Solving Time-Dependent Planning Problems

Mark Boddy Thomas Dean
Brown University Department of Computer Science
Providence, RI 02912, U.S.A.

Technical Report No. CS-89-03
February 1989

Solving Time-Dependent Planning Problems

Mark Boddy Thomas Dean*
Department of Computer Science
Brown University, Box 1910
Providence, RI 02912

Abstract

We are interested in constructing solutions to *time-dependent planning problems*: problems in which time spent planning affects the utility of the system's performance. In [4], we define a framework for constructing solutions to time-dependent planning problems, called *expectation-driven iterative refinement*. In this paper, we analyze and solve a moderately complex time-dependent planning problem involving path planning for a mobile robot, as a way of exploring a methodology for applying expectation-driven iterative refinement. The fact that we construct a solution to the proposed problem without appealing to luck or extraordinary inspiration provides evidence that expectation-driven iterative refinement is an appropriate framework for solving time-dependent planning problems.

*This work was supported by an IBM graduate fellowship, by the National Science Foundation under grant IRI-8612644, and by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Air Force Office of Scientific Research under Contract No. F49620-88-C-0132.

1 Introduction

In [4], we suggest viewing planning as a decision-making process driven by predictions or expectations regarding the future occurrence of events. Predicted events may provide constraints on the amount of time available or otherwise affect the cost of time spent planning. We say that a planning problem is *time-dependent* if the time spent planning has a cost (*i.e.*, affects the utility of the system's performance). In [4], we define a framework for constructing solutions to time-dependent planning problems that we call *expectation-driven iterative refinement.* This framework is intended to restrict the form of the solutions generated, in order to make their analysis easier. Planning using expectation-driven iterative refinement is accomplished using a set of decision procedures known as *anytime algorithms.* Anytime algorithms are defined as algorithms that return some answer for any allocation of computation time, and are expected to return better answers when given more time. A solution to a time-dependent planning problem using expectation-driven iterative refinement will consist of a set of anytime algorithms for planning, and a *deliberation-scheduling* algorithm that allocates computational resources to the set of anytime algorithms based on expectations regarding their performance.

In this paper, we explore a methodology for applying expectation-driven iterative refinement to time-dependent planning problems. As an example of what such a methodology should consist of, consider the use of divide-and-conquer as a strategy for solving combinatorial problems. Divide-and-conquer is a useful strategy for problems that can be decomposed, such that the resulting subproblems have substantially easier solutions. The literature offers a great deal of advice on how to go about decomposing a particular problem (*e.g.*, keeping the size of the subproblems approximately equal [1]). Problems that are not completely decomposable may have approximate divide-and-conquer solutions based on the introduction of appropriate assumptions (*e.g.*, Karp's partition for TSP [9]). Analogously, a methodology for expectation-driven iterative refinement will require some way of identifying planning problems as likely candidates for expectation-driven iterative refinement, and a set of methods for generating solutions.

We proceed with the development of a methodology for expectation-driven iterative refinement by applying it to a moderately complex robot planning problem.

The process of constructing a solution provides some insight into what a methodology for expectation-driven iterative refinement will look like. It also helps validate our claim that the framework itself is useful in that we construct a useful solution without appealing to luck or extraordinary inspiration. The problem we solve must be sufficiently complex to demonstrate a nontrivial application of expectation-driven iterative refinement. At the same time, the solution must be presented and analyzed in a reasonable amount of space. Additionally, we need a problem for which we can gather empirical data in order to construct the functions describing the expected value of the answers returned by our anytime algorithms. As a result of these constraints, the problem we solve is an abstraction from a "real-world" problem, in that the robot's environment is a simulation (and simplification) of the physical world.

In the next section, we present a robot planning problem, and discuss how to break the problem down into pieces suitable for implementation as anytime algorithms. The rest of the paper describes our solution in detail.

2 The robot courier

Suppose that we are in charge of designing the control program for a new model of robot courier for a delivery service in Manhattan. The function of these couriers is to pick up small parcels and deliver them to specified locations. Assuming that a courier completes all of its assigned deliveries, the main factor in evaluating the courier's performance will be the speed with which it accomplishes this task. The faster deliveries are made, the more deliveries a courier can make, and the more income it will generate.

In any world where the robot does not have complete information, the exact utility of a particular course of action cannot be determined *a priori*. If we can formulate expectations regarding the exact state of the world (possibly conditioned upon information that the robot has some way of obtaining) then it is possible to calculate the *expected* utility of a proposed course of action. This is the metric on the robot's performance that we seek to maximize. The utility of the robot's performance we define in terms of the time required to complete the entire set of deliveries (*i.e.*, the less time, the better).

How might a robot courier plan, in order to do its job better? The robot must fix upon a tour that visits all of locations on its current list of deliveries. We refer to this as *tour improvement* planning. Choosing a tour may have considerable impact on the length of time required to visit all the locations. Once the robot has an ordering for the locations, it may spend time determining how to get from one to another of them. We refer to this as *path* planning. We assume that path planning is accomplished by constructing an ordered set of *target points* between the two locations. Planning in any greater detail we ignore. Arguably, the tasks involved in navigating between target points are sufficiently routine that "planning" is not normally required.[1] In addition, whatever work is done on achieving tasks at lower levels will be local in scope (determining which door to go in, avoiding bumping into things, *etc.*) and should not greatly affect the expected utility of tour improvement or path planning.[2]

The benefit to be expected from either tour improvement or path planning will depend on the current situation. Finding a tour that avoids the area around the Lincoln Tunnel at rush hour is most likely time well spent. On the other hand, if the next location to visit is straight down 5th Avenue from the robot's current position, there may be only a minor benefit to refining a path to get there. In either case, time spent planning may hurt more than it helps if it causes the robot to delay too long. Deliberation scheduling for the robot courier consists of allocating time to algorithms for tour improvement and path planning based on the expected improvement in the robot's performance. An *optimal* deliberation schedule for a given situation is a deliberation schedule that maximizes the expected utility of the robot's performance in that situation. An optimal deliberation scheduling algorithm always generates the optimal schedule for the current situation. All the deliberation scheduling algorithms discussed in this paper are optimal.

In the next section, we describe an experimental domain within which we have implemented a solution to the robot courier problem.

---

[1] See for example [3]

[2] We are not ignoring the possibility that the robot may fail at low-level tasks so as to require additional planning. The possibility of failure is included in the expected utility of the plans generated.

3 The gridworld

One of the salient characteristics of solutions using expectation-driven iterative refinement is their reliance on expectations to support deliberation scheduling. Using expectations is necessary because we have no way of determining the precise behavior of the planning algorithms employed by the robot, apart from actually running them. Accordingly, an important part of constructing a solution to the robot-courier problem consists of gathering statistics on which to base expectations regarding the anytime algorithms for path planning and tour improvement. For the purposes of this exercise we have designed a simulated world in which the required statistics can be gathered easily. We call this simulated world the *gridworld.* The gridworld consists of a rectangular subset of the integer plane ($\mathcal{Z}^2$), where each point is a location that may be occupied by the robot, or by something else, but not both. The robot can move onto any of the 4 neighbors of the point it currently occupies, provided that neighbor is not already occupied. We define the distance between adjacent locations to be one unit. The robot has a map of the world, and is capable of finding its way from one point to another without a planned path, by keeping track of the heading of the destination as it performs a form of obstacle avoidance. Path planning helps because a planned path may be more direct.

The statistics we needed were gathered over randomly-generated instances of the gridworld with a predetermined size and a fixed probability that any given location is occupied (.2). We will refer to these instances as *standard* gridworlds.

4 Anytime algorithms

In this section, we present and analyze anytime algorithms for path planning and tour improvement in the gridworld. Before proceeding with our analysis, we introduce some terminology and notation:

- $\mathcal{L}$ is a set of *locations*, defined as points in the integer plane ($\mathcal{Z}^2$).
- $\mathcal{S} = \langle l_1, l_2, \ldots, l_m \rangle$ is a *tour*; an ordered sequence of locations in $\mathcal{L}$.

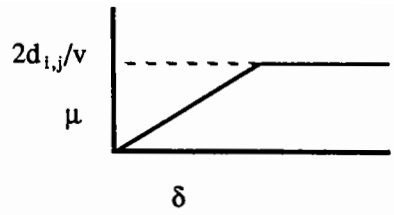

Figure 1: The expected change in travel time for path planning

- $d_{i,j}$ is the distance between locations l_i and l_j. In the gridworld with 4 neighbors, $d_{i,j}$ equals the sum of the difference in x and y coordinates between l_i and l_j.
- $\mathcal{D}_{\mathcal{S}}$ is the length of the tour $\mathcal{S}$ ($\mathcal{D}_{\mathcal{S}} = \sum_{i=1}^{m-1} d_{i,i+1}$).
- v is the mean speed of the robot, moving without incident or obstacle. We assume that the time required to attain this speed is small compared to the time the robot spends travelling between locations.
- $T_{i,j}(\delta)$ is the expected amount of time required to go from l_i to l_j, given an amount of time δ spent planning a path.
- $T_{i,j} = T_{i,j}(0)$ is the expected amount of time required to go from l_i to l_j, without a path.

4.1 Path planning

The robot plans a path from one location to another by constructing an ordered set of adjacent locations (target points), starting at the initial location and working toward the destination. The algorithm used is a heuristic search similar to an algorithm described in [10].[3] In the gridworld, the algorithm when run to completion produces a path whose length averages 1.07 * the distance between the begin and end points, and examines a number of locations averaging 1.1 * the distance from one location to the other. Since it proceeds by searching from the initial point towards the destination, it can be interrupted at any time to return a partial path

[3]Basically A*, using the distance to the goal as a heuristic evaluation function.

whose expected length is a linear function of the time spent planning, as is the expected distance remaining to the destination.

The robot can traverse a path at speed v, so that the expected time required to go from l_1 to l_2 given a completely defined path is $a * (d_{1,2}/v)$, where $a = 1.07$. The mean time required to move from one location to another without a path (*i.e.*, working around obstacles as they are encountered) turns out to be a linear function of the distance between them. For standard gridworlds, the constant involved is approximately 3. Thus $T_{i,j} = b * d_{i,j}/v$, where $b \approx 3$. In the interests of simplifying the analysis in the rest of this paper, we assume that $a = 1$ and $b = 3$.

We define $\mu(\delta)$ to be $T_{i,j} - T_{i,j}(\delta)$, the expected time saved by a partially defined path generated using an allocation of deliberation time δ. Then $\mu(\delta) = \frac{2\delta}{\tau_{pt} v}$, up to $\delta = \tau_{pt} d_{i,j}$, where τ_{pt} equals the expected amount of time required to add another target point to an existing path.[4] Figure 1 shows how the expected savings in travel time varies with the time spent planning. We call such a function a *performance profile* for the anytime algorithm. The dependent variable for a performance profile may be any characteristic of the result of an algorithm that we find significant. One anytime algorithm could have several performance profiles tracking different attributes of the results it returns.

4.2 Tour improvement

We assume that the cost of a tour equals the total distance traversed. Though we make the assumption, it is not apparent that a minimum-distance tour will take minimal time to traverse, including planning time. In fact, it is possible to construct two tours such that the longer one is expected to take less time to traverse. However, our expectations regarding the result of tour improvement are in terms of the length of the entire tour, *not* the distribution of the distances between individual locations. Thus we can make no better assumption about the individual distances than that they are all equal. Given two tours of different lengths, both assumed to have all

[4] We assume that the robot can work on a path between two locations only before it starts moving between them. This is a stronger assumption than we might like, but it can be relaxed fairly easily at the cost of complicating the deliberation scheduling algorithm (reschedule, breaking up the area to be traversed into planned and unplanned segments).

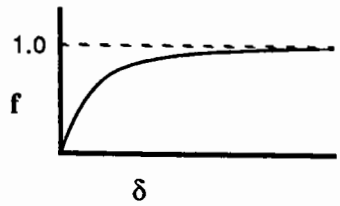

Figure 2: Performance profile for tour improvement planning

individual distances equal, the expected time to traverse the longer tour will be greater.

To construct a useful anytime algorithm, we need an initial answer whose cost is known, and some measure of how that cost is expected to improve with deliberation time. For tour improvement planning, the initial cost is the length of the initial tour (choose one at random, if necessary), and the optimal cost is the length of the optimal tour. The length of the optimal tour for n gridworld locations can be approximated by $1.2\overline{d}(n)^{.58}$, where $\overline{d}$ is the mean distance between the points on the tour, with a standard deviation of about 10% for tours of between 10 and 15 locations.

The next step is to find an anytime algorithm for tour improvement planning. Edge-exchange algorithms [13] produce tours that are progressively closer to optimal by exchanging small sets of edges (typically 2 or 3) such that the length of the overall tour decreases. It has been shown empirically that for large (100 city) tours, these algorithms average within about 8% of the optimal tour length when run to completion [11]. In the gridworld, a tour-improvement algorithm that looks for pairs of edges to exchange averages within 2% of optimal when run to completion on considerably smaller tours (between 10 and 15 points). The mean number of exchanges necessary to obtain an almost-optimal tour for between 10 and 30 locations is approximately $(.61n - 2)^{1.1}$, with a mean error of less than 1% and a standard deviation of about 10%. The mean improvement in tour length after k exchanges as a fraction of the maximum possible improvement can be approximated by a function of the form $f(k) = 1 - e^{-\lambda k}$, where λ is a function of the size of the tour. For tours of 24 locations, a value for λ of 0.29 results in a mean error and standard deviation

both under 1%. Figure 2 gives a performance profile for tour improvement, showing the expected improvement in tour length as a fraction of the maximum possible, given an allocation of deliberation time δ. Because the algorithm always looks for the best pairs of edges to exchange, the time required for any exchange is roughly constant for a given tour size.

Figure 2 is not a graph of the expected tour length over deliberation time. Given an initial tour length $\mathcal{D}_{\mathcal{S}}$, an estimated optimal tour length $\mathcal{D}_{\mathcal{S}^*}$, and an average length of time t_s to make one exchange, the expected tour length as a function of time is $g(t) = \mathcal{D}_{\mathcal{S}^*} + e^{-\lambda t/t_s}(\mathcal{D}_{\mathcal{S}} - \mathcal{D}_{\mathcal{S}^*})$. Given the assumption that the minimum-length tour will take the least time to plan for and traverse, minimizing expected tour length maximizes expected utility.

5 Deliberation scheduling

In this section, we present deliberation scheduling algorithms for three restricted versions of the robot-courier problem:

Pr-I. Deliberation scheduling for path planning for a fixed sequence of locations.

Pr-II. Deliberation scheduling for tour improvement followed by path planning for the improved tour.

Pr-III. Deliberation scheduling for a limited interleaving of tour improvement and path planning.

5.1 Pr-I

Assume a tour $\mathcal{S} = \langle l_0, l_1, \ldots, l_k \rangle$, where l_0 is the robot's current location. Let δ_i be the amount of time allocated to deliberating about the path from l_{i-1} to l_i, for $1 \leq i \leq k$. Let $\mathsf{begin}(l_i)$ and $\mathsf{end}(l_i)$ be the begin and end points, respectively, of the period of time over which the robot is travelling from l_{i-1} to l_i. This defines a time line with the following constraints:

- $\forall l_i : 1 \leq i \leq k, \mathsf{begin}(l_i) = \mathsf{end}(l_{i-1})$

- $\forall l_i : 1 \leq i \leq k, \mathsf{end}(l_i) = \mathsf{begin}(l_i) + T_{i-1,i}(\delta_i)$ — (δ_i initially 0)

We will also use the following definitions:

- Let $\delta_i^* = \tau_{pt} d_{i-1,i}$ — the maximum useful value of δ_i.
- Let $c = \frac{2}{\tau_{pt} v}$ — the slope of the function $\mu(\delta_i), 0 < \delta_i < \delta_i^*$.
- Using the results of Section 4.1:

$$\mu(\delta_i) = T_{i-1,i} - T_{i-1,i}(\delta_i) = \begin{cases} c\delta_i & \delta_i < \delta_i^* \\ 2(d/v) & \delta_i \geq \delta_i^* \end{cases}$$

The procedure given in Figure 3 is optimal in exactly the sense defined in the Introduction. In other words, the algorithm generates a deliberation schedule for $\mathcal{S}$ such that there is no other schedule for which the robot is expected to complete the tour in less time.[5] The asterisk marks a line that is important to the discussion in the Appendix, and has no other significance. The procedure works backward from the last path to be traversed, ensuring that each path is allocated the maximum useful amount of time. There may be times where the robot is predicted to be traversing a path and all later paths have already been allocated as much time as will produce any improvement. The length of the earliest such predicted idle time is the value of the variable gap. The algorithm is simplified considerably by the fact that c is the same for any path (*i.e.*, as long as $\delta_i < \delta_i^*$, it doesn't matter which δ_i gets allocated, for $1 < i < k$). That c is a constant is a property of the gridworld, not an assumption we make. Figure 4 shows a deliberation schedule constructed for an example tour. The shaded regions represent time allocated for deliberation.

5.2 Pr-II

We can extend Pr-I by allowing $\mathcal{S}$ to be reordered, but only before any path planning gets done. In other words, if the initial tour is far from optimal, we expend some effort improving it before doing path planning. Given $\mathcal{S}$, we can generate a deliberation schedule for path planning using Sched-1. Sched-1 returns the expected time to complete $\mathcal{S}$, including path planning time. We also know the expected length of

[5] Proofs that the algorithms in this section are optimal may be found in the Appendix.

```
Procedure: Sched-1( S)
  begin
        t := end(l_k)
        for j = k to 1, step −1
           begin
              if t > begin(l_j), then
                 gap := max(cδ*_j, t − begin(l_j))
                 t = begin(l_j)
               else
                 gap := 0
 (*)           δ_j := min(δ*_j, t − begin(l_1), 1/(c+1)(t − begin(l_1) + gap))
              t := t − δ_j
              gap := max(0, gap − cδ_j)
              if t = begin(l_1), goto exit
           end
  exit:: if c > 1 and gap > 0, then
           Δ := min(δ*_j − δ_j, 1/c * gap)
           δ_j := δ_j + Δ
        δ_i := 0, 1 ≤ i < j
        return Δ + Σ_1^k (T_{i−1,i}(δ_i))
  end
```

Figure 3: An algorithm for Pr-I

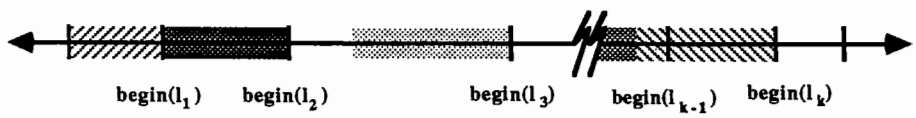

Figure 4: A deliberation schedule for path planning

the tour $\mathcal{S}'$ resulting from a given amount of time allocated to tour improvement. To calculate the expected utility of tour improvement, we need the expected time required, including path planning time, to complete $\mathcal{S}'$. We have no better information than the assumption that all the distances between locations in $\mathcal{S}'$ are equal. So the path planning problem involves k paths, all of length $d = \mathcal{D}_{\mathcal{S}'}/k$. In this case, the expected time to complete the tour given an optimal deliberation schedule can be expressed as a linear function of d (and thus of $\mathcal{D}_{\mathcal{S}'}$). We use the same definitions as in the last section. For $2 \leq i \leq k-1$, we drop the subscripts from T and δ, since they will have the same values for any such i. δ^* and d are constant for any i, $1 \leq i \leq k$. There are four cases, depending on the value of c:

- $c \geq 2$. $\delta_1 = \delta_k = \delta = \delta^*$. Then $\mathcal{T} = (k + \frac{2}{c})(d/v)$.
- $\frac{2(k-1)}{k+1} \leq c < 2$. $\delta_k = \delta = \delta^*$, $\delta_1 = (\frac{1}{c}(k+1) - \frac{2}{c^2}(k-1))(d/v)$.
  Then $\mathcal{T} = (1 + \frac{1}{c}(3k-1) - \frac{2}{c^2}(k-1))(d/v)$.
- $\frac{2}{3(k-1)} \leq c < \frac{2(k-1)}{k+1}$. $\delta_1 = 0$, $\delta_k = \delta^*$. We want δ s.t. $T + (k-2)T(\delta) - (k-2)\delta - \delta^* = 0$. Solving, $\delta = \frac{3(k-1)-\frac{2}{c}}{(k-2)(1+c)}(d/v)$. Then $\mathcal{T} = \frac{3k+c}{1+c}(d/v)$.
- $c < \frac{2}{3(k-1)}$. $\delta_1 = \delta = 0$, and $\delta_k = (k-1)T = (k-1)(3d/v)$.
  Then $\mathcal{T} = 3((1-c)(k-1)+1)(d/v)$.

The cases are distinguished by differing values of c and k. In the first case, the expected benefit from path planning is so high that each path is allocated the maximum deliberation time. In the second case, the expected benefit is slightly less, so that setting δ_1 to δ^* might make the tour take longer. In the third case, the expected benefit of path planning is low enough that any value of $\delta_1 > 0$ has a negative utility, and that only δ_k is automatically allocated the maximum amount of time. In the last case, the improvement to be expected is so small that only the last path in the tour is worth planning for, and that only while travelling among the other locations.

For a particular domain, c and v are constants. k is fixed for a given set of locations. Let $f(d)$ be the appropriate linear function of d, which we will write as αd. Let $g(\mathcal{S}, \delta')$ be the expected length of the tour returned by the anytime algorithm for tour improvement, allocated δ' units of time. Then the following procedure is an optimal deliberation scheduling algorithm.

```
Procedure: Sched-2( S)
  begin
    Choose δ'  s.t.  T = δ' + (α * g( Sδ') / k) is minimized.
    if Sched-1( S) < T, then δ' = 0.
    return δ'
  end
```

Once the tour improvement algorithm has been run, we run Sched-1 on the resulting tour. That the required minimization can be done easily is guaranteed by the fact that g' is always negative (see Section 4.2); there are no local minima.

5.3 Pr-III

Interleaving tour improvement and path planning arbitrarily can be complex; there are an exponential number of ways in which a given tour might be subdivided into pieces on which to run the tour improvement algorithm. There is a useful special case, however, for which we can construct a deliberation scheduling algorithm whose complexity is $O(n^2)$ in the number of locations. We define the scheduling problem Pr-III to consist of determining some number of locations for which we plan paths first, followed by the remaining locations, which are treated as an instance of problem Pr-II. Sched-3 is an optimal deliberation scheduling algorithm for Pr-III. Sched-3 uses the procedure Sched-3a, given in Figure 5.

```
Procedure: Sched-3( S)
  begin
    min := ∞
    for i = 0 to k − 2
      begin
        s := Sched-3a( S, i)    ;returns est. time.
        if min > s, then
            min := s, i* = i
      end
    t := Sched-1( S)
    if t ≤ s, then
      i* = k
    return i*                          ;use the minimum-time del. schedule
  end
```

```
Procedure: Sched-3a( S, i)
  begin
    δ' := end(l_i) − begin(l_1), t := end(l_i)                    ;;Allocate it all to δ'.
        for j = i to 1, step −1
          begin
            if t > begin(l_j), then
               gap := max(cδ*_j, t − begin(l_j)
               t = begin(l_j)
             else
               gap := 0
  (*)          while T'(δ' − δ_j) < c/(c+1) and
                      δ_j < min(δ*_j, t − begin(l_1), 1/(c+1)(t − begin(l_1) + gap))
               Allocate δ_j back from t.
            t := t − δ_j
            δ' := δ' − δ_j
            gap := max(0, gap − cδ_j)
            if t = begin(l_1), goto exit
          end
 exit:: if c > 1 and gap > 0, then
          begin
            while T'(δ' − Δ) < (c−1)/c and Δ < min(δ*_j − δ, 1/c * gap)
                Allocate Δ back from begin(l_1).
            δ_j := δ_j + Δ
            t := t − Δ
            δ' := δ' − Δ
          end
      while T'(δ' + Δ') > 1
        Allocate Δ' back from t.
      return Δ' + Δ + Σ_1^i T_{j−i,j}(δ_j) + T(δ' + Δ')
end
```

Figure 5: An algorithm for Pr-III

We use most of the same definitions as for the preceding algorithms. Let $\mathcal{T} = \alpha * g(\langle l_i, \ldots, l_k \rangle, \delta')$, and $\mathcal{T}'$ be the first derivative of $\mathcal{T}$ with respect to δ'. The algorithm works by assuming initially that all the travel time for the locations $l_0, \ldots, l_i$ will be used to deliberate about finding an improved tour for $l_i, \ldots, l_k$. Using a procedure almost identical to Sched-1, the algorithm then sweeps back from begin(l_i), determining for each increment of time whether spending time on tour improvement or path planning is expected to save more time in the final schedule.

6 More general problems

The robot courier problem can be generalized in several ways. For one, time spent planning or acting is accounted for in an extremely simple way. The time taken to complete the tour does not merely figure in the utility of the robot's behaviour, it *is* the (inverse) utility of the robot's behaviour. This makes the treatment of preformance profiles for planning very easy: the expected utility is a linear function of the expected time to complete the tour, which is a linear function of the time to plan for and traverse each path, plus the time allocated to tour improvement.

Situations that include deadlines are a bit more complicated; the distribution of error in performance profiles becomes significant. Consider, for example, a situation where the cost of being late is very high: if you don't make it to the altar on time, your intended spouse will take offense, and leave you. Suppose that you have two path-planning algorithms, one of which usually produces a good path very fast, but every once in a while takes a very long while and produces a very bad path. The other algorithm does not produce quite as good a path, but is effectively guaranteed to generate a path whose length is very close to that predicted by the performance profile. The expected length of the path produced by the first algorithm is better than that for the second (the earlier you arrive, the more champagne you have a chance to drink). But the expected *utility* is higher for the second path, because there is less chance of being abandoned. Deliberation scheduling for problems involving deadlines requires making tradeoffs like this; deadlines and the cost of being tardy will determine which algorithm has the higher expected utility.

There are other situations that require computing joint distributions. Perhaps an observation of a particular type can be used to predict with certainty that an

event will happen, but only approximately when. In this case we are interested in the distribution of a random variable corresponding to the value of the result we have at the time the event happens. This will depend on two distributions: one for the value of the result obtained given a particular amount of time, and the other for the time until the event occurs. We may also have expectations regarding the *occurrence* of events. Problems in this class are a subject for future research.

7 Related work

There is a growing interest in applying decision theory to AI problems. [8] outlines the current state of this effort, and points out some areas where further research might be most fruitful. Horvitz [7] investigates the general problem of applying decision theory to resource-bounded reasoning. In addition to our work on planning, the application of decision theory to the control of resource-bounded reasoning is currently being pursued in at least the areas of medical diagnosis [6], heuristic search [5], and computer vision [12]. Breese and Fehling [2] provide an abstract characterization for the decision-theoretic control of an autonomous agent, including reasoning about tradeoffs among planning, acting, and gathering information. Wellman [15] provides an extended analysis of the use of uncertain knowledge in planning, but does not discuss the issue of resource bounds on reasoning. The work in [14] does not make explicit use of decision theory, but embodies a similar philosophy to our work in that the analysis of a specific domain is used to guide the construction of strategies for controlling reasoning in that domain.

8 Conclusion

In [4], we define a framework for solutions to time-dependent planning problems called expectation-driven iterative refinement. This framework is intended to restrict the form of those solutions to facilitate decision-theoretic analysis. In the current paper, we are interested in a methodology for applying expectation-driven iterative refinement. Specifically, we want a way to identify problems as likely candidates and a body of techniques for generating solutions. We detail the process of

solving a time-dependent planning problem involving a mobile robot courier, as a way of investigating such a methodology.

The process of generating a solution provides some indication of the characteristics of the methodology we are developing. Good candidates for solution using expectation-driven iterative refinement will be problems that can be decomposed into subproblems that can be solved using anytime algorithms for which we have, or can generate, useful expectations in the form of performance profiles. There is an additional restriction, which is that the problem must be decomposed in such a way that we know how to combine the profiles for the subproblems so as to produce a performance profile for time spent on the problem as a whole.

There are several lessons to be taken from this paper on how to go about solving a time-dependent planning problem. For one, we do not "solve" either of the planning problems facing the robot. We provide approximate solutions, and some measure of the errors to be expected. A problem where a precise answer is critical would probably use different algorithms. An interesting point that arises in the generation of performance profiles is that we need a fast estimate of how good an answer is obtainable. This estimation will of necessity introduce some error (otherwise the anytime algorithm would be superfluous) and we should know how much. Assuming that the anytime algorithms can be represented by continuous performance profiles does not for this problem introduce any large error—the increments in which planning is done are small enough that the approximation is a reasonable one. The form of the solution we provide is strongly determined by the structure of the environment within which the robot operates. The environment determines the performance profiles of the anytime algorithms the robot employs for planning, and to a lesser extent the algorithms themselves.

Solving even a moderately complex time-dependent planning problem requires combining the results of several processes. In some problems, there may be several unrelated subproblems competing for processor time. For the robot courier, this would correspond to allocating planning time for the individual paths on a tour. Also, most useful anytime algorithms we know of apply to sufficiently simple problems that any interesting planning problem will require combining the results of several anytime algorithms. In order to plan for a tour, our robot courier needs to combine the results of path planning and tour improvement. These are not separate

competing processes. The expected utility of time spent path planning depends on the result of tour improvement. In our example, combining expectations for the two planning algorithms is straightforward: they can be composed, given an assumption that is justified by the information available to the robot. Other problems and other decompositions will require combining expectations in different ways.

Future work will involve further exploration of the appropriate classes of problems for solutions using expectation-driven iterative refinement, and the refinement of a methodology by which those solutions can be constructed.

References

[1] Aho, Alfred V., Hopcroft, John E., and Ullman, Jeffrey D., *The Design and Analysis of Computer Algorithms*, (Addison-Wesley, Reading, Massachusetts, 1974).

[2] Breese, John S. and Fehling, Michael R., Decision-Theoretic Control of Problem Solving: Principles and Architecture, *Proceedings of the 1988 Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN*, 1988.

[3] Chapman, David and Agre, Phil, Pengi: An Implementation of a Theory of Activity, *Proceedings AAAI-87, Seattle, Washington*, AAAI, 1987, 268–272.

[4] Dean, Thomas and Boddy, Mark, An Analysis of Time-Dependent Planning, *Proceedings AAAI-88, St. Paul, Minnesota*, AAAI, 1988, 49–54.

[5] Hanson, Othar and Meyer, Andrew, The Optimality of Satisficing Solutions, *Proceedings of the 1988 Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN*, 1988.

[6] Horvitz, Eric J., Reasoning About Beliefs and Actions Under Computational Resource Constraints, *Proceedings of the 1987 Workshop on Uncertainty in Artificial Intelligence, Seattle, Washington*, 1987.

[7] Horvitz, Eric J., Reasoning Under Varying and Uncertain Resource Constraints, *Proceedings AAAI-88, St. Paul, Minnesota*, AAAI, 1988, 111–116.

[8] Horvitz, Eric J., Breese, John S., and Henrion, Max, Decision Theory in Expert Systems and Artificial Intelligence, *Journal of Approximate Reasoning*, (1988).

[9] Karp, R., Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman Problem in the Plane, *Math. Oper. Res.*, **2** (1977) 209–224.

[10] Korf, Richard E., Real-Time Path Planning, *Proceedings DARPA Knowledge-Based Planning Workshop, Austin, TX*, 1987.

[11] Lawler, E.L., Lenstra, J.K., Rinnooy Kan, A.H.G., and Shmoys, D.B., *The Travelling Salesman Problem*, (Wiley, New York, 1985).

[12] Levitt, Tod, Binford, Thomas, Ettinger, Gil, and Gelband, Patrice, Utility-Based Control for Computer Vision, *Proceedings of the 1988 Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN*, 1988.

[13] Lin, S. and Kernighan, B. W., An Effective Heuristic for the Travelling Salesman Problem, *Operations Research*, **21** (1973) 498–516.

[14] Ow, P. S., Smith, S. F., and Thiriez, A., Reactive Plan Revision, *Proceedings AAAI-88, St. Paul, Minnesota*, AAAI, 1988.

[15] Wellman, Michael P., *Formulation of Tradeoffs in Planning Under Uncertainty*, Technical Report MIT/LCS/TR-427, MIT AI Laboratory, 1988.

A Appendix

A.1 Proof that Sched-1 is optimal

Let $\mathcal{T}$ be the total expected time for the robot to complete a tour, including planning time, for a given deliberation schedule. A deliberation schedule is optimal only if for any increment of time Δ, there is no other way that Δ could be allocated, either to another path or to no path at all, such that $\mathcal{T}$ is smaller. From Section 4.1:

$$T_{i-1,i}(\delta_i) = T_{i-1,i} - \mu(\delta_i) = \begin{cases} 3d/v - c\delta_i & \delta_i < \delta_i^* \\ d/v & \delta_i \geq \delta_i^* \end{cases}$$

We make use of the following properties of deliberation scheduling for this particular problem:

1. Any optimal deliberation schedule can be transformed into a deliberation schedule in which the robot deliberates about each path in turn, in the order they occur, without increasing $\mathcal{T}$.

2. Any optimal deliberation schedule can be transformed into a deliberation schedule in which the robot plans without moving only at the very beginning of the tour, without increasing $\mathcal{T}$.

3. For any locations l_i and l_j, the change in *travel time* in the tour with deliberation time δ_i is the same as for time allocated to δ_j, provided $\delta_i < \delta_i^*$ and $\delta_j < \delta_j^*$.

The proof will consist of tracing the execution of the algorithm, and showing that at any point, allocating deliberation time in any other way would result a larger value of $\mathcal{T}$.[6] We start with small cases, finally providing an argument for tours of arbitrary length k.

Proof:

$k = 1$:

[6]Local minima are not a problem, since μ is monotonic and continuous.

We go to exit with $t = \mathsf{begin}(l_1)$, $\mathsf{gap} = c\delta_1^*$, $j = 1$, and $\delta_1 = 0$. $\mathcal{T} = \delta_1 + T_{0,1}(\delta_1) = \delta_1 + \max(d/v, 3(d/v) - c\delta_1)$. $c\delta_1^* = 2(d/v)$. Thus $\mathcal{T}$ is minimal for $\delta_1 = \delta_1^*$ if $c > 1$ and minimal for $\delta_1 = 0$ otherwise. The code following exit will allocate δ_1 correctly.

$k = 2$.

For $j = 2$, the step marked with an asterisk will allocate δ_2 such that either:

1. $\mathsf{begin}(l_2) - \delta_2 > \mathsf{begin}(l_1)$, and $\delta_2 = \delta_2^*$, or
2. $\mathsf{begin}(l_2) - \delta_2 = \mathsf{begin}(l_1)$.

In the first case, there is no benefit to increasing δ_2. The interval between $\mathsf{begin}(l_1)$ and $\mathsf{begin}(l_2) - \delta_2$ cannot be used for deliberating about the path to $l1$. If $c > 1$, there may be some allocation made to δ_1. For $j = 1$, $\mathsf{gap} > 0$, and so δ_1 will be allocated in the code following exit so that either the idle time gap disappears, or $\delta_1 = \delta_1^*$. If $\delta_1 \neq \delta_1^*$, $\mathsf{begin}(l_1) + \delta_2 = \mathsf{begin}(l_2)$, and so $\mathcal{T} = \delta_1 + \delta_2 + T_{1,2}(\delta_2)$. Any further allocation to δ_1 will increase $\mathcal{T}$.

In the second case, we go to exit with $j = 2$. $\mathsf{gap} > 0$, so we consider increasing δ_2. Since in this case $\mathsf{begin}(l_2) - \delta_2 = \mathsf{begin}(l_1)$, $\mathcal{T} = \delta_1 + \delta_2 + T_{1,2}(\delta_2)$. Using the same argument as for $k = 1$, we see that if $c > 1$, $\mathcal{T}$ is minimal for $\delta_2 = \delta_2^*$, and otherwise δ_2 should not be increased. Any value of δ_1 greater than 0 will increase $\mathcal{T}$.

For arbitrary k, the problem is only slightly more complicated. Consider the schedule when the procedure arrives at exit: For any $i > j$, $\delta_i = \delta_i^*$. We have to justify not increasing $\delta_i, i \leq j$. But we know: $\mathcal{T} = \sum_1^j \delta_i + \sum_j^k T_{m-1,m}(\delta_m^*)$. Unless $i = j$ adding to δ_i can only increase $\mathcal{T}$. If $\mathsf{gap} > 0$ and $c > 1$, δ_j will be allocated so that either the idle time gap disappears, or $\delta_j = \delta_j^*$. Any additional allocation to δ_j would increase $\mathcal{T}$. □

A.2 Proof that Sched-2 is optimal

We need to show that the functions given in Section 4.2 do compute planning and travel time for $\mathcal{S}$ based on an optimal deliberation schedule. This can be

accomplished by verifying that they compute the same deliberation schedule as the procedure Sched-1, except in case 3, which allocates the values for $2 \leq i \leq k-1$ differently, but equivalently (recall from the proof for Sched-1 that for $2 \leq i, j \leq k-1$, time allocated to δ_i and δ_j has the same effect on $\mathcal{T}'$, provided $\delta_i < \delta_i^*$ and $\delta_j < \delta_j^*$).

Proof:

$g(\mathcal{S}, \delta')$ is the expected length of the improved tour $\mathcal{S}'$, given δ' units spent on improving tour $\mathcal{S}$. $\alpha * g(\mathcal{S}, \delta')/k$ is the best estimate obtainable (without information on the distribution of individual distances between locations) on how long it will take to traverse $\mathcal{S}'$, given an optimal deliberation schedule for path planning, which can be produced by running Sched-1 on $\mathcal{S}'$. The expected value of $\mathcal{T}$ given δ' time units spent on tour improvement is $\delta' + \alpha * g(\mathcal{S}, \delta')/k$. For the special case of $\delta' = 0$, a more precise estimate can be generated using Sched-1, because we know the distances between locations. Sched-2 thus makes the best estimate possible of the expected time to visit the set of locations, including time spent on tour improvement, for any value of δ', and picks the value such that the expected time required is minimized (expected utility is maximized). □

A.3 Proof that Sched-3 is optimal

Sched-3 checks every possible interleaving of tour improvement and path planning and traversal allowed by the statement of the problem Pr-III. It remains to be shown that Sched-3a correctly estimates the expected utility of each alternative.

Proof:

The algorithm initially allocates the time spent travelling between l_0 and l_i to tour improvement. Since this is time that figures in the expected completion time for the tour, this allocation will certainly decrease the expected completion time. What the bulk of the procedure accomplishes is to decide for each increment of time whether it would be decrease the expected completion time if that increment were spent on planning some path instead of on tour improvement. Again, the form of the function determining the expected time saved by tour improvement ensures that there is no problem with local minima, so making a local determination of how to allocate deliberation time will produce a global minimum.

The central loop in Sched-3a allocates time for path planning exactly as does Sched-1, and thus has the property that, given that path-planning has a higher expected utility, deliberation time is allocated to the correct path. What is different is the comparison of $\mathcal{T}'$ and c, to determine which type of planning to spend the time on. For any δ_m, $1 < m \leq i$, $\delta_m < \delta_m^*$, allocating an increment dt to δ_m instead of to δ' changes the expected time to complete the tour by $\mathcal{T}'\mathsf{dt} - c\mathsf{dt} + c\mathcal{T}'\mathsf{dt}$. Thus the loop marked with an asterisk takes time from tour improvement for path planning just in case that decreases the expected completion time. The code following exit first determines whether to allocate time to path planning before travelling, making the same checks as Sched-1, plus determining that the net effect of decreasing the time available for tour improvement is in the right direction. It may then allocate time to tour improvement, also before travelling, if the net effect will be to decrease the expected completion time.

Summing the resulting travel times and the time allocated to planning before the robot moves gives the expected utility for the current choice of index i. □