into the fire. By design, reflexes do little processing and return little information. When a reflex halts a bulldozer, it simply posts a flag for the cognitive component; it does not interrupt the cognitive component to explain what it did. The cognitive component doesn’t become aware of the change until it executes a regularly scheduled status-checking action. In fact, by design, nothing ever interrupts a cognitive action because the cost of saving state and switching context is prohibitive. Instead, the reflexive component of a Phoenix agent is expected to deal with situations as they arise. Most situations, such as staying parallel to a moving fire, never require the attention of the cognitive component; but even when a serious problem comes up, the reflexive component is designed to keep the agent functioning until the cognitive component finishes its current task.

The second mechanism for handling unexpected situations is error recovery and replanning. *Errors* are unexpected events that preclude the completion of an action or a plan. For example, bulldozers can travel to their designated destinations but fail to find a fire, path planning sometimes fails to generate a path, selection actions can search the plan library but fail to find a plan that satisfies all constraints, and so on. Currently, over a dozen error types can arise in Phoenix, although we don’t have plans to deal with all of them yet. The error-handling mechanism is to post on the timeline a “deal with error”

Lazy skeletal refinement responds to a complex dynamic world by postponing decisions, while grounding potential actions in a framework that accounts for data, temporal and resource interactions.

selection action, which, when executed, generates a plan for dealing with the error. Currently, error recovery involves little tinkering with the actions that are currently on the timeline, that is, no serious replanning.

Last, Phoenix agents have limited abilities for monitoring their own progress. This monitoring is accomplished by generating expectations of progress and matching them to actual progress. In the near future, this mechanism (called envelopes; see Current Status and Future Work) will enable Phoenix's cognitive components to predict failures before they occur.

In sum, planning is accomplished by adding a selection action to the timeline to search for a plan to address some conditions. Executing the selection action places an appropriate plan action or primitive action on the timeline. If this new entry is a plan action, then it expands into a plan when it is executed by putting its subactions onto the timeline with their temporal interrelationships. If it is a primitive action, execution instantiates the requisite variables, selects an execution method, and executes it. In general, a cognitive agent interleaves actions from the several plans it is working on.

This style of planning is lazy skeletal refinement, lazy because some decisions are deferred until execution time. Specifically, plans are not selected until selection actions are executed, and execution methods are selected only when an action is about to execute. This style of planning and acting is designed to be responsive to a complex dynamic world by postponing decisions while grounding potential actions in a framework (a skeletal plan) that accounts for data, temporal, and resource interactions. The combination of a reflexive and cognitive component is designed to handle time-scale mismatches inherent in an environment that requires microactions (for example, following a road) and contemplative processing such as route planning, which involves long search times and the integration of disparate data. We must stress, however, that it is too early in Phoenix's development to claim that our agent design is necessarily the best one for the Phoenix environment (see Current Status and Future Work).

The Organization of Fire-Fighting Agents in Phoenix

The fourth level of the Phoenix system is the centralized, hierarchical organization of fire-fighting agents. Because all agents have the same architecture, many other agent organizations are possible. Our centralized model is neither robust (for example, what happens if the fireboss is disabled) nor particularly sophisticated. However, it is simple, which is a great advantage in these initial phases of the project. One fireboss coordinates all fire-fighting agents' activities, sending action directives and receiving status reports, including fire sightings, position updates, and actions completed. The fireboss maintains a global view of the fire situation based on these reports, using it to choose global plans from its plan library. It communicates the actions in these plans to its agents, which then select plans from their own plan libraries to effect the specified actions. Once their plans are set in motion, agents report progress to the fireboss, from which the execution of global plans is monitored. All communication in this centralized implementation is between the fireboss and the individual agents; no cross-talk takes place among the agents.

The fireboss maintains global coherence, coordinating the available fire-fighting resources to effectively control the fire. It is responsible for all the work required to coordinate agents, such as calculating rendezvous points, deciding how to deploy available resources, and noticing when the fire is completely encircled by the fireline. The plans in its plan library are indexed by global factors, such as the size of the fire and the weather conditions. The actions in its plans are mostly concerned with coordinating and directing other agents. The fireboss's state memory records the current environmental conditions, where agents have seen fire, what actions have been taken, what agents are available, and how well global plans are progressing. The fireboss is currently implemented without any sensors, effectors, or reflexes. It is a cognitive agent that relies solely on communication for its knowledge of what develops in the outside world, although

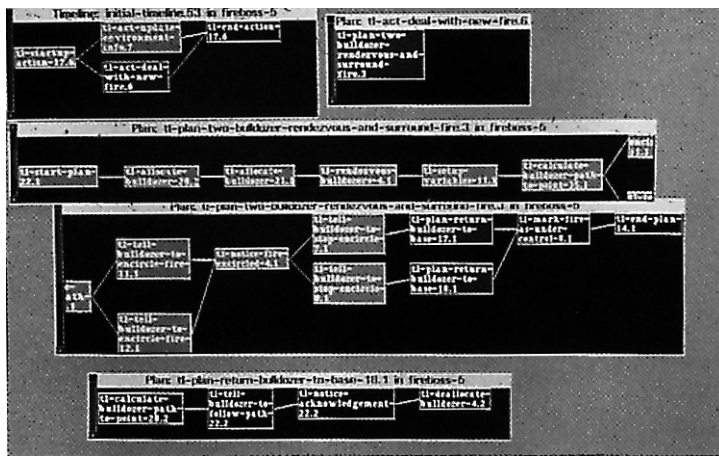


Figure 10. The Fireboss's Timeline, including Plan Expansions as Viewed from the Phoenix Desktop (headers for boxes name parent nodes).

it does have a map of the static features of Yellowstone.

Each of the other fire-fighting agents has a local view of the environment based on its own sensory input. They have access to maps of the static features in Yellowstone, such as ground cover, roads, and rivers, but only know about dynamic processes such as the fire from what they see or are told by the fireboss. Sensors have a limited radius of view, although agents are able to remember what has been perceived but is no longer in view. The fireboss's global view is only available to an agent through communication. A bulldozer is an example of an agent type. It has a movement effector that can follow roads or travel cross-country. When it lowers its blade while moving, it digs a fireline and moves more slowly. It has a sensor that sees fire within a radius of 512 meters. Another sensor picks up the contour of a fire (within its radius of view). When a bulldozer is building a fireline at the contour, it uses the follow-fire sensor in combination with the movement effector (with lowered blade) and a reflexive action that helps maintain a course parallel to the contour. As the contour changes, the contour sensor registers the change, which triggers a reflex to adjust the movement effector's course. The bulldozer's plan library has plans for simple bulldozer tasks such as following a given path or encircling a fire with a fireline.

Although all agents have the same architecture (that is, timeline, cognitive scheduler, plan library, state memory, sensors, effectors, and reflexes), they do not have the same plans, reflexes, sensors, or effectors. The dif-

ference between the fireboss and the other agents lies in their views of the world and the types of plans each knows. The lines of authority and divisions of responsibility are clear: The fireboss maintains the global picture, based on the local views of its agents, and it executes plans whose effects are to gather information, send directives to agents, and coordinate their activity through communications. In contrast, the agents execute plans whose actions program sensors and effectors, which, in turn, effect physical actions in the world. In some sense, the fireboss is a metaagent whose sensors and effectors are other agents.

An Example

We now return to the example that we introduced earlier to illustrate cognitive scheduling. In this two-bulldozer-surround plan, the fireboss instructs two bulldozers to rendezvous, then go to the fire and build a fireline around it in opposite directions. Figures 2, 3, and 4 show the progress of this plan. Each offers two views of the situation. Figure 2 shows the real world in the left pane and the fireboss's view in the right pane. Note that the fireboss is unaware of the true extent of the fire. What it knows is based on status reports from a watchtower agent (not shown). Each watchtower has a sensor programmed to look for new fires at regular time intervals. When it spots a fire, it reports the location and size to the fireboss. Based on this report and the resources available, the fireboss selects the two-bulldozer-surround plan. The first plan steps allocate the bulldozers, which ensures they are not busy with other tasks and assigns them to this plan. The next step instructs them to rendezvous so they can follow the same route to the fire. Next, the fireboss plan calculates a route for the bulldozers to follow, then executes two steps in which it communicates instructions to each bulldozer to follow the given path and encircle the fire. They are given clockwise and counterclockwise encircling directions, respectively.

After receiving its instructions, each bulldozer searches its plan library to find a plan for following its given path and encircling the fire in the given direction until it closes the fireline. Neither bulldozer knows about the other, nor does either know the full extent or precise location of the fire. Recall that the fireboss doesn't know exactly where the fire is either, so the path it supplied to the bulldozers can direct them wide of the fire or, more often, to a location that is burning. In

this example, the path given by the fireboss ends at a point engulfed in fire, so the bulldozers follow the path until they detect it. In figure 3, we see the bulldozers starting to build a line. In the fireboss's view (right pane), each one appears at the position it had reached when it made its last status report. Thus, they are at slightly different positions that are out of date with respect to their real positions in the left-hand pane.

When fire is seen, a bulldozer reflex is triggered to stop its movement effector. A cognitive action also notes that the sensor has seen fire and reprograms the sensors and effectors with the right combination of instructions to follow the fire in the direction specified by the fireboss, building a fireline as it goes. A message is sent to the fireboss to signal the start of line building. Each bulldozer stops building a line when it finds another fireline and closes it off. In figure 4, we see in the left pane that the bulldozers have almost encircled the fire. In the right pane is the view of the bulldozer encircling in the clockwise direction. Note that it only knows about the fire it has seen as it was building line. It is just coming within range of the other bulldozer (see the spot of fireline to the northwest of the bulldozer).

This simple bulldozer plan, to follow a path and encircle a fire without reference to other bulldozers, can be used by one, two, or many bulldozers. With its global view, the fireboss picks points around the fire, selects any number of bulldozers, and directs each to go to one of the points and build a fireline in a specified direction. The bulldozers only act with regard to their instructions and the local information in their field of view. If the bulldozers fail to fully encircle the fire (for whatever reason), the fireboss is responsible for noticing the failure, based on what is reported to it from watchtowers and bulldozers.

Figure 10 shows the fireboss's timeline for this plan at about the time the bulldozers complete their encirclement of the fire. The network in the top, left box is the top level of the timeline, which contains four entries and reads from left to right. The entries are start action and end action (placeholders) and two with no temporal constraint between them. The top entry is an action that periodically executes and updates state memory with new information about the environment. The bottom entry is an action that is automatically placed on the fireboss's timeline with the report of a new fire. This selection action causes a plan to be selected from the fireboss's plan library, based on the characteristics of the reported fire, and then expanded on the

timeline, as illustrated in Agent Design, Layer 3. The selected plan is shown in the top, right box, and its expansion is shown in the two middle boxes (the plan unfolds left to right and is continued in the lower, middle box). Gray actions have already been executed. These actions include allocating each bulldozer, instructing them to rendezvous, calculating a route for them to take to the fire, and (in undetermined order) instructing them to follow this route and encircle the fire. They also include noticing when the fire is completely encircled and instructing the bulldozers to stop building a fireline. These actions are necessary because the fireboss maintains the global view of the fire and must tell the bulldozers when the fire is surrounded. The bulldozers' timelines are not shown in figure 10.

The fireboss is currently executing a plan action for each bulldozer that expands into a plan for sending the bulldozer back to the base. Note that these actions are not temporally ordered in the plan, so the fireboss has to decide on an order. The timeline expansion for this plan (for one of the bulldozers) is shown in the bottom box, where the currently executing action is calculating a path for the bulldozer. Subsequent steps instruct the bulldozer to follow the path, wait for an acknowledgment that the bulldozer has reached the base, and make a mental note about the bulldozer's availability to fight a new fire.

Current Status and Future Work

The Phoenix system is a work in progress. As is clear from the preceding sections, several important aspects are handled in a rudimentary or preliminary way. Currently, five people are pursuing research and enhancing the system in the following areas.

We have about a dozen plans that directly attack fires, with as many as four bulldozers building a line at the fire front until the fire is encircled. We are starting to develop indirect attack plans that incorporate natural barriers. These plans require more knowledge and coordination on the part of the fireboss; because other agents can't see the fire unless they are close to it, the fireboss must guide their activities when they are working at a distance. The fireboss must take advantage of natural barriers when deciding where to build a fireline, which requires the ability to project the spread of the fire and the progress of the fire-fighting agents. As we develop new and more sophisticated plans, we must also enhance the mechanisms by which agents select plans. Currently, the keys for selecting

The development of Phoenix has been intimately tied to our evolving ideas about AI research methodology, and specifically to our understanding of the role of evaluation in AI research.

plans are just wind speed and the availability of fire-fighting agents, as well as some plan-specific criteria such as whether bulldozers are nearby or distant when the plan is selected. The keys have to become more discriminating, and we probably have to develop more sophisticated plan-selection mechanisms.

We have designed a general monitoring mechanism called *envelopes* that minimizes the cognitive resources devoted to monitoring yet provides early warning of plan failure. Envelopes incorporate expectations of how plans are to proceed; they represent these expectations as functions. As actual progress is reported, it is compared with these expectations, and deviations outside certain parameterized thresholds are flagged for cognitive attention. For example, if an agent must be at a certain place at a certain time, we can tell by projection whether the deadline is feasible—can the agent travel the distance in the given time? By projecting the expected time of travel (based on a parameter such as average speed for the agent on the given terrain), we can create an envelope for the travel time and use it to monitor the agent's progress. The envelope also predicts the expected arrival time based on the recent progress of the agent. Furthermore, it predicts the minimum speed at which the agent must travel over the remaining distance to arrive before the deadline. If this speed is at or approaching the top speed of the agent, then the envelope signals the planner that the deadline is in jeopardy, providing an early failure warning. Currently, we have hooks for envelopes in plans, and limited mechanisms to replan when envelopes are violated. For example, if the fire spreads quicker than anticipated, a slack time envelope is violated and the fireboss sends one or more additional bulldozers to help out.

Error recovery and replanning are implemented as cognitive actions, just like plan selection and plan expansion. When an error is detected in a plan, an action is posted to the timeline that inspects the error and attempts to fix the existing plan. Consider, for example, a failure on the part of a bulldozer to find fire at the location to which it was sent. A plan we currently have to fix the

error is to travel a little further in a specified direction, looking for the fire. A really intelligent error recovery knows when to try cheap fixes, such as modifying a destination; when to begin a search for a way to significantly modify a plan (for example, by dispatching another bulldozer); and when, as a last resort, to abandon the current plan and begin from scratch. Error recovery and replanning significantly depend on intelligent monitoring; in fact, envelopes are designed to predict errors before they happen, minimizing downtime.

We need to enhance the scheduling abilities of the cognitive component to make agents responsive to real-time demands in fire fighting. This need is particularly true for the fireboss in our implementation because it is essentially a cognitive agent. Currently, scheduling involves three actions: selecting an action to execute, deciding how much time is available, and selecting an execution method. Although these actions are charged for the time they use, they are not themselves scheduled, nor are there multiple execution methods to implement them. In short, the cognitive scheduler is a separate interpreter of the timeline. To make the scheduling of actions completely uniform, scheduling actions must themselves be scheduled. In addition, we must develop scheduling strategies, along the lines suggested in Lesser, Durfee, and Pavlin's (1988) approximate processing proposal.

To facilitate experiments with different agent designs in different environments, we have started to build a generic agent architecture. It is a collection of parameterizable structures that represent the design of parts of an agent. For example, our generic action structure includes pointers to execution methods, envelopes, and predicates that are tested before the action is selected. Generic execution methods, in turn, contain estimates of their time requirements, their prerequisites, and so on. We also have generic structures for sensors and effectors. In the near future, we will implement generic structures for strategies, including memory-access strategies and cognitive-scheduling strategies. The eventual goal is a full generic agent architecture that makes it easy to implement dif-

ferent agent designs by specifying how the agent manages its sensors, effectors, and memory, and decides what to do next.

We demonstrated one way to organize a multiagent planner in the Phoenix testbed, but the agent architecture certainly supports others. Work is under way in Victor Lesser's laboratory to build a cooperating, distributed planner for the Phoenix testbed. Although preliminary, this model assumes multiple firebosses, each with spheres of influence (geographic areas and agents) under its control, who cooperatively fight fires at their borders, loaning resources to neighbors or redrawing their boundaries to shift the work load in times of stress. Although this model is similar to the Phoenix planner in the relationship between firebosses and agents, it adds a cooperative relationship between firebosses.

Phoenix agents should learn to improve their performance. The opportunities for learning are myriad: We can learn new reflexes and chain reflexes together to learn short plan fragments. We can learn new plans from patches to failed ones. We can learn correlations between environmental conditions, such as changes in wind direction, and failures, such as bulldozers becoming trapped in the fire. Currently, we are extending the error-recovery mechanisms to learn patches to failed plans. This is one aspect of Adele Howe's (1989) dissertation work. Allen Newell recently pointed out that "you can't program SOAR" because much of its behavior emerges from sequences of locally selected chunks, and no way really exists to predict how a chunk, added by hand, makes the system behave. We have found the same to be true of actions and reflexes in Phoenix and concur with Newell that once a system attains a degree of complexity, it must learn to improve its performance on its own.

Conclusion

The development of Phoenix has been intimately tied to our evolving ideas about AI research methodology and, specifically, to our understanding of the role of evaluation in AI research (Cohen and Howe 1988, 1989). Clearly, the evaluation of Phoenix must be with respect to the goals of the project. Moreover, it must tell us not only whether we have succeeded but whether we are succeeding and why or why not. As noted in *The Phoenix Research Agenda*, the goals of Phoenix are of three kinds. Our technical goals are to build a real-time planner with learning, approximate scheduling, envelopes, and the other features noted earlier. Our scientific goal is to under-

stand how environmental characteristics influence agent design—the relationships discussed in the context of the behavioral ecology triangle (figure 5). Last, we are using Phoenix as a framework in which to develop AI methodology.

Progress toward each of these goals is evaluated differently. Phoenix is parameterized and instrumented at all its layers to facilitate evaluations of specific technical developments; for example, we can assess whether an approximate scheduling algorithm is robust against varying time pressure because we can vary time pressure while we hold other factors constant. We can run fire scenarios in dozens of conditions with dozens of variations in the algorithms used by the Phoenix planner. These experiments are scheduled to begin in the fall of 1989. They will enable us to demonstrate the utility of our technical solutions, explain why they are solutions, and discover the limits on their scope (Cohen and Howe 1988).

Clearly, these cannot be the only aims of the experiments. Although it is valuable to probe the scope and efficacy of specific techniques, such experiments might not necessarily address our scientific goals. We might show that a Phoenix planner works well in the Phoenix environment but not how the environment constrains the design of planners. Furthermore, unless we are trying to answer specific questions of this sort, experiments with techniques are unguided. Dozens of variations on the Phoenix planner exist as well as hundreds of environmental conditions in which they might be tested. To guide the search of this space, we will generate and test general rules that justify and explain the design of agents. These rules will call on functional relationships that capture trade-offs. For example, the well known store-recompute trade-off lurks in the design of the Phoenix planner: We use it to justify the decision to rely on stored plans in an environment that exerts time pressure, favoring storage over computation. Perhaps, there is a general rule here (for example, under time pressure, rely on storage over computation), or many specific variants of this rule for environments with different kinds of time pressures and agents with different kinds of store-recompute trade-offs. In any case, our scientific goal is to discover functional relationships (and to exploit those we already know such as the store-recompute trade-off) and embed them in rules for designing intelligent agents. To evaluate progress toward this goal, we need to measure the extent to which performance can be predicted. If we really understand the rela-

tionships between environmental characteristics, agents' behaviors, and agents' designs, then we should be able to predict that agents with particular designs will behave in particular ways under particular environmental conditions.

Although we are far from this goal, it is paradigmatic of the style of AI research we advocate. To evaluate the success of this methodological stance will take a long time, but if it is possible, surely there is no better aim for AI than to understand—to the point of being able to predict behavior—how to design intelligent agents in complex environments. ■

Acknowledgments

This research was supported by Defense Advanced Research Project Agency/Rome Air Development Center contract no. F30602-85-C0014; the Office of Naval Research (ONR), under a university research initiative grant, no. ONR N00014-86-K-0764; ONR under contract no. N00014-87-K-238; and a grant from Digital Equipment Corporation. We also wish to thank Paul Silvey.

References

- Agre, P. E., and Chapman, D. 1987. Pengi: An Implementation of a Theory of Activity. In Proceedings of the Sixth National Conference on Artificial Intelligence, 268-272. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Brooks, R. A. 1986. A Robust-Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation* RA-2(1) 14-23.
- Cohen, P. R. 1989. Why Knowledge Systems Research Is in Trouble and What We Can Do About It, Technical Report, 89-81, Experimental Knowledge Systems Laboratory, Dept. of Computer and Information Science, Univ. of Massachusetts.
- Cohen, P. R., and Howe, A. E. 1989. Toward AI Research Methodology: Three Case Studies in Evaluation. *IEEE Transactions on Systems, Man, and Cybernetics*. Forthcoming.
- Cohen, P. R., and Howe, A. E. 1988. How Evaluation Guides AI Research. *AI Magazine* 9(4): 35-43.
- Cohen, P. R.; Greenberg, M. L.; Hart, D.; and Howe, A. E. Forthcoming. An Introduction to Phoenix, the EKSL Fire-Fighting System, Technical Report, Experimental Knowledge Systems Laboratory, Dept. of Computer and Information Science, Univ. of Massachusetts.
- Howe, A. E. 1989. Adapting Planning to Complex Environments. Ph.D. diss., Dept. of Computer and Information Science, Univ. of Massachusetts.
- Lesser, V. R.; Durfee, E. H.; and Pavlin, J. 1988. Approximate Processing in Real-Time Problem Solving. *AI Magazine* 9(1): 49-61.

Paul R. Cohen is an associate professor and director of the Experimental Knowledge Systems Laboratory in the Department of Computer and Information Science, University of Massachusetts, Amherst. His research is concerned with how complex environments constrain the design of autonomous agents. Cohen is also an editor of the *Handbook of Artificial Intelligence, Volumes III and IV*.

Michael L. Greenberg is a graduate student in the Experimental Knowledge Systems Laboratory, Department of Computer and Information Science, University of Massachusetts, Amherst. He received a B.S. in electrical engineering from Carnegie-Mellon University in 1982 and an M.S. in computer science from the University of Massachusetts in 1986. He is currently working on his Ph.D. on designing planners for complex, real-time domains.

David M. Hart is manager of the Experimental Knowledge Systems Laboratory at the University of Massachusetts, Amherst. His research interests include knowledge-based systems, real-time dynamic planning, and control. Hart is a member of AAAI and ACM.

Adele E. Howe is a graduate student in the Experimental Knowledge Systems Laboratory, Department of Computer and Information Science, University of Massachusetts, Amherst. She received her B.S.E. in computer science and engineering from the University of Pennsylvania in 1983 and her M.S. in computer science from the University of Massachusetts in 1987. Her Ph.D. dissertation is on adapting planning to real-time, complex domains.