EXPLOITING CAPABILITY CONSTRAINTS TO SOLVE GLOBAL TWO-DIMENSIONAL PATH PLANNING PROBLEMS

R. F. Richbourg, Neil C. Rowe, Michael J. Zyda

Department of Computer Science, Naval Postgraduate School, Monterey, California

Abstract

Mobile autonomous vehicles require the capability of planning routes over ranges that are too great to be characterized by local sensor systems. Completion of this task requires some form of map data. Much work has been done concerning planning paths through local areas, those which can be scanned by on-board sensor systems. However, planning paths based on long range map data is a very different problem. Extant solution techniques require the search of discrete, node and link representations which characterize continuous, two dimensional problem environments. We assume the availability of topographic data organized into regions of homogeneous traversal cost. Given this, we present a solution technique for the long range planning problem which relies on a Snell's Law heuristic to limit a graph search for the optimal solution.

Introduction

Route planning has been an area of interest for many years, particularly in the Operations Research field [1]. Such problems typically feature a discrete number of choices for possible paths and can thus be characterized as essentially one dimensional. Recently, interest in two dimensional route planning has grown. In this domain, there is a continuum of available choices for the next movement. Such planning is appropriate for robot or human cross country routing or for determining placements for the future construction of permanent, linear features such as pipelines or roads. The simplest version of this problem is known as the Find-Path problem [2].

A solution to the classic Find-Path problem consists of a routing plan tailored to a specific environment. In the robotics area, the problem is described as, given the initial and goal locations of an object, and a set of obstacles in the environment, find a continuous path for the movement of the object from the initial location to the goal location which avoids collisions with any obstacles along the way [2]. The problem has typically been associated with planning local motion control or manipulator trajectory control for a robot in a two or three dimensional space where information is gained through integral sensor systems, primarily visual.

Another version of the classic Find-Path problem requires planning motion over ranges too great to be explored by local sensor systems. We term this problem the extended Find-Path problem. Long range planning is a useful and difficult problem which humans solve easily many times each day. Prior to initiating any movement, humans, either consciously or not, plan a particular course to be followed. Moreover, the planned course has the characteristic that some factor is optimized (or at least perceived as optimized). A typical optimization factor is that of time.

A salient difference between the classic and extended Find-Path problems concerns optimization. The classic version posits a binary view of the world. Every point in the environment is either traversable or impassable. An optimal route in such an environment can be found by planning the shortest distance path around the obstacles. This "least distance implies least cost" assumption is viable in the classic version because the range of the sensor systems generally encompasses only a uniform medium of traversable areas. The validity of this assumption can easily be established by an example (using time as a measure of route cost). The time required to walk across any portion of a floor is the same over the entire floor, given the absence of obstacles.

The "least distance implies least cost" assumption is not reasonable for the extended Find-Path problem because the increased physical range over which travel occurs generally implies that more than one traversable medium is represented. Regardless of the factor which is to be optimized, one medium can not be assumed to have the same traversal characteristics as all others, given the diversity of objects which can be found in the physical world. Thus, a new parameter is introduced into the cost computation for each possible path. In the binary case problem, the cost of traversing every route is computed by the simple formula $Cost = \sum_{i=1}^{n} d_i$ where there are n line segments in the complete path and d_i is the Euclidean distance along the i'th line segment. Considering the extended problem, the formula becomes $Cost = \sum_{i=1}^{n} c_i \times d_i$ where n and d_i have the same meaning. The new parameter, c_i reflects the cost of the i'th line segment. Also, the number of line segments is typically increased. A straight line traversing two different cost areas is represented by two line segments. The addition of cost information has the effect of invalidating the straight line hypothesis so prevalent in the binary case solution techniques.

Brief Review of Some Extant Solution Techniques

Visibility Graphs The prototypical solution technique for the binary case problem is known as the Visibility Graph or VGraph method [7]. This method relies on the knowledge that the optimal path must be either a straight line from start to goal or a series of line segments which includes obstacle vertices as turn points along the start to goal path. The VGraph technique creates a graph representation of the problem, G = (V,L), where the set V includes all obstacle vertices plus the start and goal locations. The straight line hypothesis is used to create the set L. Whenever an unobstructed line segment connects any two members of V, that segment is included in L. When the construction of the graph is complete, standard graph search techniques can be used to find the optimal path in G.

The exhaustive nature of G defeats the interacting sub-problems difficulty of the Find-Path problem. The cost is that every obstacle in the problem environment must be considered. When o obstacles are present, having a total of v vertices, the cardinality of V is $n = v + 2$. In the worst case, $\sum_{i=1}^{i=n-1} i$ line segments must be inspected to form L. This can be expensive for large problem spaces. A dynamically created graph considering only those obstacles in the general area of the optimal path is less costly. However, care must be taken to insure a global problem view so that interacting subproblems do not prevent optimal solutions.

Propagation The wavefront propagation technique is the most common method used to solve the extended problem [5,6,11,12]. In this technique, a uniformly spaced lattice is imposed on the problem representation. Each link is assigned a cost relative to the actual cost of traversing the associated real world terrain. Then, an omnidirectional graph search, similar to the propagation of a wave, is conducted. When the wavefront touches the goal, the optimal path can be retrieved by referencing backpointers or a similar technique. This method has the primary advantage that it is the only known technique which provides optimal solutions to the extended problem. However, it has significant deficiencies.

The propagation technique is very expensive. Much of the search is expended in areas that do not hold any portion of the optimal path since there is no guiding principle such as the straight line hypothesis used in the binary case problem. Also, there is a digital bias inherent in the problem representation so that the method actually returns a set of "optimal" paths, each of which must be compared to determine the unique solution [6]. That is, the method finds a set of equal cost, "stairstep" approximations to straight line solutions. These must be smoothed out to characterize the optimal path.

Homogeneous Regions One technique which attempts to avoid the exhaustive lattice of the propagation technique is the homogeneous regions approach [3]. This method relies on the knowledge that the physical world is not a series of disjoint points, but rather a grouping of similar regions (fields, lakes, etc.). The points inside each region have equivalent traversal characteristics. The method operates by representing the problem space as a grouping of polygons (homogeneous regions) and conducting a graph search between region centers. The method has deficiencies; not everything in the physical world can be accurately modeled by centrally symmetric figures and paths between region centers need not include the best routes. However, the observation that the real world is a grouping of regions, not distinct points, is important.

Desirable Solution Technique Characteristics
A solution technique for the extended Find-Path problem will have several key properties. First, there must be provisions to deal with the interaction of subproblems. Specific domain knowledge can been employed to prevent these difficulties as is the case in the VGraph model where knowledge leads to the exhaustive decomposition of the problem into a graph of obstacle vertices which can be intelligently searched. The analogous problem decomposition for the extended problem leads to the imposition of a uniform lattice structure as in the wavefront propagation technique. Again, an intelligent graph search can be conducted to find an optimal solution. The salient difference is that an unintelligent problem representation confounds the search strategy and leads to problems of representational resolution, combina-

torial explosion, accumulation of error, and a multiplicity of solutions which erroneously appear as equivalent. The homogeneous regions approach attempts to establish a more intelligent problem representation by grouping similar regions together. However, this technique also fails due to poor representational robustness (not all physical world features can be adequately modeled) and the lack of an appropriate straight line hypothesis to guide search (moving from region center to region center is inadequate). We note that a dynamically created graph can lead to greater efficiency by avoiding wasteful computations.

Humans solve similar problems routinely, providing reasonable solutions quite efficiently. Their abilities seem to rely on several principles. First, they make use of topographical knowledge and knowledge of their own capabilities to partition the problem representation into homogeneous areas with similar traversal characteristics. Secondly, humans reason at different levels of abstraction. A high level route plan might contain the step "go through Smithfield" while a lower level abstraction details the route through the town. In this sense, abstraction is as useful in combating combinatorial explosion as is the homogeneous regions premise. Humans are also opportunistic. Special cases of terrain features present opportunities for problem decomposition. This concept generalizes to the appearance of a corridor through an otherwise impenetrable obstacle. An example is the occurrence of a single road through a densely wooded and treacherous mountain area. A suitable solution technique for the extended Find-Path problem must be able to achieve opportunistic decomposition by recognizing similar situations. Another useful aspect of human reasoning is that it is multidirectional. Moreover, directionality is intelligently specified. The wavefront technique is multidirectional; however, omnidirectional search is not an intelligent strategy.

In summary, a suitable solution method for the extended problem should exhibit several properties. These include the use of a basic, guiding principle for search (such as the straight line hypothesis), domain knowledge, capability knowledge, multidirectional (at least bidirectional) and intelligent search, opportunistic decomposition, intelligent problem representation, and levels of abstraction. Also, the solution provided should be in some sense optimal. We have noted that humans quickly solve the problem, but not necessarily with optimal results. Graph theoretic techniques can provide optimal solutions, but not necessarily quickly. A suitable solution method for the extended Find-Path problem will achieve the best traits of both methods.

Proposed Solution Technique
Representing the problem as a lattice of uniformly spaced nodes produces a firm, graph theoretic basis from which an optimal solution can be obtained. However, obtaining this solution is computationally expensive. In the wavefront propagation technique, the computational cost is relative to the distance to be traveled from the start location to the goal. The number of grid squares examined here is approximately proportional to the area of a circle whose radius is the start to goal distance. Thus, the computational cost grows approximately in proportion to the square of this distance. A technique to intelligently guide the wavefront search is not readily apparent. Because of the diversity of speeds attainable over different physical terrain features and the arbitrary nature of their physical, relative placement, it is generally not possible to determine if a specific area is more favorable to the production of an optimal path than any other area, unless that area is explicitly moved to and examined. This characteristic causes the exhaustive, omnidirectional nature of the wavefront propagation search. We desire to produce a method which constrains the search required of wavefront propagation

techniques. Further, the developed method must be relatively efficient as computational excess is the primary problem to be overcome.

Normally, limiting graph searches involves finding good decision criteria for use at branch points. This does not seem appropriate for the extended problem. However, we can bound the physical size of the entire graph. Such a bound is possible and easily derived. Any solution to the problem serves as an upper bound on the cost of the optimal solution. Consider a hypothetical path from start to goal which traverses only the least desirable (i.e. highest cost) regions but avoids all obstacles. Compute the cost of this hypothetical path. Now, compute the distance which can be traversed when only the most favorable terrain type is used and the same cost is accrued. Consider this second distance to be the sum of the distances from each foci to the boundary of an ellipse. Let the start and goal locations define such foci. Clearly, the optimal path must lie entirely within this ellipse. Thus, we have created a physical bound on the area to be searched by a propagation technique. However, the bound is not tight. Remembering that computational cost changes in approximate relation to the square of the distance to be traveled, we note that obtaining the best bound possible is highly desirable.

The efficiency of human produced solutions to similar problems stems from the human ability to group many distinct points into homogeneous regions and then plan paths on a region to region basis. The resulting paths may not be optimal, but they are reasonable and can be obtained efficiently. Given this broad characterization, it is apparent that an efficient, human like search of a homogeneous region problem representation constitutes a suitable vehicle to determine a bound for the propagation techniques. Optimality is not a requirement of the bound. An efficiently produced, reasonable solution is the goal.

The two drawbacks associated with the extant homogeneous regions method are detailed above. First, we overcome the problem associated with region centers. Consider the problem of finding the optimal path in a simple case. Figure 1 depicts a situation where the requirement is to move from an initial location inside a "cheap" region to a goal also in the "cheap" region when there is a rectangular "expensive" region on the straight line path between them. The optimal path is some perturbation of a straight line which trades increased distance in the cheap region for deceased distance in the expensive region. In Figure 1, S1 and S2 represent cost reciprocals. D, d1, d2, y1, y2 and y3 represent Euclidean distances. Let T denote the cost of a path. The equation describing T is a sum of terms as below.

$$T = \frac{\sqrt{d\,1^2 + y\,1^2}}{S\,1} + \frac{\sqrt{(D - d\,1 - d\,2)^2 + y\,3^2}}{S\,2} + \frac{\sqrt{d\,2^2 + y\,2^2}}{S\,1}$$

Taking the first partial derivatives of T with respect to d1 and d2 and setting the resulting equations equal to zero produces:

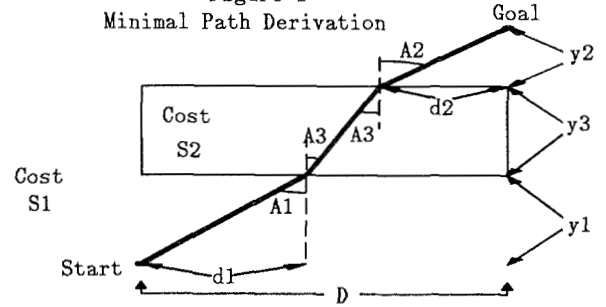$$\frac{d\,1}{S\,1\sqrt{y\,1^2 + d\,1^2}} = \frac{D - d\,1 - d\,2}{S\,2\sqrt{y\,3^2 + (D - d\,1 - d\,2)^2}}$$

$$\frac{d\,2}{S\,1\sqrt{y\,2^2 + d\,2^2}} = \frac{D - d\,1 - d\,2}{S\,2\sqrt{y\,3^2 + (D - d\,1 - d\,2)^2}}$$

In Figure 1, let A1, A2 and A3 represent angular measures. Note that:

$$sin\,(A\,1) = \frac{d\,1}{\sqrt{y\,1^2 + d\,1^2}}, \quad sin\,(A\,2) = \frac{d\,2}{\sqrt{y\,2^2 + d\,2^2}}$$

$$sin\,(A\,3) = \frac{D - d\,1 - d\,2}{\sqrt{y\,3^2 + (D - d\,1 - d\,2)^2}}$$
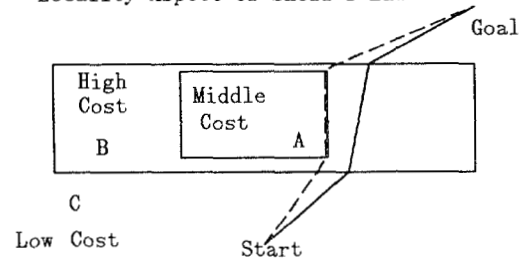

Figure 1
Minimal Path Derivation

Then, the above computations simplify to:

$$\frac{sin\,(A\,1)}{S\,1} = \frac{sin\,(A\,3)}{S\,3} = \frac{sin\,(A\,2)}{S\,1}$$

A minimum cost path is the result of satisfying this final equation.

We note that the final equation is exactly Snell's Law [4] commonly used in optics to compute the paths of light rays traveling through media of different reflective indices. There is no known, closed form equation to solve a Snell's Law problem. Iteration and search (such as bisection search) are normally used to find optimal paths. Also, Snell's Law is known to express an entirely "local" relationship. Thus, the homogeneous regions need not be strictly represented by rectangular figures, any linear boundary separating two regions is sufficient to support application of the law. This characteristic eliminates the difficulties of those homogeneous regions methods which are based upon movement from region center to region center only [3]. Another benefit of Snell's Law is that the optimal path can be computed across any number of consecutive, linear boundaries.

Figure 2
Locality Aspect of Snell's Law

The locality aspect of Snell's Law also has an adverse implication as illustrated by Figure 2. Let region C be the most favorable, region A the next best and region B be the least favorable region. Snell's Law will simply perturb the original, straight line start to goal path until an optimal path involving regions C and B is determined (as represented by the solid line). The law has no ability to consider alternate paths involving region A, even though the optimal path could easily be similar to the dashed line path. Thus, Snell's Law implicitly relies on a straight line hypothesis to propose an initial route. The law ignores more favorable, adjacent areas.

This situation can lead to a form of the same problem that effects the wavefront propagation technique, i.e. combinatorial explosion and computational excess. Deriving a tight bound by applying Snell's Law to a large number of distinct regions requires that each region within some circumscribing limit be specifically examined. The problem is somewhat less serious because the size of each individual area requiring examination is generally much larger than the areas used in a lattice representation. However, the type of tech-

nique we desire must produce a bounding solution efficiently to be valuable. It does not appear that a method which must "look everywhere" can fulfill this requirement.

This indicates that Snell's Law may not be an appropriate formula to apply. However, abandoning the law is not a realistic option. The analysis began with a mathematical, calculus based description of a minimum cost path through a single homogeneous region. The solution to the calculus problem produced Snell's Law as a result. Thus, short of finding another mathematical basis for the minimization problem itself, it seems that any characterization of the minimal cost route will have the same problems as Snell's Law. Another option is to develop a heuristic which allows a more intelligent representation of the problem space so that Snell's Law will be applicable in an efficient manner.

To derive such a heuristic, we again appeal to the human solution model. While humans use the homogeneous regions approach, they seem to rely on a very small number of region groupings. Again refer to Figure 1 where Snell's Law works perfectly. There are only three possible types of regions associated with this situation, optimal, impassable (obstacle), and regions which are traversable, but not at optimal cost. The use of only three region groupings solves the problems stemming from the local nature of Snell's Law. If a path (or path segment) traverses only an optimal region, there is no need to "look" elsewhere for a possibly better solution. If a path traverses a non-optimal region, it is only necessary to examine paths through immediately bordering regions (as well as the original path through the non-optimal region) because the border regions must be optimal. Thus, a ternary representation of the problem space facilitates the application of Snell's Law. This representation also corresponds to a human-like strategy of grouping regions as either the best (optimal), the worst (obstacle), or somewhere in between (traversable, but not at optimal cost).

Thus, we have a general method of combining human-like efficiency with the capability of graph theoretic representations to achieve optimal solutions. The requirement of the human-like planning process is to develop "as good a plan as possible" within the efficiency constraints. The "goodness" of the solution determines the degree of constraint imposed on the efforts required by an optimizing graph search procedure.

The Ternary Case Assumption

The importance of capability knowledge is established in preceding sections. Without such knowledge, homogeneous regions cannot be created. To establish capability knowledge, some specific type of motion-achieving entity must be referenced. There has been a large amount of progress towards solving the extended problem for wheeled vehicles [5]. This being given, we select legged entities as the capability model for this study. Legged, autonomous vehicles have characteristics which greatly differentiate them from wheeled vehicles [9]. First, roadways do not exhaustively define the set of most favorable terrain features for traversability. Clearly, using a human model for illustration, one can walk across an open field as easily as a road. A second characteristic of legged machines is that their total range of attainable speeds is not as great as that of wheeled vehicles. Thus, a collection of terrain features grouped into a single homogeneous region based on a generalized range of attainable speeds contains more members than a region created using wheeled vehicles as the model. Based upon these two general capability traits, as well as the discussions concerning Snell's Law, an assumption pertaining to the creation of homogeneous regions is appropriate. Accordingly, we posit the ternary case assumption: all physical terrain features can be placed into one of three disjoint equivalence classes by a traversability relation based upon legged vehicle capabilities. The first class contains optimum

(per traversability) features such as roads, open fields, non-wooded areas with little elevation change, pastures and trails. The second class contains terrain features which are traversable, but not at an optimum speed. Typical class members are sparsely wooded areas, marsh or rocky areas with marginal footing, areas with moderate elevation changes or slopes and congested areas, such as townships. The final class contains non-traversable features. This class contains such features as rivers, buildings, densely wooded areas, steep mountains and restricted areas. Given that features can be recognized (or provided a priori), an algorithmic means to utilize the ternary case assumption and solve the extended Find-Path problem is presented.

Path Planning Through Elliptical Limiting

A method which combines the advantages of human and graph theoretic models is the desired goal. We have observed that human solutions are efficient while graph theoretic methods produce optimality. This suggests a natural division of the algorithmic organization into two primary parts. The first part relies upon human inspired traits to efficiently produce a possibly non-optimal solution. This solution is used as a bound to constrain the efforts of a second stage process which is designed to ensure optimality. In all further discussion we assume that terrain features have been recognized and grouped according to the ternary case assumption. Conceptually, we assume the existence of a two dimensional plane where nonoptimal, traversable class features and obstacle class features are represented as convex polygons. Optimal class features are not specifically presented but are considered to be the "background" over which the polygons are superimposed. The linear sides of polygonal figures support application of Snell's Law.

The initial step is to perform an opportunistic decomposition of the problem by recognizing corridors through impassable regions. A knowledge based component can perform this function. Each part of this decomposition can be solved independently. An appropriate overall solution is achieved by geometrically linking the subproblem solutions together at corridor entry/exit points. Since each subproblem is independent, we only need to concern ourselves with the solution to a single problem.

The second step solution process resembles a bidirectional A^* search [10] over a dynamically created graph and uses a straight line hypothesis to guide a generate and test methodology. A homogeneous region representation is used to dynamically construct the graph. The difficulties associated with interacting subproblems are avoided by using bidirectional A^* search and selectively remembering nonobstructed links. In the case where only one obstacle lies between the start and goal there is no opportunity for subproblem conflicts to occur. The solution technique utilizes recursion to extend this situation over the entire graph.

First, propose a straight line path from start to goal. If a collision occurs on this path, determine those vertices which facilitate movement around the associated obstacle. Then, treat each obstacle vertex as the current location and generate two new paths from this point, one to the start and one to the goal. Similar actions occur when traversable, but non-optimal areas are intersected. We determine the vertices which facilitate movement around the region and treat these as new, current locations. We also compute the best path through the region in accordance with Snell's Law. This path terminates at the start, the goal, or another region intersection. Eventually, the entire process finds an unobstructed link set to both the start and goal. At this point, the process stores only those links which directly connect an obstacle vertex to either the start or goal. Once this process has been completed, all the "best" links going into the start and into

the goal have been located. We then use the A* estimate function to approximate the cost of completing the middle of each possible path. When the most favorable start and goal link pair has been determined, we treat the interior endpoints of the links as new start and goal locations and recursively apply the same technique. When the actual cost of completing the most favorable partial path is zero, the best "straight line" solution has been found.
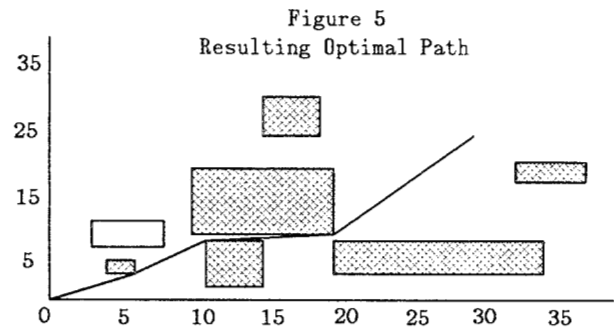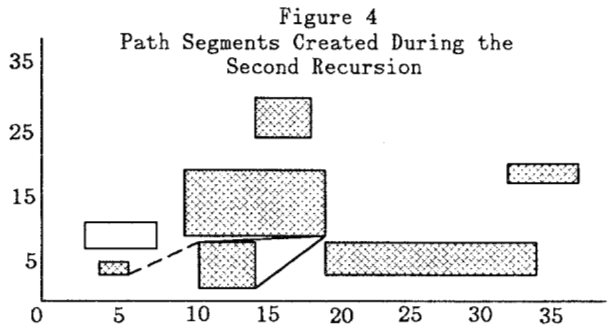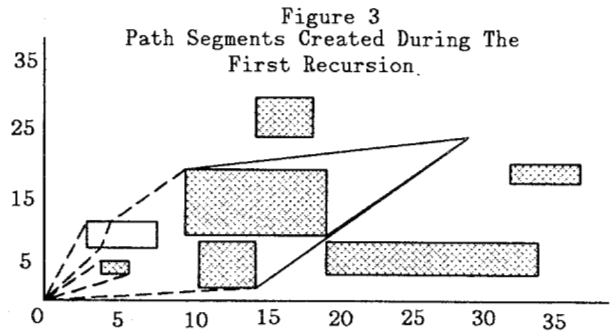
The above process requires the use of the straight line hypothesis. A form of the hypothesis based on the ternary case assumption is available. Clearly, the hypothesis is valid inside a single homogeneous region. This includes the case of going around obstacles where only the leftmost and rightmost (conceptually) vertices need to be located. The difficulty arises when passing between slow and high speed traversable areas. In this case, Snell's law can be applied to find the optimal entry/exit point on the slow/fast region boundary. Once Snell's law has found the pertinent boundary points, they can be added to the dynamically created graph as easily as obstacle vertices.

Figures 3 through 5 illustrate the operation of the second step process. In each figure, the original start is located at the origin. The goal has coordinates (30,25). The filled rectangles represent obstacles. The single, unfilled rectangle represents a non-optimal, traversable region. In Figures 3 and 4, the dashed lines represent unobstructed links to the start while the solid lines depict similar links for the goal. Figure 3 shows the solution state at the end of the first recursion where the start and goal are as listed above. There are 3 links connected to the goal and 5 for the start. Of the start links, two involve the non-optimal region. There is a straight line segment connecting the upper left vertex of this region to the start. This point is static. The second link is a path perturbed by Snell's Law which connects the start to an obstacle vertex at coordinates (10,20). The path (and path cost) is static as are the endpoints. Subsequent paths through the region require new computations. As an example, if it proves desirable to move from the vertex at (4,6) to the vertex at (10,20), a new Snell's Law path must be computed.

At the end of the first recursion, all possible pairs of start links and goal links are matched to determine the most promising pair. Note that two complete paths already exist. However, their costs are greater than the actual cost of the (0,0) to (6,4) link plus the actual cost of the (30,25) to (20,10) link plus the estimated cost of the uncompleted portion of the remaining, middle link. This particular, hypothetical path has the most favorable cost estimate of any link pair. Thus, (6,4) is selected as the new start and (20,10) the new goal to be used in the second recursion (Figure 4). The second recursion finds a complete path between these two points which also happens to have the least cost estimate of any pair of links discovered to that point. Therefore, the process is complete and the path [(0,0), (6,4), (11,9), (20,10), (30,25)] is returned as the optimal solution at the conclusion of two recursive applications (see Figure 5). This second step portion has been implemented in Prolog. The execution times recorded by the method compare very favorably to those of the VGraph technique in the solution of binary case problems.

The result of the second step process is a reasonable, perhaps nonoptimal, path through a region characterized according to the ternary case assumption. The process relied on homogeneous regions, a dynamically created graph, guidance by a form of the straight line hypothesis, planning for interaction conflicts, opportunistic decomposition, recursion to save useful results and intelligent, bidirectional search. The solution has been achieved relatively quickly and wasteful computations have been avoided. It remains to ensure optimality.

The third step of the elliptical limiting method is

Figure 3
Path Segments Created During The First Recursion.

Figure 4
Path Segments Created During the Second Recursion

Figure 5
Resulting Optimal Path

designed to provide an optimal solution. Here, a lower level of abstraction is appropriate. Assume that the homogeneous regions format can be converted (or is otherwise available) to a lattice representation such as that used in the wavefront propagation methods. Given this representation, the solution obtained in the second step can be used as a bound to constrain the efforts required for a wavefront propagation technique. Any optimal solution must have a cost less than or equal to that required by the second step solution. The cost of that solution is known. We compute the greatest distance that can be traveled at equal cost if only the best speed regions are traversed. We then construct an ellipse as discussed previously. Any optimal solution must lie entirely within this ellipse. The number of points interior to the ellipse can be reduced by eliminating any points belonging to the obstacle class as defined in the homogeneous region classification. We then conduct a wavefront propagation over the lattice of remaining points to find the optimal solution.

The third step process results in a set of start to goal paths which all require the same amount of traversal time. The final step in the elliptical limiting process determines the actual, optimal path from this set. Before the ternary grouping can be made, the boundaries of all areal features in the environment must be recognized. Note that in an optimal path, turn points only occur at feature boundaries because within a single region, the least cost path is always a straight line. To find the optimal solution, we superimpose the solutions found in the third step process over the original (unpro-

cessed, analog map feature image) problem representation. Any point in the third step solution set not corresponding to a boundary crossing can be eliminated. Moreover, the accumulation of error present in "stairstep" digital representations of straight, diagonal lines is eliminated.

Figures 6 through 8 illustrate the entire elliptical limiting process on a simple problem which includes only optimal and obstacle class regions. The optimal regions are composed of roadways and open, flat fields. Obstacle class features include a river and several buildings (the checkerboard filled rectangles). Figure 6 illustrates the initial step of the process. The optimal path must cross the bridge over the river. The result is an opportunistic decomposition which includes the formulation of additional start and goal locations. There is a new goal in the lower subproblem and a new start for the upper subproblem. The dashed line through the river represents the decomposition boundary. Figure 7 depicts the second step solution to the lower subproblem. The shaded areas are obstacle regions. The background is the optimal region. This situation is overly simplified by the absence of non-optimal traversable regions. In fact, the absence of such areas results in a simple binary problem for which the second step process will produce the exact, optimal solution. Figure 8 depicts the ellipse produced from the second step solution and the optimal path found by the graph search performed as the third step. The final step requires that the set of optimal paths found in the graph search be overlaid on the original problem representation. The only turn point in the path produced by the graph search which lies on a feature boundary corresponds to the obstacle vertex considered by the second step process. Thus, the optimal path contains the start location, a single turn point, and the goal location. This path is exactly the route depicted in Figure 7, the second step solution.

Summary

The extant extended case solution methods exhibit one or more of several, salient deficiencies. Either they fail to produce optimal solutions, they are computationally excessive, or they require heuristic, ancillary operations to determine the optimal path from a set of solutions. Elliptical limiting minimizes these deficiencies and includes new techniques. In this approach, human-like capabilities are combined with those of a machine. Also, our system includes a ternary problem representation and an efficient method to solve this case. Previously, only "intelligent" solutions to binary case problems or "brute force" techniques for the n-ary case problems have been available. Elliptical limiting features an intelligent, second step strategy usable on uniprocessor architectures and a constrained third step process which can be implemented on multiprocessors. (Propagation techniques have been implemented on parallel machines [12].) The process is also flexible. If timeliness is the paramount concern, only the first two steps need be invoked to produce a quick, reasonable solution. When optimality must be ensured, all steps can be utilized. The process involves mutually cooperating processes; each serves to ameliorate the poor performance aspects of the others. The first two steps are efficient and need not be overly concerned with optimality (much as is the human path planning process). The third step ensures optimality by searching a space which has been greatly pruned by the efforts of the earlier processes. Further, the better the result of the second step, the less effort required in the third step. Finally, the last step determines the true optimal path without resort to heuristic, ancillary operations. Our plans for the future include an implementation of the entire method so that it may be measured against competing processes as well as diverse problem representations.
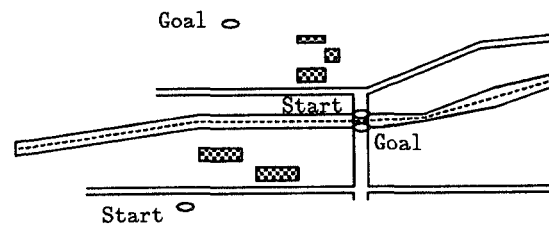
Figure 6
Problem Decomposition
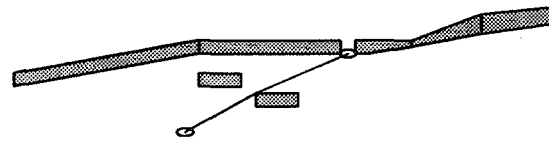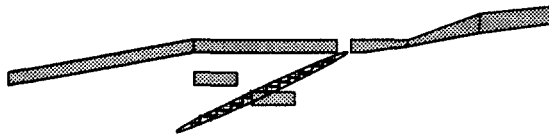

Figure 7
Second Step Solution

Figure 8
Third Step Solution and
Limiting Ellipse

References

[1] Bazara, M. and Jarvis, J., *Linear Programming and Network Flows* , John Wiley and Sons, New York, 1977.

[2] Brooks, R.A., "Solving the Find-Path Problem by Good Representation of Free Space", *IEEE Trans. on Systems, Man, and Cybernetics* , Vol. SMC-13, No. 3, March/April, 1983.

[3] Chavez, R. and Meystel, A., "A Structure of Intelligence for an Autonomous Vehicle", *IEEE International Conference on Robotics* . 1984.

[4] Finney, R. and Thomas, G., *Calculus and Analytic Geometry, 5th Edition* , Addison-Wesley, Reading, 1981.

[5] Kiersey, D. and Mitchell, J., "Planning Strategic Paths Through Variable Terrain Data", *SPIE, Vol. 485, Applications of Artificial Intelligence* , 1984.

[6] Kiersey, D., Mitchell, J., Payton, D. and Preyss, E., "Path Planning for Autonomous Vehicles", *SPIE, Vol. 485, Applications of Artificial Intelligence* , 1984.

[7] Lozano-Perez, T., "Spatial Planning: A Configuration Space Approach", *IEEE Transactions on Computers* , Vol. C-32, No. 2, Feb., 1983.

[8] Lozano-Perez, T. and Wesley, M., "An Algorithm for Planning Collision Free Paths Among Polyhedral Obstacles", *CACM* , Vol. 22, N0. 10, Oct., 1979.

[9] McGhee, R. and Iswandi, G., "Adaptive Locomotion of a Multilegged Robot Over Rough Terrain", *IEEE Transactions on Systems, Man, and Cybernetics* , April, 1979.

[10] Pearl, J., *Heuristics, Intelligent Search Strategies for Computer Problem Solving* , Addison-Wesley, Reading, 1984.

[11] Thorpe, C., "Path Relaxation: Path Planning for a Mobile Robot", *Proceedings of the National Conference on Artificial Intelligence* , Austin, Texas, Aug., 1984.

[12] Witkowski, C., "A Parallel Processor Algorithm for Robot Route Planning", *Proceedings of IJCAI-8* , Karlsruhe, FRG, August, 1983.