Quest Journals
Journal of Research in Applied Mathematics
Volume 5 ~ Issue 3 (2019) pp: 31-37
ISSN (Online): 2394-0743 ISSN (Print): 2394-0735
www.questjournals.org

Research Paper

Modelling of Dynamic Location Problem in the Context of Distributed Service Networks

Dr. Shailendra Kumar
Assistant Professor in Mathematics,
Govt. Raza P.G. College, Rampur

Abstract
This paper investigates facility location because existing static models which predict operational behavior of networks show limitations because actual networks undergo continuous changes. The Dynamic Location Problem (DLP) addresses this reality by extending classical facility location theory across multi-period planning horizons, allowing facilities to open, relocate, or close as demand landscapes change. The application of this framework to distributed service networks which require geographically separated nodes to work together for total service delivery, results in a more complex and difficult problem. The article establishes DLP mathematical formulations, which include deterministic and stochastic and robust formulations, while showing their applicability to distributed network structures. The research investigates essential modelling decisions which include time representation, uncertain demand management, relocation cost integration, and inter-node dependency modelling. The research presents solution approaches which include Lagrangian relaxation and Bender's decomposition and metaheuristics while explaining their respective computational advantages and disadvantages. These models show practical value through their applications in emergency medical services and cloud computing infrastructure and logistics networks. The article argues that modern service network design requires temporal adaptability because it represents core structure needed for contemporary service networks.

Keywords: *dynamic location problem, p-median model, distributed service networks, stochastic optimization, facility location, network design*

I. Introduction

Imagine a city that grows rapidly on its eastern edge over five years. A hospital built in the city center a decade ago now sits too far from where most people actually live. A delivery hub positioned near an industrial zone becomes less useful once that zone turns into a residential area. These are not abstract problems — they happen constantly in modern urban and regional planning, and they expose a basic flaw in how facility location has traditionally been modeled.

This is the fundamental tension that the Dynamic Location Problem tries to resolve. The classic facility location problem asks where to place facilities to best serve a set of demand points. That question, posed once, for a fixed snapshot of the world, yields elegant and well-studied answers. The p-median model, the set covering model, the maximum coverage location problem — these are mature tools, widely understood and computationally tractable. But they all share one critical assumption: demand is static and facilities, once placed, stay put. That assumption is rarely true for long.

Distributed service networks make things more interesting still. In these systems — content delivery networks, multi-depot logistics operations, regional hospital systems, emergency dispatch networks — no single facility works alone. Nodes share load, route requests between one another, and compensate for each other's failures. The position of any one node shapes the performance of every other. You cannot optimize node placements independently; you have to think about the whole system at once.

This article brings those two ideas together. The goal is to explain, concretely and rigorously, how the Dynamic Location Problem is modelled when the service network is distributed — when facilities are not isolated points serving independent demand, but interconnected nodes whose relative positions determine collective system performance. We walk through the mathematics, the modelling choices that actually matter in practice, the solution approaches researchers and practitioners use, and the real-world domains where these models are most consequential.

II. Background: Facility Location and the Case for Dynamic Models

2.1 The Classical Framework

Hakimi's (1964) formulation of the p-median problem is still the starting point for most discussions of facility location. Given a set of demand nodes I and candidate facility sites J, the p-median model places exactly p facilities to minimize total weighted service distance. Letting $y_{j} \in \{0,1\}$ indicate whether a facility is opened at site j, $x_{ij} >= 0$ the fraction of demand at i served by facility j, $d_{ij}$the service distance, and h_i the demand weight, the model is:

$$min \sum_{i \in I} \sum_{j \in J} d_{ij}.h_i.x_{ij}$$

subject to:

$$\sum_{j \epsilon J} x_{ij} = 1 \ \forall_i \in I, \quad x_{ij} \le y_j \forall_{i,j}, \quad \sum_{j \epsilon J} y_j = p, y_j \epsilon \{0,1\}$$

The model is clean, the constraints are intuitive, and for moderate-size instances, branch-and-bound handles it well. Church and ReVelle's (1974) maximal coverage location problem extends this logic: given a fixed number of facilities and a coverage radius r, maximize the total demand within range of at least one facility. Between them, these two formulations gave the field a common language and still underpin the vast majority of applied work.

2.2 Why Static Is Not Enough

The problem is not the mathematics. The problem is the assumption that demand is fixed. In practice, populations age and migrate. Infrastructure investments reshape travel times. Economic activity concentrates and disperses. A distribution center perfectly located for 2010 demand patterns may be awkwardly placed by 2020. Worse, the longer the planning horizon, the more severe the mismatch — and most infrastructure investments operate on 10-to-20-year planning cycles.

Wesolowsky (1973) was among the first to formally study multi-period location, showing that even with perfectly known demand trajectories, the optimal sequence of facility configurations is not simply a series of static solutions strung together. Relocation costs introduce genuine intertemporal trade-offs: sometimes it is worth accepting a suboptimal configuration now to avoid expensive moves later. That insight launched an entire branch of the literature that this article explores.

III. Mathematical Formulation of the Dynamic Location Problem

3.1 Decision Variables and Objective

Let $T = \{1, 2, \dots, \tau\}$ be the planning horizon. Demand at node i in period t is h_i^t. The core decision variables are:

- $y_j^t \in \{0,1\}$: facility j is open in period t
- $x_{\{ij\}}^t \ge 0$: fraction of demand at i served by facility j in period t
- $o_j^t \in \{0,1\}$: facility j is opened at the start of period t
- $c_j^t \in \{0,1\}$: facility j is closed at the start of period t

The dynamic location problem minimizes the sum of operating costs, opening costs, closing costs, and service assignment costs across all periods:

$$min \sum_{t \in T} \left[\sum_{j \in J} \left(f_j^t y_j^t + O_j^t o_j^t + C_j^t c_j^t \right) + \sum_{i \in I} \sum\nolimits_{j \epsilon J} d_{ij} h_i^t \, x_{ij}^t \right]$$

Here f_j^t is the fixed operating cost of facility j in period t, while $O_j^t$and C_j^t capture the costs of opening and closing, respectively. The two terms in the outer sum pull in opposite directions: to minimize service costs, you want to track demand closely and reconfigure often; to minimize relocation costs, you want stability.

3.2 Structural Constraints

Demand assignment and facility availability constraints carry over from the static model:

$$\sum_{j \in J} x_{ij}^t = 1 \quad \forall i \in I, \quad t \in T$$

$$x_{ij}^t \le y_j^t \quad \forall i \in I, \quad j \in J, \quad t \in T$$

The constraint that makes the dynamic model genuinely dynamic is the linking condition between consecutive periods:

$$y_j^t - y_j^{\{t-1\}} = o_j^t - c_j^t \quad \forall j \in J, \qquad t \in T \setminus \{1\}$$

This equation says that a facility's status in period t depends on its status in period t-1 plus any opening or closing decisions made at the start of t. It ties decisions across time, making the problem inherently sequential — you cannot solve period t without knowing what happened in period t-1. That dependency is both the source of the model's realism and much of its computational difficulty.

For distributed networks with finite node capacities, an additional constraint bounds the total demand each open facility can serve:

$$\sum_{i \in I} h_i^t x_{ij}^t \leq Q_j y_j^t \quad \forall j \in J, \qquad t \in T$$

where Q_j is the service capacity of facility j. These forces demand overflow to secondary facilities when the nearest node is saturated — exactly the kind of load-balancing behavior that distributed networks are designed to exploit.

3.3 Introducing Demand Uncertainty

Real distributed networks rarely evolve according to a deterministic script. Demand fluctuates, link characteristics change, and new demand zones emerge. The stochastic variant of the DLP replaces known demand values with random variables $\tilde{h}_i^t$ and minimizes expected cost:

$$min \sum_{t \in T} \left[\sum_{j \in J} \left(f_j^t y_j^t + O_j^t o_j^t + C_j^t c_j^t \right) + \mathrm{E}_{\tilde{h}} \sum_{i \in I} \sum\nolimits_{j \in J} d_{ij} \tilde{h}_i^t x_{ij}^t \right]$$

The most common approach in practice is scenario-based stochastic programming. Uncertainty is represented as a finite set of scenarios S, each with probability π_s, and the expectation becomes a weighted sum over scenarios. As shown in Figure 1, the branching structure of a scenario tree illustrates how each period's demand uncertainty compounds over a multi-period horizon.

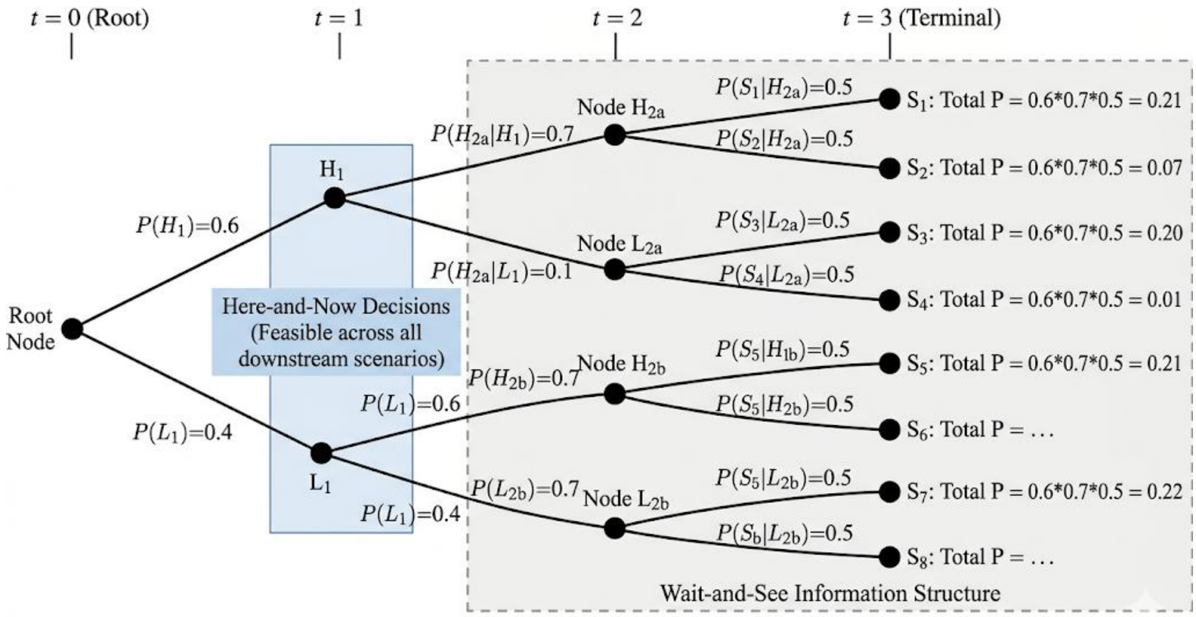

Figure 1: Scenario Tree for a Three-Period Stochastic Dynamic Location Problem

This diagram shows a scenario tree with time periods t=1, t=2, and t=3 along the horizontal axis. From the root node (t=0), the tree branches into two child nodes at t=1, each representing a different demand realization. Each of those branches further at t=2, and again at t=3, yielding eight terminal scenarios. Each branch is labeled with a transition probability, and terminal nodes show total scenario probability (product of branch probabilities along the path). The key insight is that facility location decisions made at t=1 must remain feasible across all eight downstream scenarios, illustrating the fundamental "here-and-now versus wait-and-see" trade-off in stochastic multi-period models.

Robust optimization offers an alternative view. Instead of specifying demand distributions, it assumes demand belongs to an uncertainty set $\mathcal{U}$ and seeks a solution that performs well in the worst case:

$$\min_{y,x} \max_{\tilde{h} \in \mathcal{U}} \sum_{t \in T} \left[\sum_{j \in J} f_j^t y_j^t + \sum_{i \in I} \sum_{j \in J} d_{ij} \tilde{h}_i^t x_{ij}^t \right]$$

For polyhedral uncertainty sets, the inner maximization can be dualized and the overall problem reformulated as a tractable mixed-integer linear program (Snyder, 2006). The price of robustness is conservatism — solutions are designed to survive worst-case scenarios, sometimes at significant cost to average performance.

IV. Distributed Service Networks: Structure and Key Properties

The distributed service network function as a unified system which connects multiple independent facilities through its interdependent nodes and their communication links for delivering services to customers spread across different geographical areas. The Internet content delivery networks and regional public health systems and humanitarian logistics chains and multi-depot freight operations and cloud computing platforms all exhibit this network structure.

The three characteristics of distributed service networks create distinct operational capabilities which an ordinary multi-facility system lacks.

The first point shows how the nodes work together through their system interface. The system redirects incoming requests from the overloaded node to the nearby adjacent node. The node performance depends on the activities performed by other nodes in the network system. The system operates through two processes which manage load distribution from busy nodes to less busy nodes while service standards depend on how the entire system routes its traffic between different nodes.

The second point proves that network topology functions as a fundamental element which determines network behavior. Some nodes use direct high-bandwidth connections to other nodes while other nodes rely on intermediary systems for communication. This network structure determines which routing methods users can select and how latency travels through the network and which nodes function as main traffic bottlenecks. The relocation of a facility affects both its distance to demand points and its function within the network's routing structure.

The cascade effect creates failures. When a key node goes offline — whether due to equipment failure or extreme weather or demand overload — neighboring nodes start receiving traffic which leads to their subsequent overload and system failures begin to spread. Distributed networks must treat resilience as a core design requirement instead of treating it as an enhancement added after the main design process. The location models which neglect to include failure propagation through their framework will create results which exist at their most complete level of achievement.

The DLP system functions as a distinct entity within distributed networks because its design characteristics differ from multi-facility systems that function on separate nodes. A facility move creates two major effects: it modifies network routing paths and it distributes load changes throughout all remaining network elements which results in new failure patterns for the entire system. The optimization needs to evaluate the complete system instead of just focusing on the component that undergoes relocation.

V. Modelling Frameworks: From Deterministic to Adaptive

5.1 Deterministic Multi-Period Models

The deterministic case — where demand trajectories across all periods are assumed known — provides the essential analytical scaffold. Snyder and Daskin (2005) show that even with known demand, multi-period location problems exhibit non-trivial behavior: the optimal policy sometimes involves closing a well-performing facility early to pre-position for cheaper future configurations. That kind of temporal strategy is invisible to static models.

In distributed networks, the deterministic dynamic model must be extended to capture routing behavior. The effective cost of serving demand at node i from facility j is not simply d_{ij} but a function of how load is distributed across the entire network. Congestion-aware formulations incorporate latency functions that grow with load, often following a Bureau of Public Roads (BPR) form:

$$\ell_j(v_j^t) = \ell_j^0 \left[1 + \alpha \left(\frac{v_j^t}{Q_j} \right)^{\beta} \right]$$

where v_j^t is the actual load at facility j in period t, ℓ_j^0 is free-flow latency, and α, β are calibrated parameters. This nonlinear cost structure turns the deterministic DLP into a mixed-integer nonlinear program, which is considerably harder to solve — but also far more realistic for congestion-prone networks.

5.2 Two-Stage Stochastic Models

The two-stage stochastic formulation cleanly separates what must be decided before uncertainty resolves from what can be adjusted after. First-stage decisions govern which facilities to open, operate, or close in each period — these commit resources before demand is observed. Second-stage decisions govern how to route demand, given whatever demand actually materialized. The first stage minimizes the sum of facility costs and the expected recourse cost from the second stage.

For distributed networks, the second-stage problem is a routing and assignment problem, and it can be solved independently for each demand scenario once first-stage decisions are fixed. This separability is what makes Benders decomposition so effective here: the first-stage master problem is solved iteratively, and second-stage subproblems generate feasibility and optimality cuts that progressively strengthen the master. Each cut is a linear inequality that approximates part of the recourse cost function from below.

Figure 2 shows the typical trade-off between expected cost and robustness level across three modelling approaches — deterministic, stochastic, and robust — applied to a five-node distributed service network over a ten-period horizon.

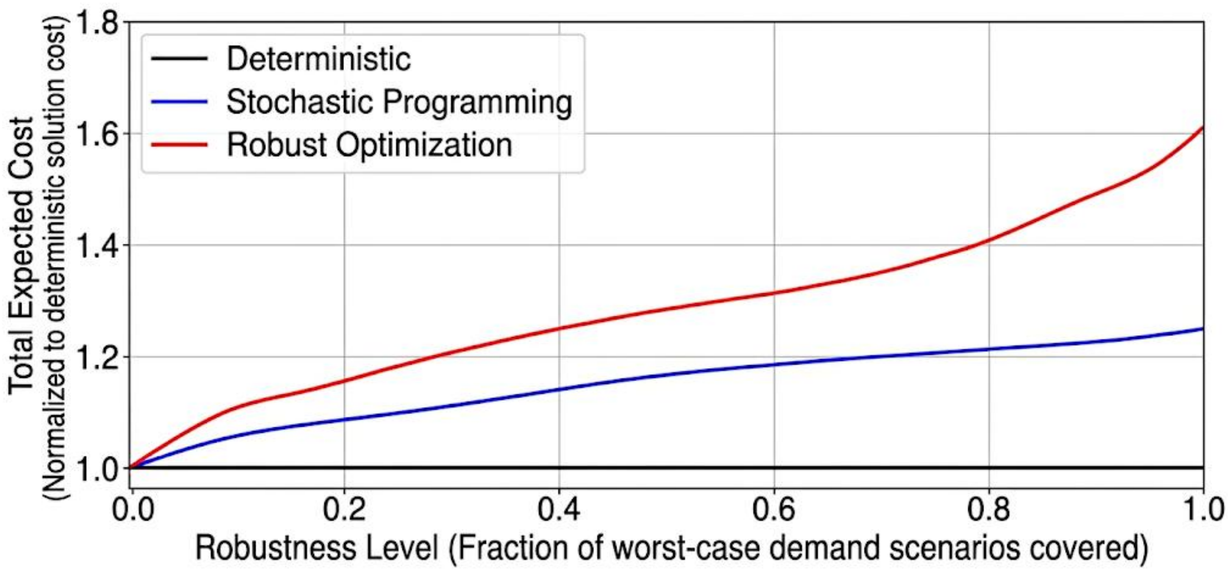

Figure 2: Cost-Robustness Trade-off Across Dynamic Location Modelling Approaches for a Five-Node Distributed Service Network

The figure displays three distinct curves which appear on a two-dimensional plot. The horizontal axis measures the robustness level which shows the solution's ability to handle worst-case demand scenarios from 0 to 1. The vertical axis shows total expected cost over a ten-period planning horizon which has been converted to the cost of the deterministic solution. The deterministic curve maintains a constant value near 1.0 because it only considers certain conditions and does not account for any uncertainties which results in no assurance of system reliability. The stochastic programming curve begins to ascend and reaches approximately 1.25 at the point of complete scenario implementation. The robust optimization curve increases at a faster rate which reaches approximately 1.6 when all elements are present. The main finding demonstrates that stochastic models deliver superior cost-robustness results when demand distributions are accurately defined but robust models establish better protection against worst-case scenarios which results in higher typical expenses.

5.3 Solution Methods and Computational Considerations

The dynamic location problem is NP-hard in general — even the static p-median problem is NP-hard on arbitrary networks (Kariv & Hakimi, 1979). Adding multiple time periods, stochastic demand, and inter-node dependencies in a distributed network multiplies the computational burden significantly.

Lagrangian relaxation is among the most effective exact-adjacent approaches. The linking constraints connecting facility configurations across time periods are dualized into the objective using Lagrange multipliersλ_j^t:

$$L(\lambda) = \min \sum_t \left[\sum_j \left(f_j^t y_j^t + O_j^t o_j^t + C_j^t c_j^t \right) + \sum_{i,j} d_{ij} h_i^t x_{ij}^t + \sum_j \lambda_j^t \left(y_j^t - y_j^{t-1} - o_j^t + c_j^t \right) \right]$$

The Lagrangian subproblem decomposes by period, making it tractable. A subgradient algorithm iteratively updates the multipliers, tightening the lower bound. Combined with Lagrangian heuristics to construct feasible upper bounds, this approach often produces solutions within one or two percent of optimality.

Benders decomposition is particularly well-matched to two-stage stochastic formulations. The master problem handles facility location decisions, while subproblems handle demand routing for each scenario. The algorithm alternates between solving a relaxed master problem and solving subproblems that generate Benders cuts. For large distributed networks with many scenarios and demand nodes, this can reduce computation time by orders of magnitude compared to solving the monolithic formulation directly.

When exact methods are too slow, metaheuristics step in. Tabu search, simulated annealing, and variable neighborhood search have all been successfully applied to dynamic location problems in the literature (Current et al., 2000; Melo, Nickel, & Saldanha-da-Gama, 2009). These methods sacrifice guaranteed optimality for practical tractability — which is often a reasonable trade-off when input data itself carries substantial uncertainty.

VI. Real-World Applications

6.1 Emergency Medical Services

EMS systems are one of the most compelling application domains. Ambulance positioning problems require placing vehicles so that any emergency call can receive a response within a critical time window — often 8 to 12 minutes. Demand patterns shift continuously through the day, and ambulances can be repositioned dynamically as calls are received. Dynamic models that update positioning recommendations in response to observed demand have been shown to improve response times significantly compared to fixed deployment plans. Brotcorne, Laporte, and Semet (2003) provide a thorough review of ambulance location models, many of which incorporate explicit temporal dynamics and probabilistic demand representations.

6.2 Cloud and Edge Computing Infrastructure

Content delivery networks face exactly the distributed dynamic location problem at scale. An operator managing hundreds of edge servers must decide periodically which servers to activate or decommission in response to shifting streaming load. Viral content creates demand spikes that are geographically concentrated but temporary — a server farm that makes economic sense during a major sports event may be unnecessarily expensive in the off-season. Opening and closing server capacity carries real costs in provisioning, licensing, and migration overhead. The DLP, with its explicit treatment of opening and closing costs, maps naturally onto these decisions. Aghezzaf (2005) discusses related capacity planning problems with uncertain demand that share the same mathematical structure.

6.3 Supply Chain and Logistics

In multi-echelon supply chains, distribution center locations are revisited every few years as trade corridors shift, markets expand, and customer expectations for delivery speed intensify. A retailer entering new markets faces a genuinely dynamic location problem embedded in a distributed logistics network where depot positions jointly determine delivery cost and speed across the entire network. Fleischmann, Ferber, and Henrich (2006) describe this kind of problem in the context of European distribution network redesign, showing that ignoring temporal dynamics leads to suboptimal solutions with measurably higher total cost.

VII. Conclusion

The Dynamic Location Problem in distributed service networks sits at the productive intersection of two mature modeling traditions — multi-period facility location and network design — and that intersection turns out to be especially rich territory. The core message is simple: when service demand evolves and facilities operate as interdependent nodes in a network, static, isolated location models are not just suboptimal — they are structurally misspecified. They optimize for a world that does not exist.

The mathematical tools for doing better are well-developed. Integer programming formulations capture the essential trade-offs between relocation costs and demand tracking. Stochastic programming and robust optimization address demand uncertainty in principled and computationally tractable ways. Lagrangian relaxation and Benders decomposition extend these formulations to problems of practical scale. Metaheuristics handle the cases where exact methods run out of steam.

What the field now needs is better integration between these models and the data environments in which real location decisions are made. Demand forecasts, failure probability estimates, cost parameters — all of these are noisy, and the models need to be more openly honest about that. Real-time adaptive methods, capable of updating location recommendations as new demand observations arrive, represent an important frontier. So does the development of multi-stakeholder frameworks that acknowledge the competitive and political dimensions of location decisions in distributed public and private service networks.

The ambulance dispatcher repositioning vehicles before the evening rush, the CDN operator spinning up edge servers ahead of a streaming peak, the logistics manager rethinking her depot network before a market expansion — all of them are solving a version of the dynamic location problem. Getting it right matters, practically and directly. The models are there. The challenge now is making them easier to trust, faster to run, and better connected to the messy, shifting realities of the networks they describe.

References

[1]. Aghezzaf, E. H. (2005). Capacity planning and warehouse location in supply chains with uncertain demands. *Journal of the Operational Research Society*, *56*(6), 730–740. https://doi.org/10.1057/palgrave.jors.2601847

[2]. Albareda-Sambola, M., Fernández, E., & Nickel, S. (2012). Multiperiod location-routing with decoupled time scales. *European Journal of Operational Research*, *217*(2), 248–258. https://doi.org/10.1016/j.ejor.2011.09.022

[3]. Brotcorne, L., Laporte, G., &Semet, F. (2003). Ambulance location and relocation models. *European Journal of Operational Research*, *147*(3), 451–463. https://doi.org/10.1016/S0377-2217(02)00364-8

[4]. Church, R., &ReVelle, C. (1974). The maximal covering location problem. *Papers in Regional Science*, *32*(1), 101–118. https://doi.org/10.1007/BF01942293

[5]. Current, J., Ratick, S., &ReVelle, C. (2000). Dynamic facility location when the total number of facilities is uncertain: A decision analytic approach. *European Journal of Operational Research*, *110*(3), 597–609. https://doi.org/10.1016/S0377-2217(97)00303-2

[6]. Drezner, Z., & Hamacher, H. W. (Eds.). (2002). *Facility location: Applications and theory*. Springer.

[7]. Erlenkotter, D. (1978). A dual-based procedure for uncapacitated facility location. *Operations Research*, *26*(6), 992–1009. https://doi.org/10.1287/opre.26.6.992

[8]. Fleischmann, B., Ferber, S., & Henrich, P. (2006). Strategic planning of BMW's global production network. *Interfaces*, *36*(3), 194–208. https://doi.org/10.1287/inte.1050.0187

[9]. Geoffrion, A. M., & Graves, G. W. (2003). Multicommodity distribution system design by Benders decomposition. *Management Science*, *20*(5), 822–844. https://doi.org/10.1287/mnsc.20.5.822

[10]. Hakimi, S. L. (1964). Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations Research*, *12*(3), 450–459. https://doi.org/10.1287/opre.12.3.450

[11]. Kariv, O., & Hakimi, S. L. (1979). An algorithmic approach to network location problems. Part II: The p-medians. *SIAM Journal on Applied Mathematics*, *37*(3), 539–560. https://doi.org/10.1137/0137041

[12]. Laporte, G., Nickel, S., & Saldanha da Gama, F. (Eds.). (2015). *Location science*. Springer.

[13]. Lim, M. K., Daskin, M. S., Bassamboo, A., & Chopra, S. (2010). A facility reliability problem: Formulation, properties, and algorithm. *Naval Research Logistics*, *57*(1), 58–70. https://doi.org/10.1002/nav.20385

[14]. Melo, M. T., Nickel, S., & Saldanha-da-Gama, F. (2009). Facility location and supply chain management: A review. *European Journal of Operational Research*, *196*(2), 401–412. https://doi.org/10.1016/j.ejor.2008.05.007

[15]. Nickel, S., & Puerto, J. (2005). *Location theory: A unified approach*. Springer.

[16]. Owen, S. H., & Daskin, M. S. (1998). Strategic facility location: A review. *European Journal of Operational Research*, *111*(3), 423–447. https://doi.org/10.1016/S0377-2217(98)00186-6

[17]. Snyder, L. V. (2006). Facility location under uncertainty: A review. *IIE Transactions*, *38*(7), 537–554. https://doi.org/10.1080/07408170500517256

[18]. Snyder, L. V., & Daskin, M. S. (2005). Reliability models for facility location: The expected failure cost case. *Transportation Science*, *39*(3), 400–416. https://doi.org/10.1287/trsc.1040.0107

[19]. Wesolowsky, G. O. (1973). Dynamic facility location. *Management Science*, *19*(11), 1241–1248. https://doi.org/10.1287/mnsc.19.11.1241