

A. Martelli and U. Montanari
 Istituto di Elaborazione dell'Informazione
 Consiglio Nazionale delle Ricerche
 Pisa, Italy

Abstract

In this paper we approach, using artificial intelligence methods, the problem of finding a minimal-cost path in a functionally weighted graph, i.e. a graph with monotone cost functions associated with the arcs. This problem is important since solving any system of functional equations in a general dynamic programming formulation can be shown equivalent to it. A general heuristic search algorithm with estimate is given, which is a nontrivial extension of algorithm A* by Hart, Nilsson and Raphael. Putting some constraints on cost functions and on the estimate, this algorithm can be simplified until the classical version, with additive cost functions, is reached.

1. Introduction

Dynamic programming is a well-known methodology for representing optimization problems in terms of the so-called functional equations and for solving them. While in the artificial intelligence literature dynamic programming is often described reductively as a static, breadth-first technique for searching cycle-free graphs, a general model for the discrete, deterministic case has been given by Karp and Held [1]. This model, called *sequential decision process* consists essentially of a finite automaton with *cost functions* associated with the transitions. If the cost functions are monotone, and if a limit condition is satisfied, searching the transition graph of the finite automaton is equivalent to solving a suitable set of functional equations. The authors have shown how both the functional equations model and the Karp and Held model can be derived from a quite primitive problem-reduction schema [2,3] which directly interprets the Bellman's concept of *problem embedding* and shows how the *optimality principle* holds if and only if the cost functions are monotone.

In this paper, we approach the problem of searching a functionally weighted graph with methods proper of artificial intelligence, in particular we extend the heuristically guided search method by Hart, Nilsson and Raphael [4]. Since here the cost of an arc is represented by a generic monotone function, the extension is nontrivial, and it may happen that no *simple* or no *finite* optimal path exists. Putting further constraints on the cost functions and on the estimate (strict monotonicity, positive monotonicity) it is then possible to simplify the general search algorithm until the classical version is reached.

The relevance of bridging the gap between dynamic programming and heuristic search methods is twofold. On one hand, the Karp and Held model provides a framework where search problems can be stated in quite general terms and (heuristic) search algorithms can be proved effective under well specified constraints. On the other hand, if good estimates can be derived from the problem domain, the solution of problems amenable to a dynamic programming representation can be greatly sped up.

2. Dynamic Programming

Dynamic programming [5] is an optimization technique based on the representation of a problem as a process which evolves from state to state. When the cost structure is appropriate, the determination of an optimal solution may be reduced to the solution of a system of *functional equations*. A general form of such equations is the following [1]

$$(2.1) \quad y_i \leq \min_{j=1, \dots, n} (\min_{j=1, \dots, n} f_{ij}(y_j), a_i) \quad i = 1, \dots, n$$

where variables y_i can assume any value in the complete lattice*

* The presentation of functional equations given here is based on a model of dynamic programming developed by the authors in [2].

$L = \mathbb{R} \cup \{+\infty\} \cup \{-\infty\}$, consisting of the real numbers together with a bottom element $-\infty$ and a top element $+\infty$; the functions $f_{ij}: L \rightarrow L$ are *monotone* (i.e. $x_1 \leq x_2 \Rightarrow f(x_1) \leq f(x_2)$), continuous from above; the values a_i are real constants.

The solution of system (2.1) is the maximal n -tuple $(\bar{y}_1, \dots, \bar{y}_n)$ satisfying the given inequalities. Although for general functions such a maximum does not exist, it is possible to prove, by using the fixpoint theorem of lattice theory, that in the monotone case the maximum exists and satisfies (2.1) with the equal sign. Furthermore, the fixpoint theory provides an iterative algorithm for computing the solution.

The relationships between dynamic programming and search algorithms can be made clear by introducing a model of dynamic programming, due to Karp and Held [1], called *sequential decision process*. This model is a finite automaton with functions associated to its transitions, but it can also be presented as a graph (the transition graph of the automaton).

We have a directed graph G with $n+1$ nodes $N_1, \dots, N_n, N_{n+1}$. A *monotone* function continuous from above $f_{ij}(x)$, called *cost function*, is associated with each arc (N_i, N_j) . The node N_{n+1} , called the *solved node*, has no outgoing arcs and has a constant c_{n+1} (its *cost*) associated with it. The graph G with cost functions is called a *functionally weighted graph*. An example is given in Fig. 1.




Fig. 1

Given a functionally weighted graph G , we can obtain from it a system of functional equations by associating a variable y_i to each node N_i ($i = 1, \dots, n$) and by writing a system of the type (2.1) where the functions f_{ij} are the cost functions of the arcs (N_i, N_j) in G . Furthermore, we put $a_i = f_{i, n+1}(c_{n+1})$, for each i such that arc (N_i, N_{n+1}) exists. Conversely, we can obtain a functionally weighted graph from a system of functional equations. For instance, the functionally weighted graph of Fig. 1 corresponds to the following system of functional equations, if $a = f_4(c_4)$:

$$y_1 = \min(f_1(y_1), f_2(y_2))$$

$$y_2 = \min(f_3(y_1), a)$$

Given a path in a functionally weighted graph from node N_i to the solved node N_{n+1} , its *cost* is obtained by applying in the reverse order to constant c_{n+1} the cost functions of the arcs of the path. The next theorem shows how the solution of the functional equations is related to the costs of the paths in the corresponding functionally weighted graph. (See [3] for a proof).

Theorem 2.1. The solution $\bar{y}_i$ of functional equations for node N_i is equal to the greatest lower bound (glb) of the costs of all paths in the

functionally weighted graph from node N_i to the solved node N_{n+1} .

This theorem shows that a dynamic programming problem can be reduced to a minimal-cost path problem, if the cost of a path is suitably defined. Therefore such a problem could be solved by using algorithms for finding minimal-cost paths in a functionally weighted graph. These algorithms are well-known in the case of additive costs, i.e. where the cost functions have the form $f(x) = x + a$ [4,6]. In the next section we give a general search algorithm which uses heuristic information to guide the search.

Conversely, this theorem says that a minimal-cost path problem can be solved by solving a system of functional equations with, for example, the iterative algorithm. Such a result is well-known in the additive case, and several iterative algorithms have been given for finding the shortest path in a graph [7].

We point out that the model of dynamic programming presented in this section is general as far as *monadic* functions are considered. However there are dynamic programming problems which are better modelled by assuming that the functional equations contain *polyadic* functions [8]. In this case our model can be generalized to an AND/OR graph [9].

3. A General Propagation Algorithm with Estimate

In the previous section we have described a general model for dynamic programming and an iterative solution algorithm. No better method is available, unless special assumptions are made on the functions f_{ij} and/or an *estimate* is supplied for the cost of all intermediate problems based on information from the problem domain. This line of development corresponds to merging the branch-and-bound idea with dynamic programming and has been originally pursued by Hart, Nilsson and Raphael [4] in the special case of additive functions with positive weights. Here we give a general algorithm for searching functionally weighted graphs and we show its validity when the estimates are lower bounds.

Let G be a functionally weighted graph with $n+1$ nodes $N_1, \dots, N_n, N_{n+1}$ and *monotone* cost functions f_{ij} . Given a path p from node N_{i1} to node N_{ik} ($p = N_{i1} N_{i2} \dots N_{ik}$) we define the cost function of path p as the (monotone) function $f_p(x)$ obtained by composing the cost functions of the arcs of p :

$$f_p(x) = f_{i1,i2}(f_{i2,i3}(\dots(f_{ik-1,ik}(x)) \dots))$$

Furthermore we define the *distance function* $d_{ij}(x)$ from node N_i to node N_j as

$$d_{ij}(x) = \text{glb}_{p \in P_{ij}} f_p(x) \quad \text{for every } x$$

where P_{ij} is the (possibly infinite) set of paths from N_i to N_j . Notice that the distance is defined, for every x , as the greatest lower bound and not simply as the minimum, because, in some cases, it can only be obtained as a limit, i.e. with an infinite path.

Finally we can define the *solution cost* c_i of a node N_i ($i=1, \dots, n$) as the glb of all values which can be achieved by applying the cost function of every path from N_i to N_{n+1} to the cost c_{n+1} of the solved node N_{n+1} :

$$c_i = d_{i,n+1}(c_{n+1})$$

We give now an algorithm for finding the solution cost c_i of a given node N_i of G , called the *top node*, and a path for which this cost is achieved, called *solution path*. Note that, when the solution cost c_i is obtained as a limit, no (finite) solution path exists. This algorithm is a *bottom-up* propagation algorithm, i.e. it propagates toward the top node starting with the solved node N_{n+1} , which, from now on, will also be called *bottom node* and denoted with N_0 . No restriction is made on the generality of the algorithm by having only one bottom node and one top node. In fact, if more than one top node were present, we could add to the graph a new node N_0 connected to each top node through an arc with the identity function associated with it, and we could consider N_0 as the top node of the new graph.

* Usually, in the artificial intelligence literature, search algorithms are given for infinite graphs. Most results of this paper can be extended to this case.

At each step, the algorithm selects a node N_i and expands it generating all of its *parents*, i.e. all nodes N_j such that (N_j, N_i) is an arc of G . The selection of the node to be expanded is made according to a *total estimate* e_i associated with each node N_i , which is an estimate of the solution cost c_i of node N_i , with the constraint that the solution path passes through node N_i . The total estimate e_i can be expressed as follows

$$e_i = h_i(g_i)$$

where g_i is an estimate of the solution cost c_i of node N_i and $h_i(x)$ is a *monotone* function which gives an estimate of the distance function $c_{ii}(x)$ from N_i to N_i . This estimate can be obtained using the heuristic information available from the problem domain, whereas g_i will be constructed step by step by the algorithm.

Our algorithm is an extension to the case of monotone cost functions of algorithm A^* by Hart, Nilsson and Raphael [4,10], which considers the case of additive cost functions with positive weights, i.e.

$$f_{ij}(x) = x + c_{ij} \quad (c_{ij} > 0)$$

for every arc (N_i, N_j) . Furthermore, algorithm A^* assumes that the bottom node N_0 has a cost $c_0 = 0$ and that the estimate $h_i(x)$ is also an additive function, i.e. $h_i(x) = x + h_i$. Thus, the total estimate becomes, in this case, $e_i = g_i + h_i$.

Since our algorithm deals with the general case of monotone cost functions, the extension is nontrivial and some new features have to be added. First of all we point out that the solution path can contain the top node N_i more than once. This happens when the distance function from the top node to itself is smaller than the identity function ($d_{ii}(x) < x$, for some x), i.e. the top node belongs to a cycle which can reduce its cost. This fact is taken into account by adding a new node N_0 to the graph and an arc (N_0, N_i) with the identity function $I(x)$ associated with it. The algorithm will stop only when this node will be encountered.

Furthermore, a solution path can contain the same node several times. Therefore, if we want to obtain the solution path by tracing back through the pointers from the top node, we must have more than one pointer for each node. This is achieved by associating with each node a *stack* of pointers, to which a new element is added every time the node is reopened.

Algorithm SEARCH

- 1) Add a new node N_0 to the graph and an arc (N_0, N_i) . Let the cost function of this arc be $f_{0i}(x) = I(x) = x$ and $h_0(x) = I(x)$.
- 2) Put the bottom node N_0 on a list called OPEN and let $g_0 \leftarrow c_0$ and $e_0 \leftarrow h_0(c_0)$.
- 3) If OPEN is empty, exit with failure; otherwise continue.
- 4) Remove from OPEN that node N_i whose total estimate e_i is smallest and put it on a list called CLOSED. In case of tie on the e values, choose the node N_i whose g_i value is smallest. Resolve further ties arbitrarily, but always in favor of node N_0 .
- 5) If N_i is node N_0 , exit with the solution path obtained by tracing back through the pointers; otherwise continue.
- 6) Expand N_i , generating all of its parents. For each parent N_j compute $g'_j \leftarrow f_{ji}(g_i)$.
- 7) Associate with the parents N_j not already on either OPEN or CLOSED the values $g_j \leftarrow g'_j$ and $e_j \leftarrow h_j(g'_j)$. Put these nodes on OPEN and put on top of their stacks a pointer back to the top of the stack of N_i .
- 8) If a parent N_j is on OPEN and $g'_j < g_j$ let

$$g_j \leftarrow g'_j \quad \text{and} \quad e_j \leftarrow h_j(g'_j).$$

Replace the pointer on top of the stack of N_j with a pointer to the top of the stack of N_i .

- 9) If a parent N_j is on CLOSED and $g'_j < g_j$, put it on OPEN and let

$$g_j \leftarrow g'_j \quad \text{and} \quad e_j \leftarrow h_j(g'_j).$$

Add on top of the stack of N_j a pointer to the top of the stack of N_i .

- 10) Go to 3.

It is possible to prove that the above algorithm is *admissible* (i.e. if it terminates cleanly it finds an optimal solution), if the estimate $h_i(x)$ is a lower bound on the distance $d_{ti}(x)$ for each node ($h_i(x) \leq d_{ti}(x)$) for each x and for each i). The proof is contained in the Appendix.

The proof is based on a lemma which states that, if the estimate $h_i(x)$ is a lower bound on $d_{ti}(x)$, then at any iteration of algorithm SEARCH there is an open node N_j with $e_j \leq c_i$. This result is not trivial since the solution cost c_i might be achieved only as a limit, and thus infinite paths have to be considered.

Notice that, in general, there is no way of deciding whether the algorithm will stop. Special classes of functions and estimates for which such a decision can be made will be presented later on.

As a further remark we point out that the bottom-up structure of our algorithm is the most natural choice for the general case. In fact, when no estimate is available, with our algorithm we propagate costs, i.e. constants, from the solved node toward the top node, whereas, by using a top-down approach, we should propagate functions. Such a difference between bottom-up and top-down is not present in the case of additive cost functions, where functions can be propagated very easily. In fact, bidirectional search algorithms have been given for this case [12].

Algorithm SEARCH can be easily extended to the case where the estimate $h_i(x)$ is not a lower bound on $d_{ti}(x)$, but the maximal amount by which $h_i(x)$ can exceed $d_{ti}(x)$ is known. This case has been studied in [13] for additive cost functions.

Assume that a strictly monotone* error function $b(x)$ is given such that $h_i(x) \leq b(d_{ti}(x))$ for every x and for every node. Algorithm SEARCH can be modified by associating with the arc (N_0, N_i) the cost function $b(x)$. It is easy to see that the new algorithm finds the solution cost of node N_i . In fact, the distance function from N_0 to every node N_i is $d_{0i}(x) = b(d_{ti}(x))$ and the estimate $h_i(x)$ can be considered as a lower bound on $d_{0i}(x)$. Therefore the algorithm finds the solution cost of node N_0 , but, since every solution path for N_0 is a solution path for N_i as well, for the strict monotonicity of b , it finds also the solution cost of N_i .

In order to see how algorithm SEARCH works, let us consider the following problem. A functionally weighted graph is given with $n+1$ nodes, where each node N_i ($i = 1, \dots, n$) corresponds to a different currency C_i . A directed arc (N_i, N_j) denotes the possibility of exchanging currency C_i into C_j . To each arc (N_i, N_j) ($i, j = 1, \dots, n$) we associate a cost function $f_{ij}(x) = a_{ij}x + b_{ij}$ ($a_{ij} > 0, b_{ij} > 0$), which means that, if we want to have the amount x of currency C_j , we have to change an amount of currency C_i proportional to x plus a fixed amount for taxes. Furthermore, the solved node N_{n+1} corresponds to some goods and, for each arc (N_i, N_{n+1}) , the value $f_{i,n+1}(c_{n+1})$ gives the cost of these goods in the currency C_i .

Our problem is that of buying the given goods with the smallest amount of currency C_1 . Notice that, if we are quite optimistic, we can hope of finding a cycle which gives us some gain, and thus the solution path may contain some cycles.

An example is given in Fig. 2. The operator/in the cost functions denotes integer division. The given estimates imply that very little heuristic information is assumed. However it is clear that they are lower bounds if the domain of the cost functions are the nonnegative integers.

In Fig. 3 we give some steps of the execution of algorithm SEARCH. In Fig. 3a we show the stacks of pointers after node N_2 has been expanded for the first time. Each element of a stack has a number associated with a pointer. This number is the g value of the node, and it gives the cost of the path obtained by following the pointers. Furthermore, the encircled numbers are the total estimate e of the open nodes. Notice that every time a node is reopened, a new element is added to the stack of pointers.

Finally, in Fig. 3b we give the last step of the algorithm. We see that the solution cost of node N_1 is 8 and the solution path, obtained by tracing through the pointers, is $N_1 N_2 N_3 N_2 N_3 N_5$. Thus we have obtained a solution path with a cycle.

We want to see now how algorithm SEARCH behaves when no estimate is available. This means that, if we do not put any restriction on the cost functions, we must have $h_i(x) = -\infty$ for each node. At each step




Fig. 2




Fig. 3a




Fig. 3b

* A function f is strictly monotone if $x_1 < x_2 \Rightarrow f(x_1) < f(x_2)$.

of the algorithm, each node N_i has a total estimate $e_i = -\infty$, except for node N_0 for which we have $e_0 = g_0$. According to the condition on ties in step 4 of the algorithm, the node whose g_i value is smallest is expanded at each step. Since N_0 is the only node with total estimate $\neq -\infty$, the algorithm will stop only after all other nodes have been closed. In this case it is clear that the solution cost of *all* nodes is determined simultaneously. Thus the estimate $h_i(x) = -\infty$ for each node can be always assumed when the above more general problem must be solved.

Note that this algorithm may be used in the case of additive cost functions, if negative weights are allowed and no estimate is available.

4. Positively Monotone Cost Functions, and Consistent and Strictly Monotone Estimates

The first case we examine is when all cost functions are positively monotone. A monotone function is called *positively monotone* iff for all arguments x we have

$$f(x) \geq x$$

Examples of positively monotone functions are as follows

$$\begin{aligned} f(x) &= x + c & c &\geq 0 \\ f(x) &= ax & a &\geq 1 \quad x \geq 0 \\ f(x) &= \max(x, c) \end{aligned}$$

It is easy to see that the composition of two positively monotone functions is still a positively monotone function, and thus the cost function of every path is also a positively monotone function. As a consequence, the estimates can also be assumed positively monotone without restrictions.

For positively monotone cost functions we have the following theorem, (For the proofs of all theorems in this section, see the Appendix).

Theorem 4.1. Let G be a functionally weighted graph with positively monotone cost functions. Then there always exists a simple solution path from the top node.

From the above theorem it follows that algorithm SEARCH can take into account only temporary solution paths without cycles. Thus the algorithm can be simplified by eliminating the stacks of pointers and by replacing them with only one pointer for each node (as in algorithm A^*).

When the estimate h is strictly monotone, looking at g values in step 4 of algorithm SEARCH for solving ties is no longer useful since $h_i(g_i) = h_j(g_j)$ implies $g_i = g_j$. Thus in this case algorithm SEARCH behaves exactly as A^* .

Note that positive monotonicity alone, while insuring termination, does not guarantee that all closed nodes are never reopened. This case is well-known in the special, classical case of a graph weighted with positive constants, with estimate [4]. Following [4], we thus introduce consistent estimates. An estimate $h_k(x)$, $k = 1, \dots, n$ is *consistent* iff for all arcs (N_i, N_j) and for all x we have

$$h_i(x) < h_j(f_{ji}(x))$$

In the case of consistent estimate, but not necessarily for positively monotone cost functions, we can prove the following theorem.

Theorem 4.2. Let G be a functionally weighted graph with consistent estimate $h_k(x)$, $k = 1, \dots, n$. Then the values of the total estimate e_i of the nodes closed in the successive iterations of algorithm SEARCH form a (possibly infinite) nondecreasing sequence.

The above result does not mean that closed nodes are never reopened. In fact a node N_i could be closed a second time with the same estimate e_i but with an improved cost g_i . (Since $\bar{g}_i < g_i$, otherwise N_i would not be reopened, the monotonicity of h_i does not allow $\bar{e}_i > e_i$).

For instance, let us consider the graph in Fig. 4. The exact distances are as follows:

$$\begin{aligned} d_{tt}(x) &= \min(x, 4) \\ d_{ta}(x) &= \min(x+4, 4) \\ d_{tb}(x) &= \min(x+3, 4) \end{aligned}$$




Fig. 4

Let us take as estimate the following functions

$$\begin{aligned} h_t(x) &= \min(x-1, 5/2) \\ h_a(x) &= \min(x+3, 5/2) \\ h_b(x) &= \min(x+1, 2) \end{aligned}$$

It is easy to see that this estimate is consistent. For instance, for the loop at node N_a we have

$$\min(x+3, 5/2) \leq \min(x/2+3, 5/2)$$

However algorithm SEARCH does not stop, even if a solution path exists, and reopens node N_a infinitely many times. The values of the costs and of the total estimates for nodes closed in successive iterations are as follows:

- 1) node N_b $g_b = 0$ $e_b = 1$
- 2) node N_t $g_t = 3$ $e_t = 2$
- 3) node N_a $g_a = 1$ $e_a = 5/2$ we do not stop since $e_0 = g_0 = 3 > e_a = 5/2$
- 4) node N_a $g_a = 1/2$ $e_a = 5/2$ "
- 5) node N_a $g_a = 1/4$ $e_a = 5/2$ "

and so on.

The following result allows to characterize the set of closed nodes which could possibly be reopened.

Theorem 4.3. Let G be a functionally weighted graph with consistent estimate $h_k(x)$, $k = 1, \dots, n$. At some stage of algorithm SEARCH, let N_i be a closed node with cost g_i and total estimate e_i *smaller* than the total estimate e_j of all open nodes. Then node N_i will never be reopened, and its solution cost is g_i .

If the estimate is consistent, the above theorem allows to add a new termination condition to algorithm SEARCH: "If the top node in closed and has a total estimate *smaller* than the node currently under closure, stop with success". This condition allows for instance to stop at the third iteration of the previous example.

If the two conditions of positive monotonicity of the cost functions and of consistency of the estimate are taken jointly, we can prove the following theorem, which is an extension of a result by Hart, Nilsson and Raphael [4].

Theorem 4.4. Let G be a functionally weighted graph with positively monotone cost functions and with a consistent estimate. Then when algorithm SEARCH closes a node N_i , its cost g_i is the solution cost of N_i . Thus algorithm SEARCH never reopens a node.

The above theorem provides an obvious termination condition for algorithm SEARCH: "Stop when the top node is closed".

The above results apply, as pointed out before, to the well-known

case of additive cost functions. However their generality allows also the application to cost functions of different form, where the standard A* algorithm cannot be used. An example is provided by the so called *secondary optimization problem* which arises when we want to find the best ordering of variable elimination in non-serial dynamic programming [14]. The problem can be stated as that of finding the minimal cost path in a weighted graph, but where the cost of a path is the maximum among all the weights of its arcs. Therefore the cost functions of this graph are of the form

$$f_{ij}(x) = \max(x, c_{ij})$$

i.e. they are positively monotone functions and algorithm SEARCH can be used for finding the solution path, if suitable estimates are provided*.

A further example of positively monotone but nonadditive cost functions is the so called *1-dimensional stock cutting problem* [1], where we want to determine the minimum number of standard pieces of material of length L from which it is possible to cut a_1 pieces of length l_1 , a_2 pieces of length l_2 , ..., and a_n pieces of length l_n . In the graph representation of this problem, the meaning of an arc is "cut a piece of length l_j " and the cost function associated with it is

$$f(x) = \begin{cases} x + l_j & \text{if } x + l_j \leq L \lceil x/L \rceil \\ L \lceil x/L \rceil + l_j & \text{otherwise} \end{cases}$$

The second case in the definition of $f(x)$ corresponds to the situation in which the current standard piece does not have enough material remaining for the next cut, and a new standard piece is required.

According to Theorem 4.2 and 4.3, if the estimate is consistent, then a closed node N_i can be reopened only with the same total estimate $\bar{e}_i = e_i$ and with a better cost $\bar{g}_i < g_i$. Thus it is clear that, if the estimate is strictly monotone, then no closed node can be reopened. Thus the same result proved in Theorem 4.4 holds also in this case.

As a conclusion, we point out that in this paper we have given algorithm SEARCH for searching functionally weighted graphs which is an extension of algorithm A*, and we have shown that, for generic monotone functions, this algorithm requires some new features such as the addition of node N_0 , the stack of pointers and the tie handling procedure in Step 4. Furthermore, we have shown that, by considering restricted classes of cost functions and estimates, simpler versions of algorithm SEARCH can be given. In particular, the validity of the trivial generalization of A* is extended either to the case of strictly monotone estimates and positively monotone cost functions or to the case of strictly monotone and consistent estimates. Properties of algorithm SEARCH for another special class of cost functions (strictly monotone cost functions) are described and proved in [11].

Appendix

Lemma 3.1. If $h_i(x) \leq d_{ii}(x)$ for every x and for every node, then at any step of the algorithm and for any path p from N_0 to N_b , there exists an open node N_j on p with $e_j \leq f_p(c_b)$.

Proof. Let N_j be the first open node we encounter by following path p backwards from N_b to N_0 . There must be such a node because otherwise N_0 is closed and the algorithm has terminated. Let p' be the part of path p from N_j to N_b . Since all nodes on p' are closed, we must have

$$g_j \leq f_{p'}(c_b).$$

Since the distance function $d_{ij}(x)$ is a monotone function and since, in general, p is not a solution path, we have

$$d_{ij}(g_j) \leq d_{ij}(f_{p'}(c_b)) \leq f_p(c_b)$$

Finally, since $h_j(x)$ is a lower bound on $d_{ij}(x)$, we have

$$e_j = h_j(g_j) \leq d_{ij}(g_j) \leq f_p(c_b) \quad \text{Q.E.D.}$$

Lemma 3.2. If $h_i(x) \leq d_{ii}(x)$ for every x and for every node, then at any step of the algorithm there exists an open node N_j with $e_j \leq c_t$.

Proof. If c_t is achieved with a finite path p , then $c_t = f_p(c_b)$ and the lemma is proved by Lemma 3.1. Assume now that c_t is obtained as a limit, and let $p_1, p_2, p_3, \dots$ be an infinite sequence of paths from N_0 to N_b such that

$$f_{p_1}(c_b) \geq f_{p_2}(c_b) \geq f_{p_3}(c_b) \geq \dots \geq c_t$$

$$\text{and } \lim_{i \rightarrow \infty} f_{p_i}(c_b) = c_t.$$

Let us apply now the construction of Lemma 3.1 to each path of the sequence and let us consider the subsequences of paths for which the same node is obtained. Among these subsequences, there must be an infinite subsequence $p_{j_1}, p_{j_2}, p_{j_3}, \dots$ such that

$$\lim_{i \rightarrow \infty} f_{p_{j_i}}(c_b) = c_t.$$

Let N_j be the node obtained with the construction of Lemma 3.1 from these paths. We have

References

1. Karp, R.M. and Held, M., Finite-state processes and dynamic programming, *SIAM J. Appl. Math.*, 15, 1967, pp. 693-718.
2. Martelli, A. and Montanari, U., On the foundations of dynamic programming, N.I. B73-11, I.E.I., Pisa, Oct. 1973, to appear in *Topics in Combinatorial Optimization*, S. Rinaldi ed., CISM, Springer Verlag.
3. Martelli, A. and Montanari, U., Dynamic programming schemata, *Proc. second Coll. on Automata, Languages and Programming*, Saarbrücken, July 1974, Springer Verlag, pp. 66-80.
4. Hart, P., Nilsson, N. and Raphael, B., A formal basis for the heuristic determination of minimum cost paths, *IEEE Trans. on Sys. Sci. Cybernetics*, SSC-4, N. 2, July 1968, pp. 100-107.
5. Bellman, R.E., *Dynamic Programming*, Princeton University Press, Princeton N.J. 1957.
6. Dijkstra, E., A note on two problems in connection with graphs, *Numerische Math.*, 1, 1959, pp. 269-271.
7. Johnson, D.B., Algorithms for shortest paths, Ph. D. Thesis, Cornell University, Ithaca, N.Y., May 1973.
8. Martelli, A. and Montanari, U., Additive AND/OR graphs, *Proc Third Int. Joint Conf. on Art. Int.*, Stanford, 1973, pp. 1-11.
9. Martelli, A. and Montanari, U., Programmazione dinamica e punto fisso, *Atti Conv. di Informatica Teorica*, Mantova, Nov. 1974, pp. 1-19 (in Italian).
10. Nilsson, N.J., *Problem-solving methods in artificial intelligence*, McGraw-Hill, N.Y., 1971.
11. Martelli, A. and Montanari, U., From dynamic programming to search algorithms with functional costs, IEL Internal Report B75-1, January 1975.
12. Pohl, I., Bi-directional and heuristic search in path problems, Ph.D. Thesis, Stanford University, Stanford, 1969.
13. Harris, L.R., The heuristic search under conditions of error, *Artificial Intelligence*, 5, N. 3, Fall 1974, pp. 217-234.
14. Brioschi, F. and Even, S., Minimizing the number of operations in certain discrete-variable optimization problems, *Operations Research*, 18, N. 1, 1970, pp. 66-81.

* Note that if estimates of the same type $h_i(x) = \max(x, c)$ are considered, the tie handling procedure in Step 4 of algorithm SEARCH is essential for Theorem 4.4 to hold. The same happens whenever the estimates are not strictly monotone.

$$e_j \leq f_{p_{ji}}(c_b) \quad i = 1, 2, 3, \dots$$

Therefore we have

$$e_j \leq c_t \quad \text{Q.E.D.}$$

Theorem 3.1. If $h_i(x) \leq d_{ti}(x)$ for every x and for every node, and if the algorithm terminates cleanly (i.e. without failure), then it finds the solution cost c_t of node N_t .

Proof. When the algorithm stops we have $e_0 = h_0(g_0) = f_{0t}(g_t) = g_t$. Assume now that e_0 is not the solution cost of N_t , i.e. $e_0 = g_t > c_t$. By lemma 3.2 we know that there existed just before termination an open node N_j with $e_j \leq c_t$. Thus, N_j , which must be different from N_0 , would have been selected for expansion rather than N_0 , contradicting the hypothesis that the algorithm terminates. Q.E.D.

Theorem 4.1. Let G be a functionally weighted graph with positively monotone cost functions. Then there always exists a simple solution path from the top node.

Proof. Given a nonsimple path P_{tb} of cost

$$P_{tb}(c_b) = p_{ti}(p_{ij}(p_{jb}(c_b)))$$

there always exists a simple path P'_{tb} of not larger cost

$$P'_{tb}(c_b) = p_{ti}(p_{ib}(c_b))$$

In fact, $p_{ii}(x)$, being the composition of positively monotone functions, is positively monotone and thus

$$P_{ib}(c_b) \leq p_{ii}(P_{ib}(c_b))$$

But $p_{ti}(x)$ is monotone, so we have

$$P'_{tb}(c_b) = p_{ti}(P_{ib}(c_b)) \leq p_{ti}(p_{ii}(P_{ib}(c_b))) = P_{tb}(c_b).$$

As a consequence, any path of minimum cost in the finite set of simple paths is optimal. Q.E.D.

Theorem 4.2. Let G be a functionally weighted graph with consistent estimate $h_k(x)$, $k = 1, \dots, n$. Then the values of the total estimate e_i of the nodes closed in the successive iterations of algorithm SEARCH form a (possibly infinite) nondecreasing sequence.

Proof. When a node N_i is closed by algorithm SEARCH its total estimate $e_i = h_i(g_i)$ by construction is not larger than the total estimates e_k of all open nodes. In the updating phase, the total estimate of some node N_j , adjacent to N_i , may be decreased as follows:

$$\bar{e}_j = h_j(f_{ji}(g_i))$$

and, if N_j was closed, it will be reopened. However, the new value $\bar{e}_j$ is not smaller than e_i by the consistency assumption. Thus, the total estimate of the next closed node, which is the minimal among the total estimates of all open nodes after updating, will not be smaller than e_i . Q.E.D.

Theorem 4.3. Let G be a functionally weighted graph with consistent estimate $h_k(x)$, $k = 1, \dots, n$. At some stage of algorithm SEARCH, let N_i be a closed node with cost g_i and total estimate e_i smaller than the total estimate e_j of all open nodes. Then node N_i will never be reopened, and its solution cost is g_i .

Proof. The cost g_i is clearly not worse than the costs of all paths visiting only closed nodes. Let now P be any path from N_i to N_b which passes through at least one open node, and let N_j be the last open node on P . We have

$$P_{ib}(c_b) = p_{ij}(p_{jb}(c_b))$$

and

$$P_{jb}(c_b) \geq g_j.$$

From the consistency assumption, extended to the part of P from N_i to N_j we have

$$h_j(g_j) \leq h_i(p_{ij}(g_j)) \leq h_i(p_{ib}(c_b)).$$

But since

$$h_i(g_i) < h_j(g_j)$$

by assumption, we have

$$h_i(g_i) < h_i(p_{ib}(c_b))$$

and finally, since h_i is monotone,

$$g_i < p_{ib}(c_b).$$

Clearly, since g_i is already the optimal cost, node N_i will never be reopened. Q.E.D.

Theorem 4.4. Let G be a functionally weighted graph with positively monotone cost functions and with a consistent estimate. Then when algorithm SEARCH closes a node N_i , its cost g_i is the solution cost of N_i . Thus algorithm SEARCH never reopens a node.

Proof. The proof is quite similar to that of Theorem 4.3. We can write again

$$P_{ib}(c_b) = p_{ij}(p_{jb}(c_b))$$

and

$$P_{jb}(c_b) \geq g_j.$$

Now if

$$h_i(g_i) < h_j(g_j)$$

the result of the proof of Theorem 4.3 is valid and the thesis follows. On the other hand, if

$$h_i(g_i) = h_j(g_j)$$

then we have

$$g_i \leq g_j$$

since N_i was closed before N_j . Thus we have

$$g_i \leq p_{ib}(c_b).$$

But since $p_{ij}(x)$ is now positively monotone we have

$$g_i \leq p_{ij}(p_{ib}(c_b)) = P_{ib}(c_b)$$

Q.E.D.