

The 1998 AI Planning Systems Competition

Drew McDermott

■ The 1998 Planning Competition at the AI Planning Systems Conference was the first of its kind. Its goal was to create planning domains that a wide variety of planning researchers could agree on to make comparison among planners more meaningful, measure overall progress in the field, and set up a framework for long-term creation of a repository of problems in a standard notation. A rules committee for the competition was created in 1997 and had long discussions on how the contest should go. One result of these discussions was the pddl notation for planning domains. This notation was used to set up a set of planning problems and get a modest problem repository started. As a result, five planning systems were able to compete when the contest took place in June 1998. All these systems solved problems in the strips framework, with some slight extensions. The attempt to find domains for other forms of planning foundered because of technical and organizational problems. In spite of this problem, the competition achieved its goals partially in that it confirmed that substantial progress had occurred in some subfields of planning, and it allowed qualitative comparison among different planning algorithms. It is urged that the competition continue to take place and to evolve.

In recent years, many subfields of AI have used competitions as a way of measuring progress and guiding research directions (Arkin 1998; MUC 1992, 1991). In a competition, researchers run their programs on a common set of problems at the same time with little tuning, and the results are compared. There are several purposes to such an exercise: It allows meaningful comparison of programs. It can provide an indication of overall progress in the field. It can provide a set of benchmark problems for others to use to compare their systems to the state of the art. It can focus attention on more realistic problems.

Of course, competitions have drawbacks. Preparing a program for a competition usually means polishing existing features and suspend-

ing work on new ones. "Realistic problems" might not be those of most interest in the long run. The Message-Understanding Competitions focused attention on information extraction from newspaper articles or more restricted media. The programs that did well on this task were those that were carefully engineered to detect and dissect messages in the target category. Programs derived from research programs with more ambitious goals, such as investigating the general theory of natural language understanding, did not work so well. One could argue that this outcome is evidence that such a general theory is, at least for now, a chimera. However, many people would disagree, arguing that in the long run, the discovery of such a theory is the whole point of studying natural language.

Until the AI Planning Systems (AIPS) Conference of 1998, there had never been a competition in the field of automated planning. The broadest definition of *planning* is reasoning about agent behavior. A system plans to the extent it predicts the consequences of alternative behaviors before selecting one. Because there are a wide range of agents, reasoning techniques, and ways of combining plan inference with plan execution, the field is broad, encompassing everything from factory scheduling to robot programming. Some of the application areas are of immediate practical interest, but others are still rather abstract.

In spite of this immediate practical interest, there haven't been many applications of planners that are actually used. The practical applications tend to lie in three categories: (1) scheduling problems, (2) plan-management problems, and (3) symbolic control problems. *Scheduling problems* are those in which the actions that must be taken are known, and the problem is to find an order in which to carry them out. *Plan-management problems* are those in which the plans tend to consist of stereotypical structures of actions, usually *hierarchical*

(longer-term actions being decomposed into structures of shorter-term ones), without many choices about how to decompose. *Symbolic control problems* are those in which the focus is on execution of plans for controlling a reactive system, where the plans were written by a human.

Applications in these areas tend to succeed because they eliminate or sharply constrain the amount of search the planning system has to do. The focus of the competition was intended to be planning problems that require a significant amount of search. In the past few years, there has been a lot of work on search-based planning algorithms and a fair amount of progress. It would be premature to say that practical algorithms exist, but planners are now finding solutions to problems that are an order of magnitude larger than those they could solve 10 years ago. It seemed as if the time was ripe to hold a competition and see how much progress had really been made and possibly push the community further in interesting directions.

Although some people (notably Manuela Veloso) had been arguing for a planning competition for years, serious talk about a competition began in 1996 and was a major topic among attendees at the Dagstuhl Workshop on Control of Search in Planning, held in October. The fourth biennial AI Planning Systems (AIPS) Conference was scheduled for June 1998, and that seemed like the obvious time to have a competition. Veloso was program cochair for the conference and asked me to chair the rules committee for the competition.

By the summer of 1997, we had assembled the following Rules Committee: Drew McDermott (chair), Yale University; Malik Ghallab, Ecole Nationale Supérieure D'ingénieur des Constructions Aéronautiques; Adele Howe, Colorado State University; Craig Knoblock, University of Southern California; Ashwin Ram, Georgia Institute of Technology; Manuela Veloso, Carnegie Mellon University; Daniel Weld, University of Washington; and David Wilkins, SRI International.

This article focuses on how the committee designed and ran the competition, with an overall summary of results. This issue of *AI Magazine* contains a companion article by those who actually competed, with Derek Long serving as overall author and editor. It focuses more on what strengths and weaknesses each planner embodied, with detailed comparison of their performance on various problems.

Several matters had to be dealt with to have a competition: First, exactly what sort of planning problems would we give people to work

on? As suggested earlier, even after eliminating scheduling and other areas, we are still left with several different types of problem area:

Classical planning: In this kind of problem, you are given an initial situation, a set of action definitions, and a proposition (goal) to be brought about. A solution is a sequence of actions that, when executed beginning in the initial situation, brings about a situation in which the goal is true. It is assumed that the planner knows everything that is true in the initial situation and knows the effect of every action.

Hierarchical planning: Here, you are given, in addition to the material of classical planning, a set of abstract actions. An abstract action cannot be executed directly but must be executed by executing an expansion (or reduction) of it in terms of less abstract actions, typically one found in a plan library. A problem can specify, in addition to a goal, an abstract action to be executed. A *solution* is a sequence of primitive actions that achieves the goal and corresponds to an expansion of the given abstract action.

Reactive planning: This classification is much vaguer, in which the assumption of perfect information is relaxed. There are many different types of reactive-planning problem, depending on what is assumed about the sensors and effectors available.

Learning in planning: This is not a problem type so much as an approach to the other types. A learning planner, not surprisingly, does better and better as it gains experience with planning problems of a given type. The style of learning usually studied is *case-based reasoning*, in which new problems are solved by adapting solutions to similar problems previously encountered.

The Rules Committee spent the summer of 1997 discussing how to proceed. It seemed clear that it would be hard to find one problem domain that would be attackable by planners in all these categories. Hence, it was decided to try to create several tracks, in which different categories of planner would compete. Regardless of how many tracks we settled on in the end, it was clear that we would need a notation to use as an input language for the competing programs; so, the committee set itself the goal of designing such a language, to be called the planning domain definition language, or PDDL.

Meanwhile, we engaged in serious discussion about (at least) four tracks: (1) a STRIPS track, classical planning with action-definition notation of the same expressive power as the STRIPS planner (Fikes and Nilsson 1971); (2) an ADL (action definition language) track, classical

planning with an enhanced notation allowing actions with quantified preconditions and context-dependent effects (Pednault 1989); (3) hierarchical planning, a classical domain but with explicitly given compound actions (canned plans) that the planner must be able to reason about (Erol, Nau, and Hendler 1994; Wilkins 1988); and (4) reactive planning, planning in a complex simulated domain, such as the PHOENIX fire-fighting simulator (Hart and Cohen 1992).

PDDL—The Planning Domain Definition Language

The PDDL language was designed to be a neutral specification of planning problems. *Neutral* means that it doesn't favor any particular planning system. The slogan we used to summarize this goal was "physics, not advice." That is, every piece of a representation would be a necessary part of the specification of what actions were possible and what their effects are. All traces of "hints" to a planning system would be eliminated. For example, an implication ($P \supset Q$) has several uses, including (1) to prove that Q is true in a world situation, adopt proving P as a subgoal, and (2) to achieve Q (make it true), adopt achieving P as a subgoal. A specification of which of these, if either, was actually a good idea would constitute advice. However, consider these two interpretations of $P \supset Q$: First, in every world situation, either P is false, or Q is true. Second, any plan that causes P to become true without Q being false is invalid no matter what other virtues it has. The second interpretation is a safety condition, which has a completely different meaning from the first. This distinction is one of "physics," although not in the traditional sense. In the first interpretation, $P \supset Q$ is automatically always satisfied. In the second interpretation, it might well be violated, but no legal plan is allowed to do so.

We anticipated that most planners would require some kind of advice because planners search large spaces and often fail without a bit of help. We didn't want to outlaw advice, just make sure it was properly labeled and accounted for, and omitted entirely from the core language.

It is harder than it sounds to create a notation that avoids advice entirely, for a variety of reasons. Historically, most planning researchers have not made the distinction, so existing input notations are full of it.

For example, there are actions that are useful only under certain disagreeable circumstances. If you need three nuts to attach a wheel to an

axle, one way to do it is to take one nut from each of the other wheels. This strategy is reasonable if all other nuts are lost. However, with this action in the database, some planners might consider deliberately throwing all the nuts away to make this action feasible. Although they would presumably reject this plan at some point, it might be useful to tell the planner that the condition "there are no unattached nuts" should be treated only as a *filter* condition, used to select among alternative actions but never to be achieved if false. Many planning researchers have developed sophisticated notations for expressing advice such as this and have made it so easy to include the advice in the action specifications that one is likely to overlook the fact that it's not actually part of the definition of what the action does.

The field in which advice has reigned supreme is *hierarchical planning*, which studies planners that assemble solutions from large canned plans stored in some kind of plan library. These plans typically look suspiciously like programs, and some of their steps are essentially procedures for setting up data structures, making sure the plan is appropriate, and so on. It is hard to extract the pure physics from a representation like that and just as hard to then represent the advice part as a separate set of hints.

Nonetheless, we felt it was important to try to define a purely physical hierarchical notation to support the proposed hierarchical-planning track for the competition. Given the way hierarchical planners are often used, this project might be viewed as of questionable sanity. The main reason to have a library of plans in the first place is to focus the planner's attention on these sequences of actions and away from others that are less likely to be useful. In other words, a plan library can be thought of as a hint library. If that's what it is, then specifying it is not part of the PDDL project. A team that wanted to use a plan library would have to treat it as an advice structure superimposed on an underlying physics specified without the use of hierarchy.

However, there is a sense in which a canned plan can be thought of as a purely physical entity, that is, when it represents a standard procedure whose functioning can't easily be expressed in more basic terms. A procedure for starting a nuclear reactor is, in principle, derivable from a more detailed specification of exactly what actions are possible and what their effects are, but this specification might not be available to the planner. It might just know that if it carries out the plan specified in the manual, then it will achieve certain things,

The field in which advice has reigned supreme is hierarchical planning, which studies planners that assemble solutions from large canned plans stored in some kind of plan library.

```

(define (domain logistics-adl)
  (:requirements :adl)
  (:types
    physobj - object
    obj vehicle - physobj
    truck airplane - vehicle
    location city - object
    airport - location)
  (:predicates
    (at ?x - physobj ?l - location)
    (in ?x - obj ?t - vehicle)
    (in-city ?l - location ?c - city)
    (loaded ?x - physobj)) ; ?x is loaded on a vehicle
  (:action load
    :parameters (?obj ?veh ?loc)
    :precondition (and
      (vehicle ?veh)
      (location ?loc)
      (at ?obj ?loc)
      (at ?veh ?loc)
      (not (loaded ?obj)))
    :effect (and
      (in ?obj ?veh)
      (loaded ?obj)))
  (:action unload
    :parameters (?obj ?veh ?loc)
    :precondition (and
      (vehicle ?veh)
      (location ?loc)
      (in ?obj ?veh)
      (at ?veh ?loc))
    :effect (and
      (not (in ?obj ?veh))
      (not (loaded ?obj))))
  (:action drive-truck
    :parameters (?truck - truck ?loc-from ?loc-to - location ?city - city)
    :precondition (and
      (at ?truck ?loc-from)
      (in-city ?loc-from ?city)
      (in-city ?loc-to ?city))
    :effect (and
      (at ?truck ?loc-to)
      (not (at ?truck ?loc-from))
      (forall (?x - obj)
        (when (and
          (in ?x ?truck)
          (and
            (not (at ?x ?loc-from))
            (at ?x ?loc-to)))))))
  (:action fly-airplane
    :parameters (?plane - airplane ?loc-from ?loc-to - airport)
    :precondition (and (at ?plane ?loc-from) )
    :effect (and
      (at ?plane ?loc-to)
      (not (at ?plane ?loc-from))
      (forall (?x - obj)
        (when (and
          (in ?x ?plane)
          (and
            (not (at ?x ?loc-from))
            (at ?x ?loc-to)))))))

```

Figure 1. The Logistics Domain—ADL Version.

and that it can interleave steps of other plans with the steps of this one provided certain constraints are honored.

From this point of view, it seemed reasonable to visualize the hierarchical-planning

track of the competition in these terms: For one or more domains, we would specify some canned plans and then stipulate that no problem solution could include any steps that were not part of an instance of one of these plans. In

concrete terms, if you had to pump out radioactive waste and restart the reactor, you weren't allowed to flick switch 203 unless one of the protocols for pumping and restarting called for this switch to be flicked at some point, and the other did not forbid it. It turned out that formalizing exactly the legal action sequences was extremely difficult, as I discuss later. The second most important desideratum in the design of PDDL was that it resemble existing input notations. Most input notations were Lisp-like for historical reasons, but beyond that, there were many divergences. The University of Washington UCPOP input language was the closest thing to a standard. However, it didn't have a simple notation for object types, used Lisp procedures for arithmetic operations, and didn't address the representation of hierarchical plans at all. PDDL was produced by adding to the UCPOP language a simple uniform typing syntax, some arithmetic facilities, and the notion of action expansion. Actions were classified as primitive or expandable. An expandable action could not be executed directly but had to be instantiated by selecting one of its methods, each a structure of actions.

Figure 1 gives an example of a simple domain, called *logistics*, one of those used in the competition. There were actually two versions, *logistics-ADL* and *logistics-STRIPS*. This version requires the use of ADL constructs such as typing and quantification. Types are indicated by hyphens. The type appears after the objects it qualifies. Variables are indicated by question marks.

However, some planners cannot handle types. In fact, for almost every aspect of problem definition, there is some planner that cannot handle it. To cope with this issue, we borrowed an idea from the UCPOP language, namely, to specify explicitly the requirements a planner would have to satisfy to handle this domain, which explains the field (:requirements :adl) that appears in the domain definition. If a planner can't handle the :adl package of requirements, then it can issue a warning when it sees this flag.

After the :requirements specification, the next field defines the :types that are specific to this domain. (Types such as object and integer are inherited by all domains.) Then there is a list of :predicates, each of which is given with its argument types.

There are four actions in this domain. Each is defined by giving a *precondition*, which must be true for the action to be feasible, and an *effect*, which specifies what happens when the action is executed. The effect is typically a conjunction. A conjunct of the form (not *p*) means

that *p* becomes false. A conjunct of the form (when *c e*) means that effect *e* happens only if condition *c* is true before the action occurs. (This is a context-dependent effect.) A conjunct of the form (forall (—vars—) *e*) means that *e* happens for every instance of the variables.

Thus, the fly-airplane action is defined as follows: You can fly an airplane from airport *A* to airport *B* if the airplane is at *A*; you can't fly any other vehicle between any other types of location. The effect of flying is that the airplane is at *B* and no longer at *A*. Furthermore, everything that is in the airplane is also at *B* and no longer at *A*. Note that we have to specify both that an object is at *B* and that it is no longer at *A*.

For comparison, figures 2 and 3 give the same domain with no extra requirements at all. We call this baseline notation the STRIPS notation because it is essentially the same as the notation used by the STRIPS planner (Fikes and Nilsson 1971). Here are the changes required to transform the ADL version into the STRIPS version:

First, types are replaced by unary predicates, which ultimately requires splitting the actions *load* and *unload* into two versions each, one dealing with trucks and the other with airplanes. Second, context-dependent effects ("whens") must be eliminated. In this case, the semantics of actions change. In the ADL version, an object is at the destination as soon as its vehicle moves there. In the STRIPS version, it is at the destination only when it is unloaded. I return to this issue of domain notation in The Contest.

Numbers are built in to all PDDL domains, but only those declaring requirement :expression-evaluation can have arithmetic expressions such as (+ ?x 1). These arithmetic expressions occur in special contexts such as (eval *e v*), where *e* is an arithmetic expression, and *v* must unify with its value. If *e* evaluates to a Boolean, then (test *e*) succeeds if and only if the value of *e* is true. (equation *e*₁ *e*₂) tries to bind variables in such a way as to make *e*₁ and *e*₂ evaluate to the same value. For example, the goal (equation (+ ?x 1) 3) can be satisfied by binding ?x to 2. Currently, that's about the only pattern that implementations are required to handle.

For example, in defining a grid world in which a robot can move between locations with integer-valued coordinates, we can use these facilities to specify what it means for two coordinates to be adjacent. Here is a piece of that specification:

```
(:axiom
  :vars (?i ?j ?i1 - integer)
  :implies (adjacent ?i ?j ?i1 ?j right)
  :context (and (equation (+ ?i 1) ?i1)
                (legal_coord ?i)
                (legal_coord ?i1)))
```

```

(define (domain logistics-strips)
  (:requirements :strips)
  (:predicates (obj ?obj)
               (truck ?truck)
               (location ?loc)
               (airplane ?airplane)
               (city ?city)
               (airport ?airport)
               (at ?obj ?loc)
               (in ?obj1 ?obj2)
               (in-city ?obj ?city))

  (:action load-truck
   :parameters (?obj ?truck ?loc)
   :precondition (and (obj ?obj)
                      (truck ?truck)
                      (location ?loc)
                      (at ?truck ?loc)
                      (at ?obj ?loc))
   :effect (and (not (at ?obj ?loc))
                (in ?obj ?truck)))

  (:action load-airplane
   :parameters (?obj ?airplane ?loc)
   :precondition (and (obj ?obj)
                      (airplane ?airplane)
                      (location ?loc)
                      (at ?obj ?loc)
                      (at ?airplane ?loc))
   :effect (and (not (at ?obj ?loc))
                (in ?obj ?airplane)))

  (:action unload-truck
   :parameters (?obj ?truck ?loc)
   :precondition (and (obj ?obj)
                      (truck ?truck)
                      (location ?loc)
                      (at ?truck ?loc)
                      (in ?obj ?truck))
   :effect (and (not (in ?obj ?truck))
                (at ?obj ?loc)))

  (:action unload-airplane
   :parameters (?obj ?airplane ?loc)
   :precondition (and (obj ?obj)
                      (airplane ?airplane)
                      (LOCATION ?loc)
                      (in ?obj ?airplane)
                      (at ?airplane ?loc))
   :effect (and (not (in ?obj ?airplane))
                (at ?obj ?loc)))

  ...

```

Figure 2. The Logistics Domain—STRIPS Version, Part 1

This example also illustrates PDDL's ability to represent axioms that delimit the meanings of symbols, such as *adjacent* and *right*. (There are three other axioms, for left, up, and down.)

There is also a notion of a term whose value changes in different situations, called *fluents*, following McCarthy and Hayes (1969). This

feature is especially useful in domains where quantities can change. The classical domain in which water can be poured from jug to jug might be defined as in figure 4. The (change *f e*) says that the value of fluent *f* changes to the value of *e* before the change. The *:functors* declaration is used to add new function-defining

```

...
(:action drive-truck
  :parameters (?truck ?loc-from ?loc-to ?city)
  :precondition (and
    (truck ?truck)
    (location ?loc-from)
    (location ?loc-to)
    (city ?city)
    (at ?truck ?loc-from)
    (in-city ?loc-from ?city)
    (in-city ?loc-to ?city))
  :effect (and
    (not (at ?truck ?loc-from))
    (at ?truck ?loc-to)))

(:action fly-airplane
  :parameters (?airplane ?loc-from ?loc-to)
  :precondition (and
    (airplane ?airplane)
    (airport ?loc-from)
    (airport ?loc-to)
    (at ?airplane ?loc-from))
  :effect (and
    (not (at ?airplane ?loc-from))
    (at ?airplane ?loc-to)))
)

```

Figure 3. The Logistics Domain—STRIPS Version, Part 2

symbols. Currently, it can be used only to define new fluent constructors; so, the type of amount is (fluent number), meaning that (amount x) is a number that varies from situation to situation.

Fluents are a natural generalization of traditional effects; instead of specifying how truth values change, they allow specification of how terms change. Without fluents, one could make the same definitions, but their meaning would be less clear. The precondition of *empty* would be

```

(and (amount-in ?jug1 ?a1)
     (amount-in ?jug2 ?a2)
     (capacity ?jug2 ?c2)
     (test >= (- ?c2 ?a2) ?a1))

```

and the effect would be

```

(and (not (amount-in ?jug1 ?a1))
     (amount-in ?jug1 0)
     (not (amount-in ?jug2 ?a2))
     (amount-in ?jug2 (+ ?a2 ?a1)))

```

Although this formulation causes little trouble in inferring the effects of a known action, it is difficult to use to constrain the arguments and preconditions of an empty action given a goal such as (and (amount-in jugB ?x) (> ?x 5)). The problem is that there may be many ways to make the first conjunct true, but the result is to leave some number of gallons in jugB. This number is either greater than five, or it isn't; there's no way to cause it to become bigger. Expressing the goal as (> (amount jugB) 5) is

much more perspicuous.

To support hierarchical planning, PDDL allows actions to be defined that are carried out by executing a structure of more primitive actions. Our nuclear-reactor example appears in figure 5. In English, "To restart reactor ? r , make sure it is not already running and not melting down, then open the two valves, toggle the switch, and close one of the valves, in that order." The *vars clause* is used to declare local variables that are inconvenient to consider as parameters of the action. Every reactor is supposed to have one auxiliary valve, one main valve, and one main switch; so, there is no need to name them as part of the action.

This definition is misleading in that it appears that there is only one way to expand an action. In general, it is possible to specify multiple methods for an action expansion, as in the fragment shown in figure 6 from a domain involving shipping packages from one place to another. Here there are two methods for carrying out (ship x l_1 l_2). The first, carrying it in a plane, works only if x is a piece of mail. The second, using a truck, works for any cargo item. The expression (in-context A :precondition p) means that p must be true before this occurrence of A for the plan to be valid.

In addition to defining domains, PDDL allows for the definition of problems. Figure 7 gives an example, used in the competition, in the logis-

```

(define (domain jug-pouring)
  (:requirements :typing :fluents)
  (:types jug)
  (:functors
    (amount ?j - jug)
    (capacity ?j - jug)
    - (fluent number))
  (:action empty
    :parameters (?jug1 ?jug2 - jug)
    :precondition (fluent-test
      (>= (- (capacity ?jug2) (amount ?jug2))
        (amount ?jug1)))
    :effect (and (change (amount ?jug1)
      0)
      (change (amount ?jug2)
        (+ (amount ?jug2)
          (amount ?jug1))))))
  ...)

```

Figure 4. The Jug-Pouring Domain.

```

(define (problem trans-1)
  (:action restart
    :parameters (?r - reactor)
    :vars (?valve1 ?valve2 - valve ?switch1 - switch)
    :precondition (and (not (running ?r))
      (not (melting-down ?r))
      (aux-valve ?r ?valve1)
      (main-valve ?r ?valve2)
      (main-switch ?r ?switch))
    :effect (running ?r)
    :expansion (series (verify-valves-shut ?r)
      (parallel (open ?valve1)
        (open ?valve2))
      (toggle ?switch)
      (close ?valve1)))

```

Figure 5. The Nuclear-Reactor Action.

tics domain. A problem is defined as a domain, a set of :objects, an initial situation, and a goal pattern to be made true. In domains with action expansions, a problem can have an :expansion field like the following:

```

(define (problem trans-1)
  (:domain transportation)
  (:init .....))
  (:goal (at truck3 detroit))
  (:expansion (ship pkg13 cincinnati)))

```

Here the planner must find a way to carry out the action (ship pkg13 cincinnati) in such a way

that (at truck3 detroit) is true when it is done.

Other features of PDDL include the following: Domains can include numeric parameters, such as the maximum coordinate in the grid world. PDDL allows the specification of timelessly true propositions, that is, facts that are present in all situations (thus saving having to enter them in all problem definitions). It allows one domain to be specified as a descendant of one or more alternative domains so that it inherits types, axioms, actions, and so on. It allows several problems to share an initial situ-

```

(:action ship
 :parameters (?pkg - cargo ?orig ?dest - location)
 :precondition (at ?pkg ?orig)
 :effect (at ?pkg ?dest)
 :expansion :methods)
(:method ship
 :parameters (?pkg - mail ?orig ?dest - location)
 :expansion (forsome (?p - airplane)
             (series (in-context
                     (load ?pkg ?p)
                     :precondition (at ?p ?orig)
                     (fly ?p ?dest)
                     (unload ?pkg ?p))))))
(:method ship
 :parameters (?pkg - cargo ?orig ?dest - location)
 :expansion (forsome (?tr - truck)
             (series (in-context
                     (load ?pkg ?tr)
                     :precondition (at ?tr ?orig)
                     (drive ?tr ?dest)
                     (unload ?pkg ?tr))))))

```

Figure 6. The Shipping Action.

ation, which needs to be written only once. One initial situation can be defined in terms of small changes to another. Action expansions can include simple iterations, arbitrary acyclic structures of actions, specification of conditions to be maintained true for some period during the plan, and more.

In addition to defining the language, we felt it was important to implement a syntax checker and a solution checker. The syntax checker could verify that domains submitted by others were valid PDDL and ensure that no feature was used unless it was declared as a requirement. It could also count the amount of advice that was given. To make advice accounting possible, we required all planner-specific annotations to be indicated by a special flag. The syntax checker could measure the size of these annotations and otherwise ignore them.

The other key piece of software was a solution checker. For the competition we wanted to be able to generate random problems. We anticipated not even knowing for many of the problems whether they had solutions or not. Some of the problems might have several solutions, some which might be quite long and involved. Having to check by hand if a solution was valid would be tedious and prone to error. We decided to automate it.

Creating a solution checker turned out to be harder than anticipated, for two reasons. The first is that to our knowledge, no one has ever

written a solution checker for a hierarchical planner, which you might find surprising. The reason is that most hierarchical planners do not treat prefabricated plans as part of the problem specification but as advice on how to solve problems. Once an action sequence has been found, the hierarchical superstructure can be dropped, and the action sequence can be checked as though it had been found without the use of canned plans. Checking an action sequence is easy: Just do a little deduction to verify that every action in the sequence is feasible at the point where it is to be executed and that the goal is true in the situation that results from executing the last action.

Now suppose you add a serious requirement that the action sequence not just be legal but also instantiate the :expansion given as part of the problem definition. The result is to superimpose a “parsing” problem on top of the standard deductive problem. That is, the solution checker must find a way to group the actions into a hierarchical structure to instantiate the given expansion, in a way analogous to the way a natural language parser groups words into phrases. However, the problem is much more difficult for several reasons. PDDL allows quantifiers in expansions of the form (forsome v A) and (foreach v C A). These occur in an action sequence if the right kind of instances of the action expansion A occur. In the case of a forsome, there must be one instance; in the

```

(define (problem log-x-2)
  (:domain logistics-adl)
  (:objects
    package5 package4 package3 package2 package1 - obj
    city10 city9 city8 city7 city6 city5 city4 city3 city2 city1 - city
    truck10 truck9 truck8 truck7 truck6 truck5 truck4 truck3 truck2 truck1 - truck
    plane4 plane3 plane2 plane1 - airplane
    city10-1 city9-1 city8-1 city7-1 city6-1 city5-1 city4-1 city3-1 city2-1 city1-1
  - location
    city10-2 city9-2 city8-2 city7-2 city6-2 city5-2 city4-2 city3-2 city2-2 city1-2
  - airport)
  (:init (in-city city10-2 city10) (in-city city10-1 city10)
    (in-city city9-2 city9) (in-city city9-1 city9)
    (in-city city8-2 city8) (in-city city8-1 city8)
    (in-city city7-2 city7) (in-city city7-1 city7)
    (in-city city6-2 city6) (in-city city6-1 city6)
    (in-city city5-2 city5) (in-city city5-1 city5)
    (in-city city4-2 city4) (in-city city4-1 city4)
    (in-city city3-2 city3) (in-city city3-1 city3)
    (in-city city2-2 city2) (in-city city2-1 city2)
    (in-city city1-2 city1) (in-city city1-1 city1)
    (at plane4 city3-2)
    (at plane3 city7-2) (at plane2 city3-2)
    (at plane1 city6-2) (at truck10 city10-1)
    (at truck9 city9-1) (at truck8 city8-1)
    (at truck7 city7-1) (at truck6 city6-1)
    (at truck5 city5-1) (at truck4 city4-1)
    (at truck3 city3-1) (at truck2 city2-1)
    (at truck1 city1-1) (at package5 city1-2)
    (at package4 city7-2) (at package3 city3-2)
    (at package2 city10-1) (at package1 city2-2))
  (:goal (and (at package5 city4-2)
    (at package4 city6-1)
    (at package3 city1-1)
    (at package2 city9-2)
    (at package1 city3-1))))

```

Figure 7. The Logistics Problem LOG-X-2.

case of a foreach, there must be a set of instances satisfying condition *C*. In addition, two action expansions could, unless constrained otherwise, be interleaved in an arbitrary order, and the same primitive action could occur as a part of more than one complex action.

It soon became clear that the problem of solution checking was going to be intractable unless the checker was given some hints. A solution to a problem with hierarchical expansions was going to have to include a specification of exactly which higher-level actions occurred where. Even with this change, the algorithm took a long time to develop, and it never was completely debugged.

Fortunately, or unfortunately, the lack of a

stable algorithm for checking solutions never became a problem because no contestants appeared for this part of the competition. We corresponded with several potential entrants, but none of them got over the hurdles in the way of taking part. The main problem was, we believe, that the semantics of hierarchical planning have never been clarified to the point where everyone in this area can be said to be working on the same problem. Our attempt to create a “lowest common denominator” notation succeeded only in creating a new notation that matched no one’s expectations. In addition, the hierarchical planning community is used to thinking of library plans as advice structures, which was a drastic departure from our assumption that the basic content of the plan

library contained no advice, only “physics.”

Trying to make this assumption actually work was extremely difficult. The problem is that no one has ever figured out how to reconcile the semantics of hierarchical plans with the semantics of primitive actions. Ordinary action sequences satisfy a straightforward compositionality property: If you know the preconditions and effects of two sequences of actions S_1 and S_2 , then you can infer the preconditions and effects of S_1 followed by S_2 . Hierarchical plans do not have this property, at least not obviously.

To take a simple example, consider the action restart, described in figure 5, for restarting a nuclear reactor. The action sequence $\langle \text{verify-valves-shut } r2 \rangle, \langle \text{open } v30 \rangle$ does not by itself restart the reactor, and neither does the action sequence $\langle \text{open } v29 \rangle, \langle \text{toggle } sw53 \rangle, \langle \text{close } v28 \rangle$, but the two together do, assuming that all the relevant preconditions are satisfied. In other words, a conditional effect of the second sequence is to restart the reactor, in a situation where the first sequence has (just? recently?) been executed. Conditional effects are not unusual in classical planning (Pednault 1989), but they normally take the form of an effect that becomes true if and only if a certain secondary precondition was true before the action. Expressing this sort of conditional effect is the job of the when effect clause in PDDL. With hierarchical plans, we get a new kind of implicit precondition, that a certain standard plan be “in progress.”

Suppose that an action sequence looks like $\langle \text{open } v29 \rangle, \langle \text{toggle } sw42 \rangle, \langle \text{close } v28 \rangle$. Is it legal? Does it cause (running $r2$) to become true if $v29$ is the auxiliary valve of $r2$ and $sw42$ is its main switch, and so on? One might think that the answer is “obviously not” because two actions are missing, namely (verify-valves-shut $r2$) and (open $v31$), assuming $v31$ is the main valve of $r2$. However, to make this inference requires one to assume that these actions did not occur before the action sequence we explicitly mentioned.

As it turned out, in the end, these complexities did not affect the actual competition. I describe them in such detail to save future researchers from rediscovering them.

The Contest

The competition took place in June 1998, but the contestants spent several months preparing for it. Each of them had to alter the front end of their system to accept problems expressed in PDDL. Because the language was brand new, this process was iterative: Changes to the notation

were suggested and sometimes incorporated before the problems were specified.

Even more important, the form of the contest had to be worked out, and several sample problems had to be released to give the contestants a clue about what their planners had to be capable of. A repository of problems was begun at Yale, using as a nucleus the repository developed by the UCPOP group at the University of Washington. Contestants were invited to submit new problems, and several did. In addition, some new domains were invented by the keeper of the repository, me.

Bargaining over Expressiveness

Over the six months leading to the actual competition, there was an intricate negotiation involving the committee and the community of potential contestants. The committee wanted to encourage the research community to try new things; the community wanted the committee to focus on the areas their planners did well in.

In the case of hierarchical planners, this tension proved fatal. The researchers with hierarchical planners lost interest rapidly as it became clear how great the distance was between PDDL and the kind of input their planners expected. Many of the researchers in this community think of their planners as a cross between a programming language and a knowledge-acquisition system. They have developed elaborate notations for capturing domain knowledge in the form of rules that push the planner toward particular kinds of solution. Unfortunately, PDDL defines all such notations as advice. To adapt these systems to PDDL would require factoring their input into two parts: (1) a physics part that is identical for all planners and (2) an advice part that controls how the planner reacts to the physics. The difficulty of doing this separation under the time constraints proved to be insurmountable. After a few exploratory conversations, all the hierarchical-planning researchers dropped out.

In the case of classical planning, the committee assumed at first that the ADL track was where most of the interest would be. ADL had been around since the mid-1980s, and several existing systems had been able to handle problems expressed in this format. However, much of the progress on planning algorithms in the 1990s has been based on what might be called *propositional planning*, in which variables are eliminated from planning problems by generating up front every instance of every term that might be needed. With the variables gone, all the machinery for matching literals and recording codesignation constraints is not

```

(:action move
  :parameters ( ?v - vehicle ?n1 ?n2 - node
                ?f1 ?f2 - pseudo-number)
  :precondition (and
                (loc ?v ?n1)
                (conn ?n1 ?n2)
                (fuel-at ?n1 ?f1)
                (just-less ?f2 ?f1))
  :effect (and
          (not (loc ?v ?n1))
          (loc ?v ?n2)
          (not (fuel-at ?n1 ?f1))
          (fuel-at ?n1 ?f2)))

```

Figure 8. Action Definition Using the Old Trick of Adding Arguments to the Action.

needed (Chapman 1987), and the search can focus on constraints among action terms and atomic formulas. The search process is simpler and can afford to explore a lot more possibilities. The resulting algorithms offer a significant improvement over older approaches in many cases.

Unfortunately, this power has been attained by sacrificing some expressivity. The propositional-planning researchers have focused on the STRIPS notation for the time being, and their planners lack the ability to handle problems involving numbers, nonatomic terms, and quantifiers. Some of them also have trouble with *context-dependent effects*, where the effects of an action depend on the circumstances in which it is executed. These limitations might sound serious, but in many cases, one can work around them, at the cost of using cumbersome notational tricks. For example, one of the domains we came up with for the competition was called the mystery domain. We called it that to conceal its underlying structure and make it harder to give planners advice about it. The domain actually concerned a transportation network through which vehicles could move carrying cargoes. A vehicle could move from one node to a neighboring node if there was fuel at the originating node. In PDDL,

```

(:action move
  :parameters (?v - vehicle ?n1 ?n2 - node)
  :precondition (and (loc ?v ?n1)
                    (conn ?n1 ?n2)
                    (fluent-test (> (fuel ?n1) 0)))
  :effect (and (not (loc ?v ?n1))
              (loc ?v ?n2)
              (change (fuel ?n1)
                     (- (fuel ?n1) 1))))

```

In English, you can move a vehicle from node 1 to node 2 if the vehicle is at node 1, node 1 is connected to node 2, and there is a nonzero amount of fuel at node 1. The effect of

the move is for the vehicle to be at node 2 and for there to be one less unit of fuel at node 1.

There are two problems with this action definition: It involves numbers, and it involves a context-dependent effect (the amount of fuel afterward depends on the amount of fuel before). In PDDL, a term like $(fuel\ n)$ defines a fluent, a term whose value changes from node to node. The notation $(change\ f\ a)$ means that the value of fluent f after the action is equal to the value of expression a before the action.

It might seem as if this domain were simply off limits to any planner that can't handle numbers or context-dependent effects, but in fact, there are ways to work around these limitations. Fuel amounts start off as nonnegative integers, never change except to become smaller by 1, and never become negative, so only a predictable set of natural numbers will occur in a given problem. Hence, for every problem, we can supply a set of constants $num_0, num_1, \dots, num_K$, where K is the largest number that occurs in the problem statement, and we can include among the initial conditions

```

(:init (just-less num0 num1)
       (just-less num1 num2)
       ...
       (just-less num9 num10)
       ...)

```

(in the case where $K = 10$). These constants are declared as *pseudonumbers*:

```

(:objects num0 num1...num10 - pseudo-number)

```

This eliminates the numbers; the next step is to eliminate the context-dependent effect. We use the old trick of adding arguments to the action. With this change, our action definition becomes as shown in figure 8. Suppose veh_{29} is at $node_{101}$, which has 3 units of fuel. Instead of saying, for example, "the action $(move\ veh_{29}\ node_{101}\ node_{63})$ changes the fuel at $node_{101}$ from 3 to 2," we say, "the action

```

(:action feast
  :parameters (?v ?n1 ?n2 ?f1 ?f2)
  :precondition (and
    (craves ?v ?n1)
    (food ?n1)
    (pleasure ?v)
    (eats ?n1 ?n2)
    (food ?n2)
    (locale ?n1 ?f2)
    (attacks ?f1 ?f2))
  :effect (and
    (not (craves ?v ?n1))
    (craves ?v ?n2)
    (not (locale ?n1 ?f2))
    (locale ?n1 ?f1)))

```

Figure 9. Action Definition in the STRIPS Style.

```

(:action feast
  :parameters (?v - pleasure ?n1 ?n2 - food)
  :vars (?f1 ?f2 - province)
  :precondition (and
    (craves ?v ?n1)
    (eats ?n1 ?n2)
    (locale ?n1 ?f2)
    (attacks ?f1 ?f2))
  :effect (and
    (not (craves ?v ?n1))
    (craves ?v ?n2)
    (not (locale ?n1 ?f2))
    (locale ?n1 ?f1)))

```

Figure 10. Action Definition in the ADL Style.

(move veh29 node101 node63 num3 num2) is the only feasible action of form (move veh29 node101 node63).”

It is somewhat discouraging that after 30 years of research, we are back to the notational restrictions we started with. However, it did have one benefit. We wanted to disguise the mystery domain, and all this verbosity helped do that. We labeled nodes as *foods*, vehicles as *pleasures*, and cargo objects as *emotions*. Instead of a single sort of pseudonumber, we introduced one, called *provinces*, for fuel and another, called *planets*, for space on vehicles. (loc v n) became (craves v n); (conn n_1 n_2) became (eats n_1 n_2); (fuel v k) became (local v k). The just-less predicate for numbers became (attacks k_1 k_2). Moving was called *feasting*. Thus, in the STRIPS style, our action definition becomes as shown in figure 9.

There was also an ADL version, albeit without numbers, in which types and context-dependent effects were allowed. In this version, the action was defined as shown in figure 10.

Participants

As a result of bargaining between the committee and the contestants, we arrived at a STRIPS track and an ADL track, neither of which could handle problems with numbers. By April 1998, we had two contestants who were planning to enter the ADL track, and eight who were planning to enter the STRIPS track. No one wanted to enter the hierarchical-planning track, and the other track never got off the ground. Nonetheless, we were happy with what we had. The contestants were putting a tremendous amount of work into altering their planners to take the PDDL notation. Unfortunately, for some of them, the work was just too much, and three dropped out in the weeks leading to the competition. The final participants in the STRIPS track were IPP (Jana Köhler, University of Freiburg); BLACKBOX (Henry Kautz and Bart Selman, ATT Labs and Cornell University); HSP (Hector Geffner and Blai Bonet, Simon Bolivar University); and STAN (Derek Long and Maria

Fox, Durham University). The two participants in the ADL track were Köhler's IPP and SGP (Corin Anderson, University of Washington). All the planners were written in C/C++ except SGP, which was written in Lisp.

As explained in the companion paper, all these systems except HSP were based to some extent on the GRAPHPLAN algorithm (Blum and Furst 1995). HSP was based on heuristic search guided by means-ends analysis. In addition, BLACKBOX used satisfiability testing. All the systems avoided repeated variable substitution by generating all required instances of propositions and action terms at the beginning. This lack of diversity in current research directions in classical planning means either that GRAPHPLAN on variable-free terms really is the best approach to planning or that the summer of 1998 happened to coincide with the peak of a particularly intense fad. What was particularly striking was the complete absence of partial-order, or nonlinear, planning (Weld 1994). A few years ago, many people thought that the superiority of partial-order techniques had been proven conclusively (Barrett and Weld 1994). It seems doubtful that the arguments in its favor were all wrong, and it would be interesting to see partial-order planners compete in future competitions.

Scoring

In parallel with the design of domains, we were also designing the scoring mechanism, which proved to be a difficult challenge, one that we never really solved. At first, we thought the biggest issue was going to be how to penalize a planning system for taking advice. Some members of the community feel that there is nothing wrong with advice; it was even suggested that a planner be rewarded for being able to take it. However, most people agreed that if planner *A* requires a lot more advice than planner *B* to solve a problem, *B* should win if it does almost as well as *A*.

Here is the scoring algorithm we proposed: The basic idea was to give each planner *j* a score on problem *i* equal to

$$(N_i - R_{ij}) W_i$$

where N_i is the number of planners competing on problem *i*; R_{ij} is the rank of planner *j* on problem *i* (0 for best program, $N - 1$ for worst, as explained later); and W_i is the difficulty of a problem, defined as

$$W_i = \text{median}_j T_{ij} / \sum_m \text{median}_n T_{mn}$$

where T_{bl} is the time taken by planner *l* on problem *b*.

Here is our method for computing R_{ij} : Rank all planners lexicographically as follows: (1) correctness, (2) advice, and (3) performance.

The most important dimension is correctness. There are two possible outcomes for planner *j* on problem *i* in order of decreasing possibility of winning: either it stops and reports a correct answer, or it doesn't. In other words, either it prints a correct solution or returns "no solution" when there isn't one, or it prints an incorrect solution, returns "no solution" when there is one, or never stops and has to be stopped by hand.

The second dimension is advice. Define

$$A_{i,j} = a_{i,j} + A_{D_i,j} / N(D_i)$$

where a_{ij} is the size of the advice given to planner *j* for problem *i*;

$$A_{D_j}$$

is the size of the advice given to planner *j* for domain *D* (D_i is the domain of problem *i*); and $N(D)$ is the number of problems in domain *D*. We planned to measure the size of a piece of advice by counting the number of symbols in it.

The third dimension is performance. If a problem has no solution, you measure the central processing unit time of planner *j* on problem *i*, or T_{ij} . If it has a solution, and planner *j* finds a solution, then we will replace T_{ij} with $T_{ij} (L_{ij})^h$, where L_{ij} is the length of the solution and $h = 0.5$. Length is defined as number of steps, regardless of whether some could be executed in parallel. (If $L = 0$, we treat the solution as of length 1.)

Three comments can be made. First, the idea was to take solution length into account but to discount it so that it broke a tie between planners only if they had comparable solution times. If planner P_1 is 10 times slower than planner P_2 , it would have to produce a plan 100 times shorter to win. If P_2 produces a plan twice as long, it must run in 70 percent of the time P_1 takes to beat it. The formula reflects the classical presupposition that the existence of a plan is more important than its size. (Making the exponent *h* bigger would make length more relevant.)

Second, if a planner required advice, it could never beat a planner that solves the same problems with no advice; so, it's worth giving a certain amount of advice only when you bet that no one will be able to solve the problem with less.

Third, for some of the more difficult machine-generated problems, we might not really know if there is a solution. In this case, if no planner finds a solution, we will assume that "no solution" is the correct answer. If a planner has to be stopped by hand, then it will be taken to have returned "no solution" after the amount of run time it actually spent (as

close as that can be estimated).

Unfortunately, this scoring function, in spite of its arcane complexity, failed to match everyone's judgment about what was to be measured, as I discuss later. Also, it turned out that our preoccupation with advice was misplaced. None of the competitors ever used any advice at all. A few years ago, almost every planner would have required a lot of advice, and it is remarkable how big a change had occurred.

Domains

In the last month before the competition, everyone involved put in a tremendous amount of work, making sure that every planner worked on the sample problems. Contestants were invited to contribute problem domains of their own. The idea was to allow each of them to benefit from their areas of strength by having at least one domain where they knew they would do well. Domains were submitted by the IPP group, the SGP group, and the BLACKBOX group. The final lineup of domains was as follows:

Mystery: The mystery domain described earlier defined three actions: (1) loading something on a vehicle, (2) unloading it, and (3) moving the vehicle.

Mystery prime: This is the mystery domain with one extra action, the ability to squirt a unit of fuel from any node to any other node, provided the originating node has at least two units. The contestants knew that a modified mystery domain was coming but did not actually see it until the first day of the competition.

Movie: In this domain, the goal is always the same (to have lots of snacks to watch a movie). There are seven actions, including rewind-movie and get-chips, but the number of constants increases with the problem number. Some planners have combinatorial problems in such cases. This domain was created by Corin Anderson.

Gripper: Here, a robot must move a set of balls from one room to another, being able to grip two balls at a time, one in each gripper. There are three actions: (1) move, (2) pick, and (3) drop. Most planners explore all possible combinations of balls in grippers, overlooking the fact that all combinations are equivalent and giving rise to an unnecessary combinatorial explosion. (Contributed by Jana Köhler.)

Logistics: There are several cities, each containing several locations, some of which are airports. There are also trucks, which can drive within a single city, and airplanes, which can fly between airports. The goal is to get some packages from various locations to various new locations. (Created by Bart Selman and Henry

Kautz, based on an earlier domain by Manuela Veloso.) Figure 1 gives the complete ADL version of the logistics domain, developed by me from Selman and Kautz's STRIPS version. (This is not precisely the version of the domain used in the competition; this version had an unimportant bug that has been removed.) The STRIPS version has six action definitions instead of four because an action with context-dependent effects has to be split into different versions.

Grid: There is a square grid of locations. A robot can move one grid square at a time horizontally and vertically. If a square is locked, the robot can move to it only by unlocking it, which requires having a key of the same shape as the lock. The keys must be fetched and can themselves be in locked locations. Only one object can be carried at a time. The goal is to get objects from various locations to various new locations. The ADL version of the domain has four actions, and the STRIPS version has five. (Created by Jana Kohler, based on an earlier domain of mine.)

Assembly: The goal is to assemble a complex object made of subassemblies. There are four actions: (1) commit resource assembly, (2) release resource assembly, (3) assemble part assembly, and (4) remove part assembly. The sequence of steps must obey a given partial order. In addition, through poor engineering design, many subassemblies must be installed temporarily in one assembly, then removed and given a permanent home in another. There was no STRIPS version of this domain.

It would be pleasant if we could claim that these domains covered the entire range of what planners can handle, that these domains represent approximations of real-world problems planners will eventually solve, or that within each domain the problems are typical. Unfortunately, we can make none of these claims. Two of the domains, movie and gripper, were submitted because problems in these domains were thought to be difficult for some planners to solve, even though the problems are easy for humans. The other domains were chosen because it seemed, based on experience and informal experimentation, that it was possible to create hard problems in them. However, as is now well known (Kirkpatrick and Selman 1994), it can be tricky to generate random problems that are hard. Randomly generated problems tend to be either extremely easy or impossible. (Some "impossible" problems are actually quite easy because many programs can quickly verify that they are impossible.) The zone in between the subspace of easy problems and the subspace of impossible problems has been compared to a phase transition in a phys-

Domain	Number	Average Objects	Average Inits	Smallest	Largest
Assembly	30	48	118	67	270
Gripper (ADL)	20	25	26	13	89
Gripper (STRIPS)		27	53	23	137
Logistics (ADL)	30	171	155	58	960
Logistics (STRIPS)		171	342	96	1470
Movie (ADL)	30	98	3	28	173
Movie (STRIPS)		98	99	51	341
Mystery (ADL)	30	44	82	46	233
Mystery (STRIPS)		44	126	64	317

Table 1. Problem Sizes—Round 1.

For each domain, the number of problems, the average number of objects a problem, and the average number of “inits” (propositions true in the initial situation) are shown. The columns labeled *Smallest* and *Largest* give the combined size (objects + inits) for the smallest problem in the domain and the largest. The mystery prime domain had exactly the same problems as the mystery domain, so they are not listed separately here.

ical system. Analyzing a domain to figure out where the phase transitions are is not easy, and we did not attempt it for any of the domains in the competition. As a result, some of the randomly generated problems are too easy, and some are too hard. However, it does seem that many are about right.

Results

The competition took place at Carnegie Mellon University (CMU) in June 1998, at the same time as AIPS-98. We owe a debt of gratitude to the CMU staff, especially Bob McDivett, who got the computers running and made sure that they were all identical. (The computers were 233-megahertz PENTIUM-based PC compatibles, with 128 megabytes of primary memory, running the LINUX operating system.) The contestants and I arrived early at the conference to get their systems up and running. The next few days were an intense but exhilarating effort. In the end, we had to write quite a bit more code and rethink our scoring function completely.

There were to be two rounds in the competition. The first was designed to allow contestants to get used to the environment and the problem domains. They would be allowed to run their programs several times and make changes in between. The programs that did best in round 1 would be allowed to advance to round 2, where the rules were more stringent. Some new domains would be introduced. Each planner could be run exactly once, with no tuning. Round 2 would take place in real time, as the conference proceeded, with programs’

performance announced as they finished.

Everything went reasonably smoothly through round 1. We had a total of 170 problems, drawn from the assembly, gripper, logistics, movie, and mystery domains. All 170 appeared in the ADL track; in the STRIPS track, the assembly problems were omitted, leaving 140. Table 1 summarizes the data on problem sizes. The grid domain was reserved for round 2. Contestants worked through Monday, June 8, at 5 PM, when we declared round 1 complete. On the ADL track, IPP outperformed SGP so convincingly that it was declared the winner. Both programs did well, but SGP was written in Lisp and rarely matched the raw speed of the other systems, which were written in C or C++.

The results for the STRIPS track were not at all clear. For one thing, we failed to anticipate that several of the contestants would simply not try to solve some of the problems. If their planner failed on almost all the easiest 10 problems in a domain, they didn’t see the point of letting it grind forever on the next 20. The scoring function as originally designed gave one point to a program that tried a problem, failed, and took longer than any other program that tried and failed. It gave zero to a program that didn’t try.

An even worse problem was that one planner (STAN) spent an hour each on the more difficult gripper problems before giving up on them, which was much longer than any other planner spent on any problem. Only two planners, STAN and HSP, tried the difficult gripper problems, so the median time to solve them was large, and these gripper problems ended up carrying a large fraction of the weight. HSP

Domain	Number	Average Objects	Average Inits	Smallest	Largest
Logistics (STRIPS)	5	33	67	63	159
Mystery Prime (STRIPS)	5	40	106	82	237
Grid (STRIPS)	5	66	328	209	613

Table 2. Problem Sizes—Round 2.

spent less time than STAN and actually solved the problems (although not optimally), so it got a higher overall score on round 1 than anything else. The HSP team deserves credit for solving these problems, but it seems clear that the scoring function's judgment of their importance disagrees with intuition. At the end of round 1, therefore, all four contestants in the STRIPS track could argue that their systems had done well. A total of 88 problems had been solved by at least one planner.

HSP solved more problems than any other system and found the shortest solution more often. BLACKBOX had the shortest average time on problems it attempted, but IPP had the shortest solution time on more problems. STAN was in second place for shortest solution time and second place for overall score.

The committee was unhappy with the holes that had been revealed in the scoring function. We tried to achieve a consensus on what to replace it with and finally gave up. We decided to let all the programs advance to round 2; measure their performance as well as possible; and let history judge who, if anyone, did the best. Once this decision was made, round 2 was a lot of fun. We continued to observe the same pattern as in round 1, that different planners succeeded in different ways.

For round 2, we used the grid, logistics, and mystery prime domains, all in their STRIPS versions. There were a total of 15 problems, of which 12 were solved by at least one program. Table 2 shows the sizes of these problems. We tried to generate problems that the planning systems could be expected to handle in the time allotted.

Tables 3 and 4 summarize how hard the problems were that some planner could solve. For each domain, we show the largest problem that any planner could solve and the problem whose shortest known solution is longer than that of any other problem. These figures should not be taken too seriously. For one thing, the

fact that a problem was not solved by any planner might mean that it has no solution (more later). In addition, the fact that no planner finds a short solution to a problem does not mean that there isn't one. To give a concrete idea of the performance of the planners, figure 7 contains the definition of problem STRIPS-LOG-X-2, which occurred during round 1 of the competition. It mentions 49 objects and 68 inits, for a total size of 147. Both STAN and BLACKBOX found 32-step plans to solve this problem; HSP found a 44-step plan. BLACKBOX's plan is shown in figure 11.

Tables 5, 6, and 7 give the results for both rounds. The planners are sorted in alphabetical order in each table. These results are not exactly the same as those presented at the conference because of some minor glitches that muddied the presentation. Three of IPP's solutions were checked and found to be wrong because of a trivial bug in the output printer, which caused all occurrences of one particular action to be garbled. In these tables, we have counted these as successes.

There are two important caveats about these data: First, we measure the length of the plan found by counting the total number of steps in it. However, for many of the planners, this measure might not be appropriate. Planners such as BLACKBOX, IPP, and STAN find the plan with the shortest *parallel length*, in which several steps are counted as taking one time unit if they occur as a substring at some point in the plan, and the substring could have occurred in any order. The plan with the shortest parallel length might not be the plan with the least number of steps. Which of these numbers is a better measure of plan quality is not always obvious.

If no planner found a solution to a problem, it simply doesn't enter into our statistics. However, in some cases, some of the planners were able to prove there was no solution. I discuss this further later. For further analysis of results, see the companion paper.

Domain	Largest Solved	Plan Length	Solution Time	Longest Solution	Solution Time
Gripper (ADL)	29	47	225730	47	225730
Gripper (STRIPS)	137	165	33210	165	33210
Logistics (ADL)	109	26	17400	26	17400
Logistics (STRIPS)	180	112	788914	112	788914
Movie (ADL)	173	7	50	—	—
Movie (STRIPS)	341	7	40	—	—
Mystery (ADL)	159	4	17810	13	9280
Mystery (STRIPS)	304	16	1789	16	789
Mystery Prime (ADL)	131	10	24240	12	1960
Mystery Prime (STRIPS)	214	4	7141	11	5214

Table 3. Hardest Problems—Round 1.

For each domain, data are shown for the hardest problems solved. On the left side of the table, data are displayed for the largest problem solved. Call it *B*. *Largest solved* is the size of *B*, defined as the sum objects + inits. *Plan length* is the length of the shortest solution of *B* found, and *solution time* is the time the fastest planner took to find the solution. On the right side of the table, data are displayed for the problem whose shortest solution was longest. Call it *L*. *Longest solution* gives the length of the shortest solution of *L* found, and *solution time* is the time the fastest planner took to find the solution. Times are in milliseconds.

Domain	Largest Solved	Plan Length	Solution Time	Longest Solution	Solution Time
Logistics (STRIPS)	159	31	66170	59	170394
Mystery Prime (STRIPS)	237	6	4991	7	2537
Grid (STRIPS)	209	14	2505	14	2505

Table 4. Hardest Problems—Round 2.

For each domain, data are shown for the hardest problems solved. On the left side of the table, data are displayed for the largest problem solved. Call it *B*. *Largest solved* is the size of *B*, defined as the sum objects + inits. *Plan length* is the length of the shortest solution of *B* found, and *solution time* is the time the fastest planner took to find the solution. On the right side of the table, data are displayed for the problem whose shortest solution was longest. Call it *L*. *Longest solution* gives the length of the shortest solution of *L* found, and *solution time* is the time the fastest planner took to find the solution. Times are in milliseconds.

Conclusions

The planning competition was a valuable exercise. For the first time, researchers had to compete against each other on exactly the same problems. The PDDL notation made this possible, and hopefully, it will continue to serve this role. The PDDL syntax checker and solution checker, as well as all the problems and results from the competition, can be found at www.cs.yale.edu/~dvm.

We encourage researchers to compare their planning systems against the programs that competed. However, the existence of this repository is only a first step toward a comprehensive set of benchmark problems for automated planners. We encourage others to submit candidate benchmarks to drew.mcdermott@yale.edu or to the competition committee for the 2000 Planning Competition, chaired by Fahiem Bacchus.

The competition documented a dramatic increase in the speed of planning algorithms. Some of the problems in the competition had solutions of 30 steps or more, extracted from problems with dozens of propositions to deal with. Ten years ago planners required significant amounts of domain-specific advice to achieve performance like this; the current generation requires no advice at all.

To an extent, this gain has been won by restricting the types of problem that can be worked on. Most of the planners could handle STRIPS-style problems and not much else. Some of these restrictions are only temporary, and we would urge the planning community to explore ways of removing them. However, the focus on classical planning, where perfect information is assumed, seems to be an intrinsic constraint on most planners currently being developed. There's nothing wrong with this focus, but if the current research is really leading to powerful algorithms, it will soon be time to show them working on realistic classical-planning problems. For example, combining features of the logistics and mystery domains would get us close to real-world transportation planning with capacity constraints. Perhaps this is a reasonable target for the community to aim for.

There is a remarkable divergence of opinion on whether and how planners should take advice', that is, domain- or problem-specific guidelines that are not strictly necessary in defining the domain but constrain the search for plans. Some researchers feel that any need for advice is a weakness, and others think that there is unlikely to be a general-purpose planning algorithm that solves all realistic problems. Thus, the issue is not whether domain-

```
((load-airplane package4 plane3 city7-2)
(load-truck package2 truck10 city10-1)
(load-airplane package3 plane2 city3-2)
(drive-truck truck1 city1-1 city1-2 city1)
(fly-airplane plane4 city3-2 city2-2)
(fly-airplane plane1 city6-2 city10-2)
(drive-truck truck6 city6-1 city6-2 city6)
(drive-truck truck3 city3-1 city3-2 city3)
(load-airplane package1 plane4 city2-2)
(drive-truck truck10 city10-1 city10-2 city10)
(fly-airplane plane3 city7-2 city6-2)
(fly-airplane plane2 city3-2 city1-2)
(unload-airplane package4 plane3 city6-2)
(unload-truck package2 truck10 city10-2)
(fly-airplane plane4 city2-2 city3-2)
(unload-airplane package3 plane2 city1-2)
(load-airplane package5 plane2 city1-2)
(fly-airplane plane2 city1-2 city4-2)
(load-airplane package2 plane1 city10-2)
(unload-airplane package1 plane4 city3-2)
(load-truck package3 truck1 city1-2)
(load-truck package4 truck6 city6-2)
(drive-truck truck6 city6-2 city6-1 city6)
(unload-airplane package5 plane2 city4-2)
(fly-airplane plane1 city10-2 city9-2)
(drive-truck truck1 city1-2 city1-1 city1)
(load-truck package1 truck3 city3-2)
(unload-airplane package2 plane1 city9-2)
(drive-truck truck3 city3-2 city3-1 city3)
(unload-truck package3 truck1 city1-1)
(unload-truck package4 truck6 city6-1)
(unload-truck package1 truck3 city3-1))
```

Figure 11. BLACKBOX's Solution to the Problem of Figure 7.

specific heuristics are necessary but how easy or natural it is to tell a planner about them. In this competition, the first camp decided the issue, probably because I am in this camp. However, none of the competitors could resist the temptation to ask for a "little bit" of "trivial" advice. For example, some of the propositional planners must set a bound on the length of a plan before doing a search. If the bound is unknown, they must search for it by starting with a short bound and extending it after each failure. A serious argument was made that this number should be given to the planner in advance. The antiadvice people succeeded in arguing that this one number constituted an enormous hint. (Among other things, it tells the planner that a solution exists.) In other cases, the disagreement became an impasse. For

Planner	Average Time	Problems Solved	Fastest	Shortest
IPP	21396	69	68	68
SGP	14343	38	5	35

Table 5. Results for Round 1—ADL Track.

Planner	Average Time	Problems Solved	Fastest	Shortest
BLACKBOX	1498	63	16	55
HSP	35483	82	19	61
IPP	7408	63	29	49
STAN	55413	64	24	47

Table 6. Results for Round 1—STRIPS Track.

A total of 88 problems were solved by at least 1 planner. Times are in milliseconds.

Planner	Average Time	Problems Solved	Fastest	Shortest
BLACKBOX	2464	8	3	6
HSP	25875	9	1	5
IPP	17375	11	3	8
STAN	1334	7	5	4

Table 7. Results for Round 2—STRIPS Track.

A total of 12 problems were solved by at least 1 planner. Times are in milliseconds.

example, many hierarchical planners (especially SIPE [WILKINS 1988]) give deductive rules a procedural interpretation, so that they are used in only one direction. $P \supset Q$ is interpreted to mean that Q is caused by P , which then gets generalized to “execute Q whenever P becomes true.” The use of rules of this sort can be a very directed way of getting a planner to do something, and once you get used to thinking in these terms, it’s hard to recast problems as pure physics plus some extra advice.

We hope the competition will become a regular part of the AIPS conference, thereby continuing to exert pressure on the planning community. In particular, we hope the next competition deals with the following issues:

There should be competitions

involving hierarchical planning but also run-time (reactive) planning and decision-theoretic planning. The issue of plan quality should be addressed more carefully. Different problem domains have different definitions of optimality, and these definitions should be made explicit in the domain or problem definitions. Plan length (sequential or parallel) is only a crude measure of plan cost; more realistic measures are needed. In many domains, the problem of finding an optimal plan is much more difficult than the problem of finding a feasible one, and in these cases, it might be desirable to give planners a big bonus for coming close to the optimum.

Some planners can prove that a problem has no solution; others just work for a while and then declare that

they can’t solve it. The former should be rewarded. The problem is to verify that a planner’s output is correct. In the case where it outputs a proposed solution as a sequence of steps, this sequence can be simulated to see if it is feasible and actually brings about the desired situation. The issue is what it should output in the case where it proves that there is no solution. One possibility is that it can produce a formal proof in first-order logic, in which case, the solution checker would just be a proof checker. However, most systems that prove there is no solution do not currently produce such a proof, and it would probably be a major pain to give them the ability to produce one.

Many of the problems in the competition were produced by random problem generators. It is difficult to ensure that randomly generated problems are interesting, that is, not easy or impossible to solve. One way to get around this difficulty is to make problem design a more important part of the competition and encourage participants (and others) to produce tricky problems.

A more strenuous effort should be made to accommodate planners that require domain-specific advice. Ideally, a reward mechanism should be found that gives points for ease of advising, and more creative ways should be sought of combining PDDL with planner-specific advice.

Acknowledgments

Thanks to the members of the committee and the contestants for making this event possible. Thanks to the CMU staff for technical support. Thanks to Blai Bonet, Henry Kautz, Manuela Veloso, David Wilkins, and the referees for comments on this article. Funding for the competition was provided by the Defense Advanced Research Projects Agency.

References

- Arkin, R. C. 1998. The 1997 AAI Mobile Robot Competition and Exhibition. *AI Magazine* 19(3): 13–17.
- Barrett, A., and Weld, D. S. 1994. Partial-Order Planning: Evaluating Possible Efficiency Gains. *Artificial Intelligence* 67(1): 71–112.
- Blum, A. L., and Furst, M. L. 1995. Fast Planning through Planning Graph Analy-

Running out of space on your bookshelf? Tired of leafing through thousands of pages for that one paper?

Electronic Versions of *ALL* AAAI Proceedings Are Now Available!

■ 1980–1997 AAI–IAAI Proceedings CD-ROM Set

This archival CD-ROM set consists of all the papers from the 1980, 1982, 1983, 1984, 1986, 1987, 1988, 1990, 1991, 1992, 1993, 1994, 1996, and 1997 AAAI National Conferences as well as the papers from the 1993, 1994, 1995, 1996, and 1997 Innovative Applications of Artificial Intelligence Conferences. The set consists of more than 2,200 papers in PDF format and is readable by most Windows, Macintosh, and many UNIX platforms.

4 CDs, \$250.00

■ 1998–1999 AAAI–IAAI Proceedings CD-ROM (Special Bonus: KDD–95, KDD96, KDD97, and KDD-98)

This CD-ROM set consists of all the papers from the 1998 and 1999 National Conferences on Artificial Intelligence as well as the 1998 and 1999 Innovative Applications of Artificial Intelligence Conferences. Also included on this CD are the 1995, 1996, 1997, and 1998 Knowledge Discovery and Data Mining Conference papers. The single CD contains approximately 260 technical papers in PDF format and is readable by most Windows, Macintosh, and many UNIX platforms.

1 CD, \$50.00

To order, call 650-328-3123 or send e-mail to orders@aaai.org

prices subject to change without notice

AAAI PRESS 445 BURGESS DRIVE MENLO PARK, CALIFORNIA 94025 USA

sis. In Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, 1636–1642. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Chapman, D. 1987. Planning for Conjunctive Goals. *Artificial Intelligence* 32(3): 333–377.

Erol, K.; Nau, D.; and Hendler, J. 1994. HTN Planning: Complexity and Expressivity. In Proceedings of the Twelfth National Conference on Artificial Intelligence. Menlo Park, Calif.: American Association for Artificial Intelligence.

Fikes, R., and Nilsson, N. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence* 2:189–208.

Hart, D. M., and Cohen, P. R. 1992. Predicting and Explaining Success and Task Duration in the PHOENIX Planner. In *Proceedings of the First International Conference on AI Planning Systems*, 106–115. San Francisco, Calif.: Morgan Kaufmann.

Kirkpatrick, S., and Selman, B. 1969. Criti-

cal Behavior in the Satisfiability of Random Boolean Expressions. *Science* 264(5163):1297–1301.

McCarthy, J., and Hayes, P. 1969. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In *Machine Intelligence, Volume 4*, eds. B. Meltzer and D. Michie, 463–502. Edinburgh, U.K.: Edinburgh University Press.

MUC. 1992. *Proceedings of the Fourth Message-Understanding Conference*. San Francisco, Calif.: Morgan Kaufmann.

MUC. 1991. *Proceedings of the Third Message-Understanding Conference*. San Francisco, Calif.: Morgan Kaufmann.

Pednault, E. P. 1989. ADL: Exploring the Middle Ground between STRIPS and the Situation Calculus. In *Proceedings of the Conference on Knowledge Representation and Reasoning, Volume 1*, 324–332. San Francisco, Calif.: Morgan Kaufmann.

Weld, D. 1994. An Introduction to Least Commitment Planning. *AI Magazine* 15(4): 27–61.

Wilkins, D. 1988. *Practical Planning: Extend-*

ing the Classical AI Planning Paradigm. San Francisco, Calif.: Morgan Kaufmann.



Drew McDermott is a professor of computer science at Yale University. He was educated at MIT, where he received a Ph.D. in 1976. He was chair of the Yale Computer Science Department from 1991 to

1996. He is coauthor of two textbooks in AI, serves on the editorial board of *Artificial Intelligence*, and is also a Fellow of the American Association for Artificial Intelligence. McDermott's research is in robot navigation, planning, and interagent communications. In earlier incarnations he did work on knowledge representation and nonmonotonic logic. He is currently working on a book about machine consciousness.

