Annals of Operations Research 40(1992)125–151

FORECAST HORIZONS AND DYNAMIC FACILITY LOCATION PLANNING

Mark S. DASKIN
Department of Civil Engineering, Northwestern University, Evanston, IL 60208, USA

Wallace J. HOPP
Department of Indiustrial Engineering, Northwestern University, Evanston, IL 60208, USA

Benjamin MEDINA
United Airlines, Corporate Research and Development, Executive Offices-EXOEB, P.O. Box 66100, Chicago, IL 60666, USA

Abstract

We consider a dynamic facility location model in which the objective is to find a planning horizon, τ^*, and a first period decision, X_1^*, such that X_1^* is a first period decision for at least one optimal policy for all problems with planning horizons equal to or longer than τ^*. In other words, we seek a planning horizon, τ^*, such that conditions after τ^* do not influence the choice of the optimal initial decision, X_1^*. We call τ^* a *forecast horizon* and X_1^* an *optimal initial decision*. For the dynamic uncapacitated fixed charge location problem, we show that simple conditions exist such that the initial decision depends on the length of the planning horizon. Thus, a strictly optimal forecast horizon and initial policy may not exist. We therefore introduce the concepts of e-optimal forecast horizons and ε-optimal initial solutions. Our computational experience inicates that such solutions can be found for practical problems. Although computing ε-optimal forecast horizons and initial decisions can be cumbersome, this approach offers the potential for making significantly better decisions than those generated by other approaches. To illustrate this, we show that the use of the scenario planning approach can lead to the adoption of the worst possible initial decision under conditions of future uncertainty. On the basis of our results, it appears that the forecast horizon approach offers an attractive tool for making dynamic location decisions.

1. Introduction

The location of facilities in both private and public sector systems critically affects the ability of these systems to deliver the requisite services. Facility location decisions are *long-term strategic decisions* in that they impact on such lower-level shorter-term decisions as resource allocation, vehicle dispatching, and vehicle repositioning [1]. In general, implementing a facility location strategy involves making a series of opening/closing and demand allocation decisions over time.

In the absence of large fixed costs associated with opening, closing and operating most facilities, the optimal policy over time might well be to site facilities optimally in each time period independently of (1) where they were located in prior periods and (2) where they are anticipated to be needed in the future. However, often these costs are large. For example, the start-up costs associated with an automobile assembly plant are likely to be several *billion* dollars before the first finished vehicle rolls off the line. This precludes frequent relocation of facilities and makes it essential that location decisions made today account for anticipated future conditions.

Since future conditions are typically uncertain, and forecasts are frequently unreliable and subject to revision, accounting for future conditions can be very difficult. Furthermore, there is typically no prespecified *exogeneous* time horizon beyond which conditions can be ignored. Given this, and observing that the first period decision is the only one that must be implemented immediately (i.e. without the benefit of future information or analysis), we argue that the objective of dynamic facility location planning should be *to find an optimal or near optimal first period decision for the location problem over an infinitely long planning horizon.*

This definition of the dynamic facility location problem warrants further explication. By a *decision* in a given time period, we generally mean the choice of existing facilities to close and new facilities to open at the beginning of the period and the assignment of demands to facilities during the period. A sequence of decisions, one for each period, will be called a *policy*. An *optimal policy* for either a finite or infinite horizon problem will be a policy whose adoption minimizes the present value of all present and future costs. An *optimal initial decision* is a first period decision that is part of an optimal policy. Our computational approach is to find an *endogenously* determined *forecast horizon* and an *initial decision* such that for all planning horizons of duration greater than or equal to the length of the forecast horizon, an optimal (or near optimal) policy begins with the specified initial decision (see Bean and Smith [3]). If such a forecast horizon can be identified, an optimal initial decision can be computed by solving a sufficiently long finite horizon problem. This forecast horizon approach will prevent the use of too short a planning horizon, which leads to myopic decisions and can easily happen under the common procedure of using an exogeneous planning horizon.

Because most location problems of practical interest are NP-complete, finding a provably optimal solution may require an excessive amount of computational effort for even moderate sized problems. Also, as shown below, there are simple problems for which a finite forecast horizon does not exist if we require the selected initial decision to be *strictly* optimal for all planning horizons of duration greater than or equal to the finite forecast horizon. For these reasons, we will generally concentrate on finding a *near optimal* first period decision.

The remainder of the paper is organized as follows. In section 2, we briefly review the literature on dynamic facility location modelling and forecast horizons. In section 3, we formalize the notion of a forecast horizon and an empirical

ε-optimal forecast horizon for a general location model. We also show via a simple example that a strictly optimal initial decision may not exist for any finite forecast horizon. This, together with the NP-completeness of most location problems, justifies our search for near optimal initial decisions and forecast horizons. Section 4 summarizes a series of computational experiments designed to test whether or not empirical ε-optimal forecast horizons and initial solutions can be found for practical problems. In section 5, we briefly review the scenario planning approach for dealing with uncertainty and show via a small example that the use of this approach may lead to the adoption of the *worst* possible initial decision. In section 6, we summarize our results and present recommendations for future work.

2. Review of related literature

Despite the strategic and long term nature of facility location decisions, most location models adopt a static approach. This is evident from the virtual absence of any discussion of dynamic models in recent texts on location modelling [12, 16, 21, 23, 25]. Adoption of a static approach simplifies the problem, but is justified only if the future is expected to replicate the present. In most cases, future conditions will differ from present conditions, and the use of dynamic models is warranted.

Multi-period or dynamic facility location problems may be thought of as extensions of static models in which a temporal dimension has been appended. In all such models of which we are aware, the length of the planning period is an exogenous input and the model seeks an optimal policy for the planning period. Early work in this area was performed by Ballou [2] in the context of locating and relocating a single warehouse over a period of years to maximize the net present value of a profit stream that depended on the location of the warehouse. He employed dynamic programming with backward recursion. The procedure he proposes must be viewed as heuristic in that the set of candidate sites is limited to the set of sites which correspond to solutions to the static problems for each period. This restricts the size of the state space. As noted by Ballou, however, "The danger in this assumption, of course, is that some other set of alternatives might possibly yield higher profits (p.275)."

Erlenkotter [9] reviews a variety of other heuristic procedures for dynamic facility sizing/location/allocation problems. In these problems, the objective is to determine when and where to add capacity so that demands in all periods may be satisfied and so that the total cost of all capacity additions and distribution costs are minimized.

Optimization based approaches to the problem have been developed by a number of researchers. Roodman and Schwarz [30] consider a fixed charge facility location problem in which sites must remain closed once a decision is made to close them. Such a model is appropriate in the face of a declining or shrinking market. In a subsequent note they extend the model to allow initially open facilities to be closed once and initially closed facilities to be opened once [30]. The models in both papers are solved using branch and bound. Wesolowsky and Truscott [36] consider a dynamic extension of the P-median problem [14, 15] that incorporates fixed costs associated with opening and closing facilities in each period. (Alternatively, the model may be considered a dynamic fixed charge facility location problem with a constraint on the number of facilities that must be open in each period.) The model constrains the number of location changes permitted in each period to be less than a specified value and also forces the total number of open sites in each period to equal a given value. These constraints, by restricting the growth of the state space, allow the authors to use dynamic programming to solve the problem. The authors also solve small instances of the problem using a standard branch and bound integer programming package. Dynamic extensions of the set covering location problem and the maximum covering location problem have been considered by Chrissis et al. [5], Gunawardane [13] and Schilling [31].

Van Roy and Erlenkotter [35] formulated a dynamic uncapacitated facility location model that precludes relocation of facilities. An initially open facility may be closed at most once and an initially closed facility may be opened at most once during the exogenously specified planning horizon. In its treatment of opening and closing options for facility sites, the model is therefore similar to that of Roodman and Schwarz [29] and is less general than that of Wesolowsky and Truscott [36]. Van Roy and Erlenkotter develop a dual-based solution algorithm coupled with a primal-dual adjustment procedure and (if necessary) a branch and bound algorithm. The approach, known as DYNALOC, is a modification of an earlier approach (DUALOC) developed by Erlenkotter [8] for the static uncapacitated facility location problem. Frantzeskakis and Watson-Gandy [10] allow facilities to be relocated during the planning period. They develop a solution approach that utilizes a state-space relaxation of a dynamic programming formulation of the problem to provide bounds for a branch and bound algorithm.

All of the models reviewed so far treat future conditions as being known with certainty. To handle uncertain futures, Schilling [32] formulated extensions of the set covering [34] and maximum covering [6] models that incorporate multiple future scenarios. In an extension of the set covering model, for example, the objective is to maximize the number of sites that are common to all future scenarios subject to constraints requiring all demand nodes to be covered under all scenarios and constraints that limit the number of allowable sites in each future scenario. Schilling also proposes and illustrates the use of a multiobjective maximum covering model in which trade-offs are derived between the number of sites that are common to a number of scenarios and the maximum percentage degradation of coverage (taken over all scenarios) that results from constraining the solutions to contain the specified number of sites in common. By forcing the solutions to include common sites and by constructing the common sites first, Schilling argues that a choice between alternative scenarios and the implied siting plan can be deferred. This approach is termed the *scenario planning approach*. In section 5 we explore some of the implications of the use of this approach.

All of the dynamic facility location models outlined above assume an exogenously specified planning period. The theory of forecast horizons argues that the length of the planning horizon should be determined endogenously. The optimal planning horizon is the earliest time T such that for all planning horizons of length T or more, the first period decision is the same for at least one optimal policy for each planning horizon. Bean and Smith [3] provide conditions for the existence of forecast horizons for deterministic problems. Bès and Sethi [4] and Hopp et al. [19] give existence conditions for stochastic problems. The theory of forecast horizons has been usefully applied in a number of areas, including: production scheduling problems [22, 26], capacity expansion planning [11, 17], equipment replacement [33], and R&D investment [18]. In most cases, the optimal decision in each period can often be characterized by a single variable. By way of contrast, the optimal decision in location models is characterized by a vector of binary variables, thus making the location problems inherently more difficult to solve. In the following section we outline a forecast horizon approach to the dynamic facility location problem.

3. Definition of forecast horizons in location models

We now formalize the notion of *forecast horizons* in the context of location models. Consider a dynamic facility location model and let the set of location decisions that must be made *at the beginning* of period t be represented by a vector, X_t. Given a set of open facilities in period t (which result from the location decisions made at the beginning of the first t periods along with the specification of the initially open and closed facilities) let Y_t denote the vector of decision variables related to the allocation of demands to facilities during period t. Similarly let:

$$Y = \{Y_t\} \qquad t = 1, 2, \ldots, \tau$$

and

$$\mathbf{X}_T = \{X_1, \ldots, X_T\},$$

where τ is the length of the planning period (i.e. the number of periods) in the model and $T \le \tau$. For completeness, we define $X_0 = \phi$. Note that τ may be infinite, whereas T will always be finite and will generally be either 0 or 1 as discussed below.

Without loss of generality, let $\mathbf{P}(\tau : T, \mathbf{X}_T)$ denote a *minimization* problem defined over a planning horizon composed of τ periods in which the location decisions for the first T periods are *given*. Note that in each of those periods we still need to find optimal *allocation* decisions given by the decision variables, Y_t. Thus, the optimization is over the sets of variables, $X_{T+1}, \ldots, X_\tau$ and Y. In particular, the variables $X_0, X_1, \ldots, X_T$ are inputs where X_0 specifies the *initial conditions*; i.e. the initially open and closed facilities. Clearly, the minimization is only over Y if $T = \tau$. When $T = 0$, there are no fixed decision variables though initial conditions specified by X_0 may exist. We will sometimes refer to this problem as $\mathbf{P}(\tau:0)$ or $\mathbf{P}(\tau)$.

With this notation, the *infinite horizon problem with fixed initial conditions* may be defined as problem $\mathbf{P}(\infty:0)$.

Let $V^*(\tau:T,X_T)$, $V^L(\tau:T,X_T)$, $V^U(\tau:T,X_T)$ be the optimal value, a lower bound, and an upper bound on the value of the objective function for $\mathbf{P}(\tau:T,X_T)$. For simplicity, we let $V^*(\tau)$ be the optimal value of the objective function for $\mathbf{P}(\tau:0)$ or $\mathbf{P}(\tau)$. $V^L(\tau)$ and $V^U(\tau)$ are defined similarly to be lower and upper bounds on the value of the objective function for problem $\mathbf{P}(\tau:0)$ or $\mathbf{P}(\tau)$. Note that we consider bounds on the objective function (as well as the optimal value) because finding the optimal value is often very difficult. Finding bounds may be dramatically easier.

Let X_T define a T-period location policy such that

$$V^L(\tau) \le V^*(\tau:T,X_T) \le V^U(\tau). \tag{1}$$

That is, the optimal value of the problem if we begin with the location plan defined by X_T is between some lower and upper bounds on the unconstrained problem (i.e. the problem with only the existing facilities specified). If (1) holds for all $\tau \ge \tau^*$ and

$$V^U(\tau) - V^L(\tau) \le \varepsilon$$

for all $\tau \ge \tau^*$, then X_T defines an *ε-optimal T-period policy for the ∞-period problem* and τ^* is an *ε-optimal T-period forecast horizon.* In general, we will be interested in finding an ε-optimal 1-period policy (i.e. a 1-period decision) for the ∞-period problem and an ε-optimal 1-period forecast horizon. We will refer to these quantities respectively as an *ε-optimal initial policy* (or decision) and an *ε-optimal forecast horizon.*

The search for an ε-optimal forecast horizon and an ε-optimal initial decision is justified by the fact that many of the facility location problems of practical interest are NP-complete. Thus, proving that a policy is strictly optimal for any finite duration planning period could require an inordinate amount of computation. For example, suppose $\mathbf{P}(1:0)$ is an uncapacitated fixed charge facility location problem (with no initially open facilities). This problem is NP-complete [27]. Clearly, $\mathbf{P}(\tau:0)$ will also be NP-complete as will $\mathbf{P}(\tau:T,X_T)$ as long as $T<\tau$.

The search for ε-optimal solutions is further justified by the following simple example that shows that a strictly optimal solution may not exist even for very simple problems. The example is similar to that outlined by Bean and Smith [3].

Consider a two node problem in which no facility exists prior to the beginning of the first period. Associated with each node i and each time period t is a fixed cost of establishing a facility at the site (f_{it}), a fixed cost of closing an open facility at the site (c_{it}), a fixed per period operating cost at the site assuming an open facility exists at the site (a_{it}), and costs of transporting goods from one site i to another site j which we denote by d_{ijt}. Note that the transport costs are *total* transport costs between nodes in each period and not per unit costs. We can think of these costs as being the product of a unit cost, which is independent of the quantity being shipped, and the quantity being shipped. Each node is both a demand node and a candidate site for a facility. Let:

$$a_{1t} = \begin{cases} \alpha - [\beta/(1+\gamma)], & \text{for } t = 1; \\ \alpha, & \text{for } t = 2, 3, \ldots; \end{cases}$$

$$a_{2t} = \alpha + (-1)^t \beta \qquad \forall t;$$

$$f_{it} = 0 \qquad \forall i, t;$$

$$c_{it} = \infty \qquad \forall i, t;$$

$$d_{11t} = d_{22t} = 0 \qquad \forall t;$$

$$d_{12t} = d_{21t} = 1 \qquad \forall t,$$

where r is the interest rate per period and $\gamma = (1+r)^{-1}$ is the one-period discount factor.

We assume that $\alpha - \beta > 1$. This assumption means that the (smallest) cost of operating a facility is always greater than the cost of supplying a node from a facility located at the other node. Therefore, it will be optimal to locate at only one facility in period 1. Since the cost of closing a facility is very large ($c_{it} = \infty$), the facility opened in period 1 will remain open for all time periods and the other candidate site will never be used. Since the transport costs are symmetric for all time periods, they may be excluded from the optimization. Thus, the problem becomes one of finding the site with the smallest discounted operating costs (a_{it}).

Let $C_i(t)$ be the present value of the operating costs if we locate at node i from period 1 through period t. We then have:

$$\begin{aligned} C_1(t) &= \sum_{n=1}^{t} \gamma^n \alpha - \beta\gamma/(1+\gamma) \\ &= \alpha(\gamma - \gamma^{t+1})/(1-\gamma) - \beta\gamma/(1+\gamma), \end{aligned}$$

$$\begin{aligned} C_2(t) &= \sum_{n=1}^{t} \gamma^n (\alpha + (-1)^n \beta) \\ &= \alpha(\gamma - \gamma^{t+1})/(1-\gamma) - \beta\gamma/(1+\gamma) + (-1)^t \gamma^{t+1} \beta/(1+\gamma) \\ &= C_1(t) + (-1)^t \gamma^{t+1} \beta/(1+\gamma). \end{aligned}$$

For odd values of t, $C_2(t) < C_1(t)$, while for even values of t, $C_1(t) < C_2(t)$. Therefore, *the optimal initial decision depends on the length of the planning period* in this example. Thus, for this simple problem, an optimal forecast horizon does not exist and the initial policy does not stabilize to that of selecting any particular node. Nevertheless, we note that as $t \to \infty$ the two costs converge. Thus, for sufficiently large values of t an ε-optimal initial decision will exist for all non-zero values of ε. In particular, since the absolute value of the difference between $C_1(t)$ and $C_2(t)$ is $\gamma^{t+1}\beta/(1+\gamma)$, and since this difference is a decreasing function of t,

$$t^* = \lceil \{[\ln \varepsilon + \ln(1+\gamma) - \ln \beta]/\ln \gamma - \} \rceil$$

is an ε-optimal forecast horizon for this problem where $\lceil x \rceil$ is the smallest integer greater than or equal to x. The optimal decision for a problem of this duration (i.e. select location 1 if t^* is even and location 2 if t^* is odd) is an ε-optimal initial decision. We also note that if $c_{1t} = c_{2t} = 0$ for all values of t (i.e. the *closing* costs are all 0), then the optimal policy would be to locate at node 2 in the odd numbered periods and at node 1 in the even numbered periods. This suggests that, in general, policies that prohibit opening and closing of facilities multiple times may be suboptimal. In this example, if $c_{1t} = c_{2t} = 0$ for all values of t, then the problem decomposes into t single period (trivial) location problems.

We also note that if the opening and closing costs in the example are changed to

$$f_{1t} = f_{2t} > \gamma/(1-\gamma) \qquad \forall t,$$

$$c_{it} = 0 \qquad \forall i, t,$$

then it is still optimal to locate at only one of the two sites. This is so because $\gamma/(1-\gamma)$ is the present value of all future transport costs from either site to the other. Thus, the cost of establishing the second facility exceeds the total discounted cost of shipping from one node to the other. Since the costs of establishing the facilities are identical in all other periods (and in particular in period 1), the problem again reduces to that of selecting the single site that minimizes the discounted per period operating costs (a_{it}). Finally, a_{it} may be thought of as the cost of supplying node i from a facility located at node i in period t. In other words, the problem remains the same if we set $d_{11t} = a_{1t}$ and $d_{22t} = a_{2t}$ in each period t and then assume that all operating costs are zero in each period.

This example illustrates a general problem associated with finding either exact or ε-optimal initial decisions and forecast horizons. Specifically, we need to be able to show that the initial decision is either optimal or ε-optimal for *all* planning horizons equal to or longer than the forecast horizon. Clearly, we would like to do so without having to solve an infinite number of problems. In the remainder of this section we introduce two approaches to dealing with this issue. In the first approach, we assume that all costs are bounded from above by a given value. This allows us to develop worst case bounds on the costs in any period and, as a result, an upper bound on the total cost beyond any given time. In the second approach, we ignore the costs beyond a given time and introduce the notion of empirical ε-optimal decisions and forecast horizons.

We concentrate on a cost minimization problem, where the costs considered include fixed costs of opening and closing facilities in each time period, fixed annual operating costs of any open facilities, and transport costs between demand sites and open facilities. Furthermore, we assume that each such cost (associated with each candidate site in the case of fixed costs and with each pair of sites in the case of transport costs) in each time period is bounded from above by a constant M. In the case of the transport costs, we assume that M is a bound on the *total* cost of delivering the required goods to a demand site in any period and not a bound on the *per unit* cost of delivering to the site. The total cost of delivering goods to any demand site in any period is clearly less than or equal to the product of the demand at that location and the largest unit shipment cost between a candidate facility site and that location. All costs are assumed to be non-negative. Finally, assume that all facilities are uncapacitated and that there are N demand sites.

Since we will have at most N sites open during any period (because facilities are uncapacitated), MN is an upper bound on the per period operating costs for the system in any period. Similarly, the maximum transport cost from these facilities to the demand sites is also MN. Therefore, $2MN$ is an upper bound on the total per-period operating plus transport cost for the system. (Note that by incurring a one-time closing cost for $N-1$ facilities, we can reduce the sum of these upper bounds from $2MN$ to $M(N+1)$. Furthermore, if the candidate facility locations correspond to demand sites and if the cost of transporting material from a facility at node i to demands at node i is modelled as being 0 as is often done, then the bound may be reduced to MN, the upper bound on the per period operating cost, without assuming that any facilities are closed.) Thus, while $2MN$ will be used in all of the computations below, tighter bounds may be used in cases such as those outlined above.

An upper bound on the cost of operating the system from time τ (the end of the planning period) to some future time t with *any* configuration of facilities at time τ is given by:

$$\sum_{n=\tau+1}^{t} 2MN\gamma^n = 2MN\gamma^{\tau+1}(1-\gamma^{t-\tau})/(1-\gamma).$$

Thus, for $t \geq \tau$

$$\begin{aligned} V^U(t) &\leq V^U(\tau:1,X_1) + \sum_{n=\tau+1}^{t} 2MN\gamma^n \\ &= V^U(\tau:1,X_1) + 2MN\gamma^{\tau+1}(1-\gamma^{t-\tau})/(1-\gamma). \end{aligned} \tag{2}$$

Inequality (2) simply states that one possible upper bound on the optimal solution to the t-period problem, $V^U(t)$, may be obtained by summing an upper bound on the

τ-period problem ($\tau < t$) with some fixed (possibly suboptimal) initial decision, X_1, and an upper bound on costs in periods $\tau+1$ through t.

Any optimal or ε-optimal plan for a t-period problem must begin with *some* initial decision X_1. Let this initial policy be $X_1^*(t)$. Note that we do not know this policy. In fact, $X_1^*(\infty)$, the optimal initial decision for the infinite horizon problem, is exactly what we are looking for. Nevertheless, it is then true that for $t \geq \tau$:

$$\begin{aligned} V^L(t) &\geq V^L(\tau : 1, X_1^*(t)) \\ &\geq \min_{X_1} V^L(\tau : 1, X_1) \\ &= V^L(\tau). \end{aligned} \tag{3}$$

The first inequality follows from the fact that the lower bound on the solution to a τ-period problem beginning with any initial decision, X_1, must be less than a (similarly constructed) lower bound on a t-period problem beginning with the same policy, since $t \geq \tau$ and thus additional costs will be incurred during periods $\tau+1$ through t. Thus, for any initial policy, $V^L(t : 1, X_1) \geq V^L(\tau : 1, X_1)$. Clearly this also holds for policy $X_1^*(t)$. Therefore $V^L(t : 1, X_1^*(t)) \geq V^L(\tau : 1, X_1^*(t))$. The first inequality follows by noting that $V^L(t : 1, X_1^*(t)) = V^L(t)$ (by definition). The second inequality follows from the fact that the lower bound on the τ-horizon problem beginning with some specified initial decision $X_1^*(t)$ must be at least as large as the lower bound associated with the τ-horizon problem in which we allow *any* initial policy X_1. This latter quantity is given by $V^L(\tau)$. In other words, the lower bound on a constrained τ-period problem $\mathbf{P}(\tau : T, X_T)$ should be greater than or equal to the lower bound on the unconstrained τ-period problem $\mathbf{P}(\tau : 0)$. Note that inequality (3) also follows from the fact that $V^L(t)$ increases monotonically with t.

Combining (2) and (3) we have for $t \geq \tau$,

$$V^U(t) - V^L(t) \leq V^U(\tau : 1, X_1) + 2MN\gamma^{\tau+1}(1-\gamma^{t-\tau})/(1-\gamma) - V^L(\tau). \tag{4}$$

Thus, if we can find an initial decision X_1 and a planning horizon τ such that

$$V^U(\tau : 1, X_1) + 2MN\gamma^{\tau+1}(1-\gamma^{t-\tau})/(1-\gamma) - V^L(\tau) \leq \varepsilon \qquad \forall t, \tag{5}$$

then X_1 constitutes an ε-optimal initial decision and τ is a forecast horizon for the problem. Inequality (5) may be rearranged to yield:

$$V^U(\tau : 1, X_1) - V^L(\tau) \leq \varepsilon - 2MN\gamma^{\tau+1}(1-\gamma^{t-\tau})/(1-\gamma) \qquad \forall t. \tag{6}$$

Since (6) must hold for all values of t, for $t = \infty$, inequality (6) becomes:

$$V^U(\tau : 1, X_1) - V^L(\tau) \leq \varepsilon - 2MN\gamma^{\tau+1}/(1-\gamma). \tag{7}$$

Since γ is less than 1, for any non-zero value of ε there is a sufficiently large value of τ such that the right hand side of (7) is non-negative.

Let X_1^{**} be an initial decision such that

$$V^L(\tau) \leq V^F(\tau : 1, X_1^{**}) \leq V^U(\tau), \tag{8}$$

where $V^F(\tau : 1, X_1^{**})$ is the value of the objective function for a *feasible* solution to the τ-period problem beginning with the initial decision X_1^{**}. Finding such an initial decision should not, in general, be overly difficult since the upper bound for the problem, $V^U(t)$, will usually be the objective function of the best known τ-period policy. In that case, X_1^{**} may be taken to be the initial period decision that is part of that policy. If in addition to (8), we have

$$V^U(\tau) - V^L(\tau) \leq \varepsilon - 2MN\gamma^{\tau+1}/(1-\gamma), \tag{9}$$

then the initial decision X_1^{**} is an ε-optimal initial decision and τ is an ε-optimal forecast horizon. Our ability to compute $V^U(t)$ and $V^L(t)$, upper and lower bounds on the optimal t-period problems, depends on the specific location problem under study. Inequality (9), however, implies that our ability to find ε-optimal decisions and forecast horizons is limited only by our ability to solve the τ-horizon dynamic facility location problem sufficiently accurately. Again, this may be a non-trivial problem in theory since most such problems of interest are NP-complete. However, algorithms that work very well in practice exist for many such problems (e.g. Van Roy and Erlenkotter [35] for the uncapacitated dynamic facility location problem with facilities being allowed to be opened or closed only once).

In addition to the need to be able to solve the finite horizon dynamic facility location problem sufficiently well, however, there are two additional problems associated with the use of (9). First, it requires the identification of M as an upper bound on all costs in all periods. This bound is likely to be quite large. Second, because this bound is likely to be large, the forecast horizon τ is likely to be very large. In short, we will need to be able to solve a very large finite horizon problem to be able to use inequality (9). Alternate approaches are required for practical problems.

Before discussing one alternate approach, we note that the last term of (9) corresponds to the upper bound on the sum of the discounted costs of operating the system beginning in period $\tau + 1$ for *any* configuration of facilities. The early period costs (those for periods close to $\tau + 1$) clearly contribute more to this bound than do the later period values due to the discounting of the costs. This bound may be reduced by taking *any* site configuration (e.g. the one found at the end of period t) and computing the *actual* discounted operating costs and the actual discounted transport costs for these early costs. The actual operating costs in any period can be no more than MN as noted above and will generally be considerably less than that value. Computing the transport costs in any period for a given set of demands and a given siting plan is trivial in the uncapacitated case since each demand node is assigned to the facility that can serve it at least cost. As in the case of the operating costs, we expect the actual transport costs to be significantly less than MN in each period. If in the summation representing the upper bound on the remaining costs, we replace the term $2MN$ by the actual remaining (operating and transport) costs associated with some configuration of sites for the early periods following period τ, we can reduce the value of τ for which the right hand side of (8) will be positive, thereby possibly shortening the ε-optimal forecast horizon. We are currently pursuing other approaches to bounding the costs beyond period τ.

In what follows we suggest an *empirical* approach to examining whether or not forecast horizons are likely to exist for practical problems. Consider again a dynamic facility location problem whose objective function is to be minimized. Suppose we have solved the problem for all planning horizons between 1 and τ periods long. If for all problems with planning horizons between t_1 and $t_1' = \tau$ the initial decision is the same *and* if the ratio of the difference between the upper and lower bounds to the lower bound in each such period is less than ε, then we say that t_1 constitutes an *empirical ε-optimal forecast horizon* and that the corresponding initial decision is an *empirical ε-optimal initial decision.* If there exists some sequence of periods from t_2 to t_2' (with $t_2' < \tau$) such that the conditions above are satisfied for this interval, but are not satisfied for period $t_2' + 1$, then we refer to the corresponding initial solution as a *secondary solution.* Similarly, if a secondary solution exists *and* there exists some sequence from period t_3 to t_3' (with t_3' less than the value of t_2 associated with the secondary solution) such that the conditions above are satisfied and the initial decision differs from the secondary solution, then the corresponding initial solution will be referred to as a *tertiary solution.* Alternatively, a tertiary solution may be thought of as a secondary solution defined over an interval from t_3 to t_3' (with $t_3' < t_2$) that differs from the secondary solution defined over the interval from t_2 to t_2'.

Formally, let $X_1(t)$ be the initial solution found for the t-period dynamic location problem and, as before, let $V^U(t)$ and $V^L(t)$ be upper and lower bounds on the value of the objective function for the t-period problem. Define

$$R(t) = [V^U(t) - V^L(t)]/V^L(t)$$

to be the gap between the upper and lower bounds on the solution to the t-period problem. Having solved the problem with planning horizons from 1 to τ periods long, we say we have:

(1) an *empirical ε-optimal initial solution* denoted by X_1^1 and an *empirical ε-optimal forecast horizon*, t_1, if

$$R(t) \leq \varepsilon \qquad t \geq t_l$$

and

$$X_1(t_1) = X_1(t_1 + 1) = \ldots = X_1(\tau) = X_1^1;$$

(2) a *secondary initial solution*, X_1^2, if

$$R(t) \leq \varepsilon \qquad t_2 \leq t \leq t_2' < \tau$$

and $$X_1(t_2) = X_1(t_2 + 1) = \ldots = X_1(t_2') = X_1^2$$

and either $X_1^2 \neq X_1^1$ (if a primary solution exists),

or $R(t_2' + 1) > \varepsilon$,

or both;

(3) and a *tertiary initial solution*, X_1^3, if we have a secondary solution and

$$R(t) \leq \varepsilon \qquad t_3 \leq t \leq t_3' < t_2$$

and $$X_1(t_3) = X_1(t_3 + 1) = \ldots = X_1(t_3') = X_1^3 \neq X_1^2.$$

In the following section we investigate the existence of empirical ε-optimal decisions, secondary solutions and tertiary solutions.

4. Empirical results

To investigate the existence and behavior of empirical ε-optimal forecast horizons, ε-optimal initial decisions, secondary solutions and tertiary solutions and to understand better the importance of these different types of solutions we solved a fixed charge uncapacitated dynamic facility location model with increasingly long planning horizons for a variety of input conditions. This section summarizes the results of this set of experiments.

The dynamic facility location model that we employ may be formulated as follows. We begin by defining the following indices, sets and inputs:

i = an index on facility sites;

I_o = the set of all initially closed *candidate* facility sites which can be *opened*;

I_c = the set of all initially open *candidate* facility sites which can be *closed*;

I = $I_o \cup I_c$ = the set of all *candidate* facility sites;

j = an index on demand nodes;

J = the set of all demand nodes;

t = an index on time;

T = the index of the last time period for which siting decisions are given exogenously;

τ = the number of time periods in the planning horizon;

c_{it} = the sum of all fixed costs in periods t to τ (the end of the planning horizon) if a facility is opened at site i at the start of period t for $i \in I_o$;

c_{it} = the sum of all fixed costs of operating a facility at node i from period 1 through period t plus the fixed closing cost of a facility at i at the end of period t for $i \in I_c$;

d_{ijt} = the total cost of shipping the material demanded at node j from node i to node j in period t.

We assume that all costs are positive. In addition, we define the following decision variables:

$$x_{it} = \begin{cases} 1, & \text{if the facility at site } i \text{ is } \textit{opened} \text{ at the beginning of period } t \text{ for } i \in I_o; \\ 0, & \text{otherwise.} \end{cases}$$

$$x_{it} = \begin{cases} 1, & \text{if the facility at site } i \text{ is } \textit{closed} \text{ at the end of period } t \text{ for } i \in I_c; \\ 0, & \text{otherwise,} \end{cases}$$

y_{ijt} = the *fraction* of the demand node j that is satisfied by the facility at i during period t.

X_t = $\{x_{it}\}$.

Y_t = $\{y_{ijt}\}$.

With this notation, we define the following optimization problem:

P$(\tau : T, X_T)$

$$\text{minimize} \quad \sum_{t=1}^{\tau} \sum_{i \in I} c_{it} x_{it} + \sum_{t=1}^{\tau} \sum_{i \in I} \sum_{j \in J} d_{ijt} y_{ijt} \tag{10}$$

$$\text{subject to} \quad \sum_{i \in I} y_{ijt} = 1 \qquad \forall j, t, \tag{11}$$

$$\sum_{\alpha=1}^{t} x_{i\alpha} - y_{ijt} \geq 0 \qquad \forall i \in I_o, \ \forall j, t, \tag{12}$$

$$\sum_{\beta=t}^{\tau} x_{i\beta} - y_{ijt} \geq 0 \qquad \forall i \in I_c, \ \forall j, t, \tag{13}$$

$$x_{it} \in \{0,1\} \qquad \forall i, t, \tag{14}$$

$$y_{ijt} \geq 0 \qquad \forall i, j, t. \tag{15}$$

Note that $X_1, \ldots, X_T$ are inputs and that the minimization is over $X_{T+1} \ldots X_\tau$ and Y. Nevertheless, we include the contribution of $X_1, \ldots, X_T$ in the value of the objective function.

The objective function (10) minimizes the sum of all fixed costs of opening new facilities and closing existing facilities as well as the cost of transporting material between facilities and demand nodes. Constraint (11) stipulates that the entire demand of each node j must be completely satisfied in each period t. Constraints (12) and (13) ensure that demands at j are not assigned to a facility at node i in period t unless a facility is open at node i in period t. Specifically, constraint (12) applies to candidate facility sites for opening and states that demands at node j in priod t can be assigned to a facility at node i if and only if a facility at node i was opened in one of the periods up to and including period t. Similarly, constraint (13) applies to those sites that are initially open and states that demands at node j in period t can be assigned to an already open facility at node i if and only if the facility at node i is closed at the end of period t or a subsequent period. Note that the model assumes that all initially open facilities will eventually be closed. Any initially open facility that remains open throughout all τ periods will be closed at the end of period τ thereby allowing demands to be assigned to the facility throughout the entire planning horizon. Constraints (14) are integrality constraints on the location variables x_{it} and constraints (15) are non-negativity constraints on the flow variables y_{ijt}.

Note that while this formulation does not explicitly prohibit facilities from being opened (or closed) multiple times, (i.e. there is no constraint stipulating $\sum_{t=1}^{\tau} x_{it} \leq 1$), in fact no facility will be opened (or closed) more than once. This results from the non-negativity of the cost coefficients c_{it} and the fact that constraints (12) and (13) imply that $\sum_{t=1}^{\tau} x_{it} \leq 1$ for all nodes i. Thus, problem $\mathrm{P}(\tau : T, X_T)$ is an uncapacitated fixed charge dynamic facility location model in which candidate facilities may be opened or closed at most once. Thus, the model can be solved using the dual ascent algorithm (DYNALOC) developed by Van Roy and Erlenkotter [35]. This algorithm was coded and used in the series of runs outlined below (see also [24]).

All model runs used a modified form of the 15 node network presented in Current et al. [7]. The modification involved changing the intra-nodal distances from 0 to 5 units. With this distance matrix, three basic cases (labelled A, B, and C in table 1) were defined. For each case, fixed opening, closing, and annual operating costs were selected for each candidate site and each time period from uniform distributions whose limits are shown in table 1. For each scenario, values of a_0 and b_0 were chosen uniformly on the indicated ranges to be used in generating demands for each node in each peiod. This was done by randomly (and independently) selecting a demand function from the six functional forms given in table 2 and selecting the c and r parameters from the ranges indicated in the table.

Each case was analyzed using interest rates of 3%, 6%, and 9%. We refer to a combination of a case and an interest rate as a *scenario*. Each scenario was solved for odd planning horizons of between 1 and 25 periods (i.e. 1, 3, 5, . . . , 25 periods). This resulted in the solution of a total of 117 dynamic location problems in which the initial decisions were unconstrained. In the results reported below, the value of

Table 1

Summary of input values*.

Input factor	Case A	Case B	Case C
Existing facilities?	No	No	Yes at 1, 2, 3, 7
Per unit transport cost	0.10	0.10	0.10
Max. periodic fixed cost	15,000	15,000	1,000
Min. periodic fixed cost	10,000	10,000	100
Max. opening cost	250,000	150,000	300,000
Min. opening cost	175,000	100,000	200,000
Max. closing cost	Not applicable	Not applicable	15,000
Min. opening cost	Not applicable	Not applicable	10,000
Max. a_0, b_0	500	500	750
Min. a_0, b_0	350	350	350

*All three cases were run at three interest rates: 3%, 6%, and 9%.

Table 2

Demand data.

Case	Demand function*	c Parameter range
1	$d_t = r_t(a_0 + b_0 \sin(t/c))$	$0 \le c \le 10$
2	$d_t = r_t(a_0 + b_0 e^{-ct})$	$0 \le c \le 1$
3	$d_t = r_t(a_0 + b_0 e^{ct})$	$0 \le c \le 1$
4	$d_t = r_t(a_0 + b_0 \ln(-ct))$	$0 \le c \le 10$
5	$d_t = r_t(a_0 - b_0 e^{ct})$	$0 \le c \le 1$
6	$d_t = r_t(d_{t-1} + cd_{t-1})$	$-0.02 \le c \le 0.02$

*r_t is chosen uniformly on [0,0.1] for all cases.

ε used in defining the existence of an empirical ε-optimal forecast horizon, and empirical ε-optimal initial decision, secondary solutions and tertiary solutions was 0.02.

Figure 1 summarizes the results for the case A and a 3% interest rate scenario. The top part of the figure shows the initial solution that was obtained as a function of the length of the planning horizon, while the bottom part of the figure, plots the percentage gap between the lower and upper bounds ($R(t)$ as a percentage) as a function of the length of the planning horizon. From period 15 on, the gap was under 2% and the initial solution remained the same. Thus, we identify period 15 as the *empirical 0.02-optimal forecast horizon* and locations 3, 10, and 13 as constituting the *empirical 0.02-optimal initial solution.* During periods 7 and 9, the

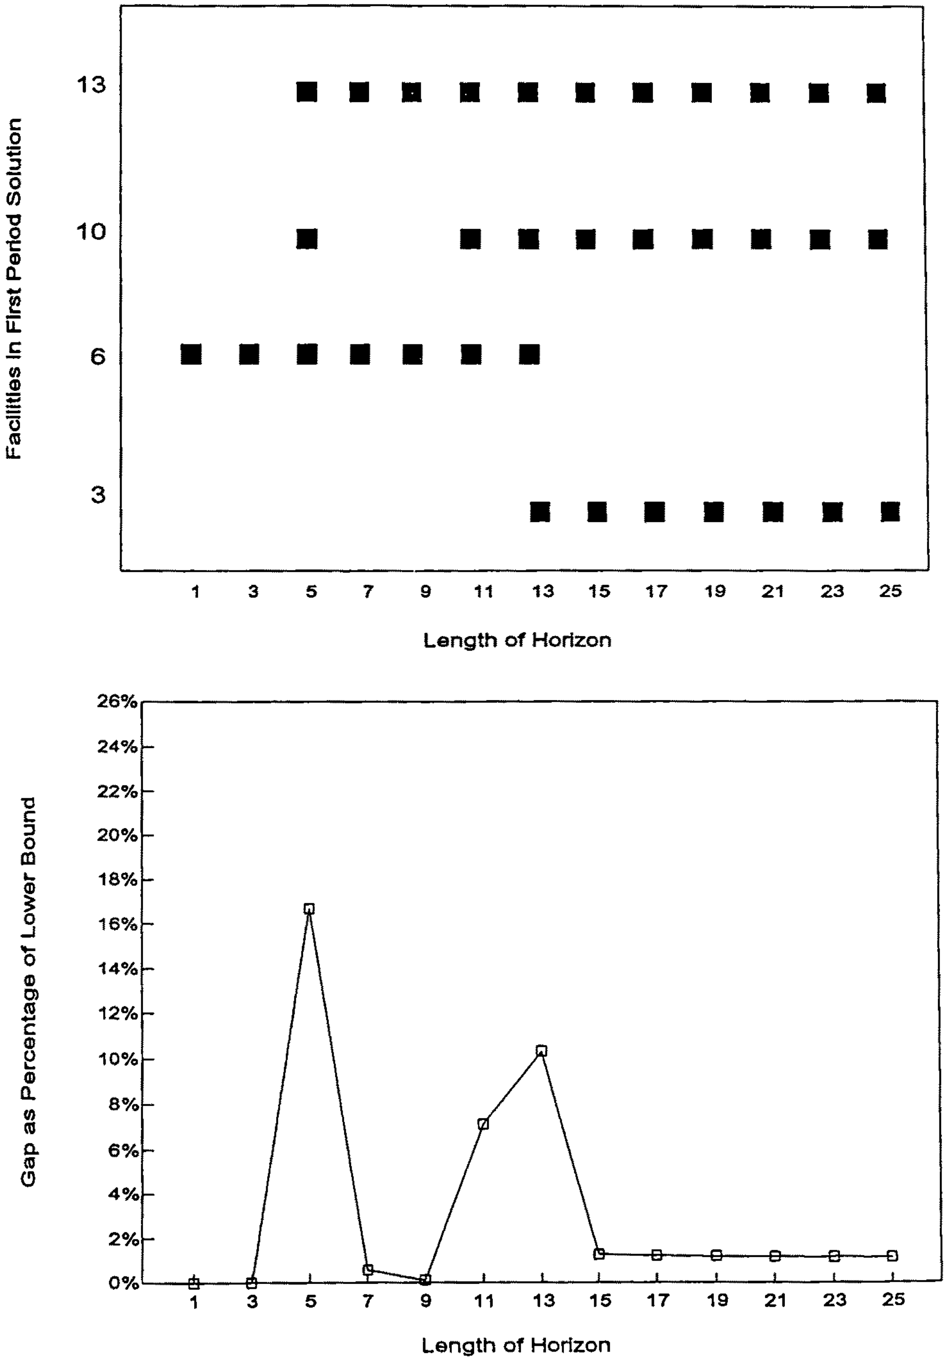

Fig. 1. Finite horizon solutions: case A, 3% interest rate.

gap is again less than 2% and the initial solution was the same for both runs. Thus, locations 6 and 13 constitute a *secondary solution* for this scenario. Finally, during periods 1 and 3, the gap is under 2% and the initial solution is identical for both periods. Thus, location 6 is identified as a *tertiary solution* for this scenario.

Tables 3 and 4 summarize the results for all 9 scenarios. As indicated, 0.02-optimal forecast horizons and initial solutions were identified in 6 of the 9 scenarios.

Table 3

Summary of periods during which solutions of different types exist.

Case	Interest rate	Forecast horizon period	Secondary Soln. periods	Tertiary Soln. periods
A	3	15	7–9	1–3
	6	21	7–9	1–3
	9	G.S.	7–11	L.C.
B	3	17	7–13	G.S.
	6	G.S.	9–15	G.S.
	9	G.S.	9–17	G.S.
C	3	19	15–17	1–7
	6	19	1–9	P.S.
	9	21	1–11	P.S.

G.S: gap size requirement exceeded; L.C: locations change though gap size requirement met; P.S: previous solution precludes existence of this solution type.

Table 4

Summary of locations of different solution types*.

Case	Interest rate	Forecast horizon initial locations	Secondary soln. initial locations	Tertiary soln. initial locations
A	3	3, 10, 13	6, 13	6
	6	3, 10, 13	6, 13	6
	9	N.A.	6, 13	N.A.
B	3	1, 6, 13	2, 6, 13	N.A.
	6	N.A.	2, 6, 13	N.A.
	9	N.A.	1, 6, 13	N.A.
C☆	3	5, 8, 10, 12, 14	5, 8, 12, 14	13
	6	5, 8, 12, 14	13	N.A.
	9	8, 12, 14	13	N.A.

*N.A. = Not Applicable.

☆Initial locations for case C are additional sites that are added in the first period to the existing sites (1, 2, 3, 7). No sites were closed in the first period under any of the scenarios.

Secondary solutions were identified in all 9 scenarios; and tertiary solutions were found for 3 of the 9 scenarios. In two of the scenarios (cases C with interest rates of 6 and 9 percent), tertiary solutions were not identified because the secondary solution begins in period 1, thereby precluding the existence of a tertiary solution. In the other cases in which either a 0.02-optimal initial solution or a tertiary solution was not found, either the ratio of the difference between the bounds to the lower bound exceeded 0.02 (indicated by G.S.) or the gap size criterion was satisfied but the first period solutions changed each period (indicated by L.C.).

The empirical 0.02-optimal initial solution is an approximation to the infinite horizon optimal initial solution. The use of either solution should result in low (or optimal) total discounted costs if a sufficiently long planning period is used. However, if a shorter planning period is employed, the use of either solution may result in higher costs than would be obtained if a different initial solution were employed. To examine this phenomenon, we resolved all of the problems on which an empirical 0.02-optimal initial solution was identified, *forcing* this solution into the model during period 1 (and allowing the model to optimize the solution from then on). The ratio of the cost of this solution (identified as the FHS solution) to the upper bound in the unconstrained case (UB) gives an indication of the short term cost of using the initial solution that is good for the long term instead of that which might be better in the short term. Figure 2 plots the results of this analysis for case A with an interest rate of 3%. As expected, the ratio is (close to) 1.0 beginning with the

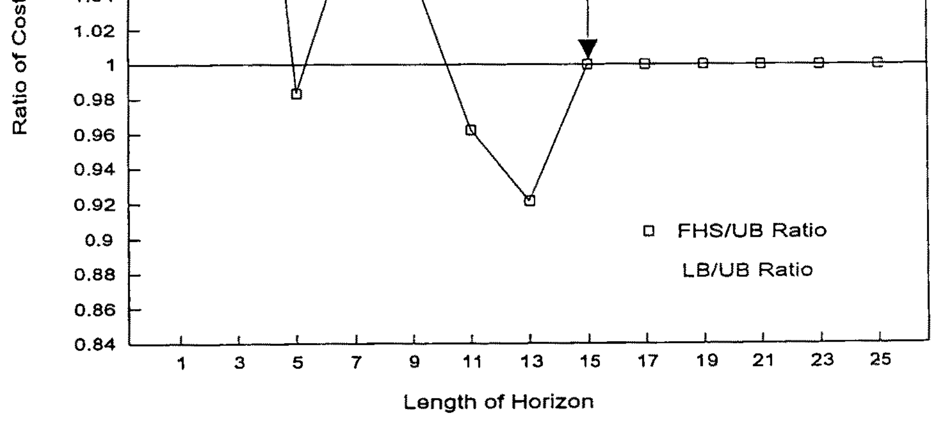

Fig. 2. Comparison of FHS to bounds: case A, 3% interest rate.

empirical forecast horizon. It is worth noting, however, that from period 5 on, the penalty for using this initial solution is less than 10 percent. In fact, in periods 5, 11 and 13, constraining the initial solution to be that of the empirical 0.02-optimal initial solution results in a *better* solution than is obtained in the absence of this constraint. Note that it is during these periods that the upper bound is considerably greater than the lower bound for the unconstrained problem (see fig. 1). The reason for being able to do better in the constrained case than we could in the unconstrained case is that Van Roy and Erlenkotter's algorithm results in heuristic solutions unless it is embedded in a branch and bound procedure. We did not employ a branch and bound procedure. (The heuristic nature of the algorithm also means that the FHS/UB ratio need not equal 1.0 for all periods after the empirical 0.02-optimal forecast horizon.)

The empirical 0.02-optimal initial solution (FHS), secondary (SS) solution, and tertiary solution (TS) may all be considered viable candidate initial solutions. Given that both an empirical 0.02-optimal initial solution and a secondary solution existed for a scenario, we reran the model using the two solutions in the first period and allowing the model to optimize the locations beyond that time. The ratio of the objective functions for these two runs for any particular planning horizon indicates the relative penalty associated with using either siting plan as the initial solution. For short planning horizons, we expect the secondary solution to do better while the FHS solution should do better for longer planning horizons. Figure 3 plots the results of this analysis for case A with an interest rate of 3%. The FHS solution was

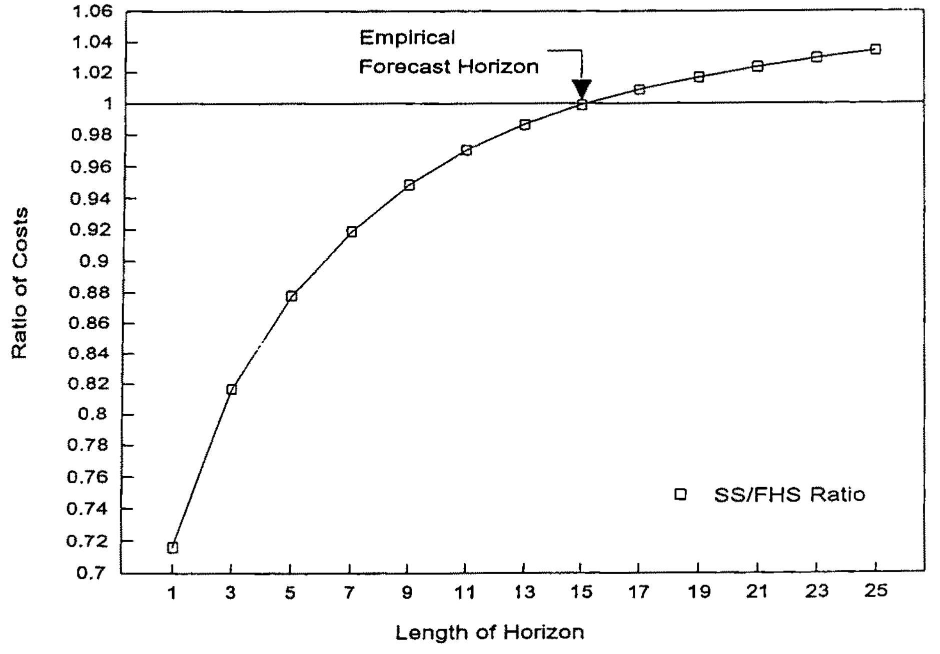

Fig. 3. Comparison of FHS and SS: case A, 3% interest rate.

better than the SS solution as an initial decision for all planning horizons equal to or longer than the empirical forecast horizon. In other scenarios, the FHS solution was better for even shorter planning horizons and in some cases was better than the SS solution during part of the period associated with the secondary solution. Again, this results from the heuristic nature of the algorithm and is an indication of the fact that the FHS solution would have been a better initial decision for planning horizons of this duration, but that the algorithm was incapable of finding this solution. Nevertheless, in all six scenarios in which both an empirical 0.02-optimal initial decision and a secondary solution were identified, the penalty associated with using the FHS solution with short planning horizons was significantly larger than the penalty associated with using the SS solution with long planning horizons. If significant uncertainty surrounds the inputs associated with future periods (during which the FHS becomes the preferred initial decision), it may be better to adopt the secondary solution as the initial solution.

In three scenarios, an empirical 0.02-optimal initial solution was not identified. In all three scenarios, however, a secondary solution was found. To examine how well the secondary solution performed as an initial decision, we forced this solution into the model for the initial period and reran the model. Figure 4 (for case B with an interest rate of 6%) is illustrative of the results. The secondary solution was identified for periods 9–15 (see table 3). However, fig. 4 shows that if we force the secondary solution into the model during the initial period, the objective function is within 2% of the lower bound from *period 7 on*. This suggests that the secondary

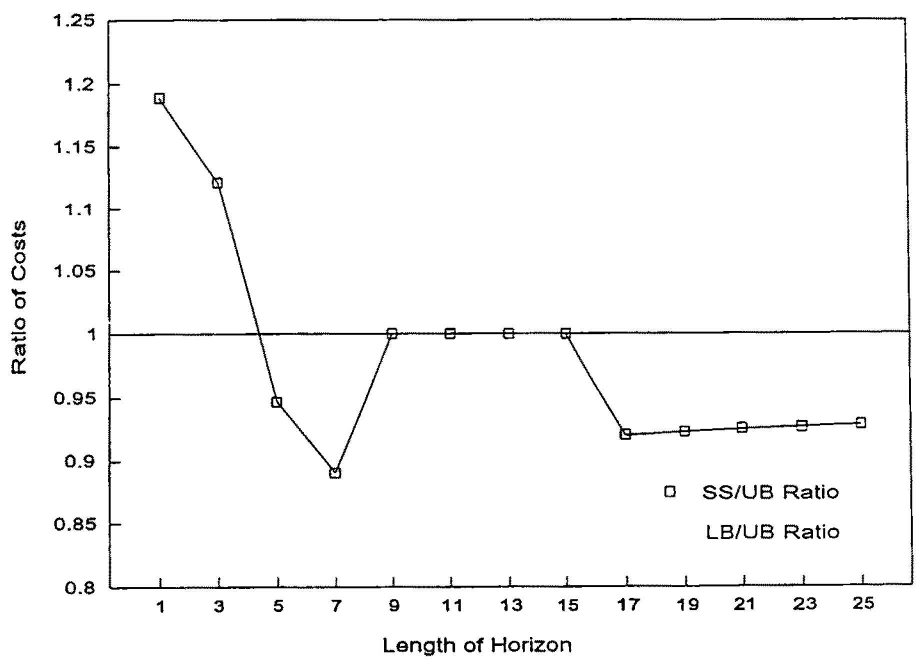

Fig. 4. Comparison of SS to bounds: case B, 6% interest rate.

solution is actually an empirical 0.02-optimal initial decision and that the associated empirical 0.02-optimal forecast horizon is period 7. Similar results were obtained for the other 2 scenarios in which the unconstrained model runs failed to identify an empirical 0.02-optimal solution.

Finally, we examined the results of forcing the tertiary solution into the model as the initial decision. Figure 5 (for case A with an interest rate of 3%) is typical of the results obtained. If the planning horizon is greater than 5 periods, either the secondary solution or the empirical 0.02-optimal initial solution (FHS) is better than the tertiary solution as an initial decision. The fact that the tertiary solution is better than the FHS solution for periods 5, 11, and 13 is again the result of the heuristic nature of the solution procedure and is not surprising in light of the rather large gaps between the upper and lower bounds for these planning horizons, as shown in fig.1.

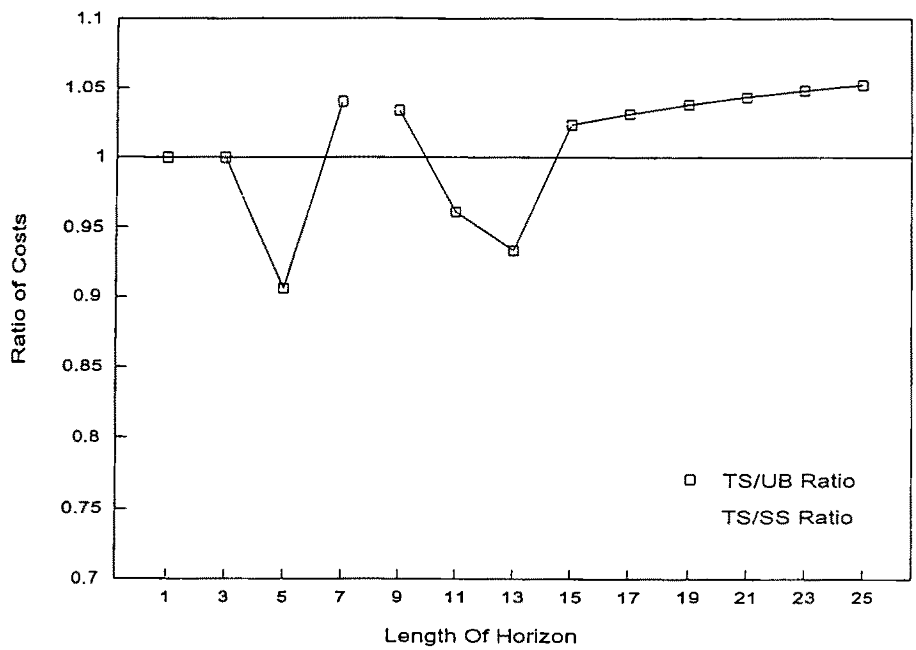

Fig. 5. Performance of TS: case A, 3% interest rate.

In summary, the empirical findings outlined above and summarized in greater detail in Medina [24] indicate that good initial decisions can be found and empirical ε-optimal forecast horizons identified despite the fact that inputs can be constructed for which a strictly optimal forecast horizon and initial decision will never exist.

5. A note on scenario planning

Sections 3 and 4 deal with the problem of finding a forecast horizon and an optimal (or near optimal) initial solution in the case of future demands that are known with *certainty*. In this section, we highlight some of the problems associated with an alternate approach that has been suggested for the problem with *uncertain* future demands.

As noted in section 2, Schilling [32] formulated an approach, using objectives, in which the number of facilities that are common to all future scenarios is constrained to be at least a given value. He was then able to develop tradeoff diagrams indicating the maximum degree of degredation in the objective function under any of the scenarios as a function of the number of required common facilities. He argued that if the degree of degredation was small, planners might want to build the common facilities first. In so doing, planners could defer having to decide which future scenario is most likely until additional information is obtained and could defer deciding which of the facilities that are unique to a particular scenario should be built. This approach is often called the *scenario planning approach.*

In this section, we present an example that indicates that the adoption of the scenario planning approach may not only lead to suboptimal decisions but can actually lead to the adoption of the *worst possible alternative.* The network used is shown in fig. 6 along with the per unit inter-nodal transport costs for all periods.

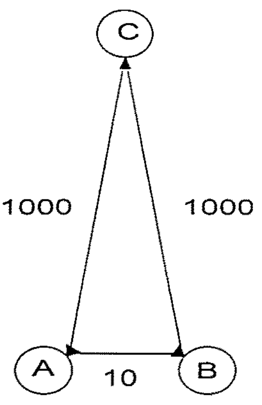

Fig. 6. Example transportation network.

All intra-nodal unit transport costs are taken to be 1. Table 5 presents the demands in six periods at each node under each of two scenarios. Note that the two scenarios are identical except that the demands for nodes A and B are inverted. A fixed cost of 50,000 is charged for the construction of a facility at any node in any period. Facilities are assumed to be uncapacitated.

Under scenario 1, the optimal solution is to build facilities as nodes A and C in period 1. The total discounted cost of this strategy is 158,263. By symmetry, the optimal solution under scenario 2 is to build facilities at nodes B and C with a total discounted cost of 158,263. Node C is common to both solutions. Thus, the scenario planning approach would call for the construction of a facility at node C

Table 5

Example problem demands*–Scenario 1☆

Time Period	Node A	Node B	Node C
1	1000	750	750
2	1050	788	788
3	1103	827	827
4	1158	868	868
5	1216	911	911
6	1277	957	957

* Demands grow approximately 5% per period.

☆ Scenario 2 demands are identical to those listed above except that the demands for nodes A and B are *inverted.*

in period 1. If such a policy is adopted and scenario 1 evolves, a facility will be added at node A in period 2. The total cost of this strategy, however, is 1,898,306– *almost 12 times the optimal cost.* Table 6 summarizes the relative cost of other strategies under each of the two scenarios. In particular, if we adopt the optimal scenario 1 solution (locate at A and C in period 1) and scenario 2 evolves, the cost is only 9% greater than it would have been had we been able to select the optimal scenario 2 strategy. Also, if we build only one facility in period 1, building at either of the non-common sites (A or B) in period 1 is preferable to building at the

Table 6

Relative costs of alternative solution strategies*.

Facilities built in period 1	Facilities built in period 2	Relative cost under scenario 1☆	Relative cost under scenario 2☆
A,C	–	1.000	1.090
B,C	–	1.090	1.000
A	C	5.725	5.815
B	C	5.815	5.725
C	A	11.995	12.070
C	B	12.070	11.995

* Values shown are the ratios of the objective function using the indicated construction strategy to the optimal objective function for the particular scenario. Thus, the ratio of the cost of building at location A in period 1 and location C in period 2 to the optimal cost under scenario 1 is 5.725:1.

☆ The optimal cost strategy under scenario 1 is to build at locations A and C in period 1 at a total discounted cost of 158,263. The optimal cost strategy inder scenario 2 is to build at locations B and C in period 1 at a total discounted cost of 158,263.

common site, even if the site selected for construction in period 1 is not a part of the optimal solution for the scenario that ultimately evolves. Since node C is isolated and is a small demand node under either of the two scenarios, building at that node in the first period and delaying the decision as to where the second facility should be located results in exorbitant first period transport costs and is *the worst possible initial decision.*

While the example outlined above applies to the problem of minimizing the present value of future costs and the scenario planning approach was proposed for covering models, it is clear that similar problems may arise in the application of the approach to covering models. In short, planners need to be concerned not only with the final configuration of sites but also with the time staging of facility additions (and deletions) that will enable the system to reach the final configuration in an optimal manner.

6. Summary, conclusions and future work

In this paper we have proposed an approach to dynamic facility location modelling that does not view the planning horizon as an exogenous input. Rather, we suggest finding a forecast horizon such that the initial optimal (or near optimal decision is invariant with respect to conditions beyond the forecast horizon. We showed that simple examples can be constructed in which a strictly optimal (finite) forecast horizon and initial decision do not exist. This together with the NP-completeness of most of the location models of practical interest justified our search for ε-optimal forecast horizons and initial solutions. For this relaxed problem we proposed a crude method of identifying ε-optimal solutions which involved bounding the costs in all future periods. Next we defined an empirical ε-optimal solution as one that is ε-optimal if costs beyond the forecast horizon are ignored instead of approximated or bounded. Empirical tests showed that good initial decisions and empirical ε-optimal forecast horizons could be found for small problems. Finally, we showed that the scenario planning approach for dealing with uncertain future conditions can lead to the adoption of the worst possible initial decision. Thus its use must be accompanied by extreme care and caution.

Much additional work remains. First, the empirical study reported on above is limited. More extensive empirical testing with bigger networks and with other location models should be conducted to determine whether or not the notions of empirical ε-optimal forecast horizons and initial solutions are of practical value. Second, methodological work aimed at finding ε-optimal forecast horizons and initial solutions should be conducted. The authors are currently engaged in work in this latter arena.

Acknowledgements

The first and third authors were supported in part by a National Science Foundation Presidential Young Investigator Award, grant 8352298 and a matching grant from the General Motors Corporation. The first author was also supported by a Fulbright Research Grant. The work of the second author was supported in part by National Science Foundation Grants ECS-8619732 and DDM-8905638. This support is gratefully acknowledged.

References

[1] N. Ahituv and O. Berman, *Operations Management of Distributed Service Networks: A Practical Quantitative Approach* (Plenum Press, New York, 1988).
[2] R.H. Ballou, Dynamic warehouse location analysis, J. Marketing Res. 5(1968)271–276.
[3] J.C. Bean and R.L. Smith, Conditions for the existence of planning horizons, Math. Oper. Res. 9(1984)391–401.
[4] C. Bès and S. Sethi, Concepts of forecast and decision horizons: Applications to dynamic stochastic optimization problems, Math. Oper. Res. 13(1988)295–310.
[5] J.W. Chrissis, R.P. Davis and D.M. Miller, The dynamic set covering problem, Appl. Math. Progr. 6(1982)2–6.
[6] R. Church and C. ReVelle, The maximal covering location problem, Papers Regional Sci. Assoc. 32(1974)101–118.
[7] J. Current, C. ReVelle and J. Cohon, The shortest covering path problem: An application of locational constraints, J. Regional Sci. 24(1984)161-183.
[8] D. Erlenkotter, A dual based procedure for uncapacitated facility location, Oper. Res. 26(1978) 992–1009.
[9] D. Erlenkotter, A comparative study of approaches to dynamic location problems, Euro J. Oper. Res. 6(1981)133–143.
[10] M. Frantzeskakis and C.D.T. Watson-Gandy, The use of state space relaxation for the dynamic facility location problem, Ann. Oper. Res. 18(1989)189–212.
[11] J. Friedenfelds, *Capacity Expansion* (North-Holland, New York, 1981).
[12] A. Ghosh and G. Rushton (eds.), *Spatial Analysis and Location-Allocation Models* (Van Nostrand Reinhold, New York, 1987).
[13] G. Gunawardane, Dynamic versions of set covering type public facility location problems, Euro. J. Oper. Res. 10(1982)190–195.
[14] S.L. Hakimi, Optimum locations of switching centers and the absolute centers and medians of a graph, Oper. Res. 12(1964)450–459.
[15] S.L. Hakimi, Optimum distribution of switching centers in a communication network and some related graph theoretic problems, Oper. Res. 13(1965)462–475.
[16] G.Y. Handler and P.B. Mirchandani, *Location on Networks: Theory and Algorithms* (M.I.T. Press, Cambridge, MA, 1979).
[17] J. Higle, J.C. Bean and R.L. Smith, Capacity expansion under stochastic demands, Technical Report 84-28, Dept. Industrial and Operations Engineering, University of Michigan (1984).
[18] W.J. Hopp, A sequential model of R&D investment over an unbounded time horizon, Manag Sci. 33(1987)500–508.
[19] W.J. Hopp, J.C. Bean and R.L. Smith, A new optimality criterion for non-homogeneous Markov decision processes, Oper. Res. 35(1987)875–883.
[20] R.E. Hughes and W.B. Powell, Mitigating end effects in the dynamic vehicle allocation model, Manag. Sci. 34(1988)859–879.
[21] A.P. Hurter and J.S. Martinich, *Facility Location and Theory of Production* (Kluwer Academic, Boston, MA, 1989).
[22] P. Kleindorfer and H. Kunreuther, Stochastic horizons for the aggregate planning problem, Manag. Sci. 24(1978)485–497.

[23] R.F. Love, J.G. Morris and G.O. Wesolowsky, *Facilities Location: Models and Methods* (North-Holland, New York, 1988).
[24] B. Medina, Dynamic facility location modeling, M.S. Thesis, Department of Civil Engineering, Northwestern University, Evanston, IL (1989).
[25] P.B. Mirchandani and R.L. Francis (eds.), *Discrete Location Theory* (Wiley, New York, 1991).
[26] T.E. Morton, Universal planning horizons for generalized convex production scheduling, Oper. Res. 26(1978)1046–1058.
[27] G.L. Nemhauser and L.A. Wolsey, *Integer and Combinatorial Optimization* (Wiley, New York, 1988).
[28] W.B. Powell, A comparative review of alternative algorithms for the dynamic vehicle allocation problem, in: *Vehicle Routing: Methods and Studies*, ed. B.L. Golden and A.A. Assad (North-Holland, New York, 1988).
[29] G.M. Roodman and L.B. Schwarz, Optimal and heuristic facility phase-out strategies, AIIE Trans. 7(1975)177–184.
[30] G.M. Roodman and L.B. Schwarz, Extensions of the multi-period facility phase-out model: New procedures and application to a phase-in/phase-out problem, AIIE Trans. 9(1977)103–107.
[31] D.A. Schilling, Dynamic location modelling for public-sector facilities: A multicriteria approach, Dec. Sci. 11(1980)714–724.
[32] D.A. Schilling, Strategic facility planning: The analysis of options, Dec. Sci. 13(1982)1–14.
[33] S. Sethi and S. Chand, Planning horizon procedures for machine replacement models, Manag. Sci. 25(1979)140–151.
[34] C. Toregas, R. Swain, C. ReVelle and L. Bergmann, The location of emergency service facilities, Oper. Res. 22(1971)1363–1373.
[35] T. Van Roy and D. Erlenkotter, A dual based procedure for dynamic facility location, Manag. Sci. 28(1982)1091–1105.
[36] G.O. Wesolowsky and W.G. Truscott, The multiperiod location–allocation problem with relocation of facilities, Manag. Sci. 22(1975)57–65.