Theory and Methodology

Applying the flow-capturing location-allocation model to an authentic network: Edmonton, Canada

M. John Hodgson [a,b], K.E. Rosing [c,*], A. Leontien G. Storrier [a]

[a] *Department of Earth and Atmospheric Sciences, University of Alberta, Edmonton, Canada T6G 2H4*
[b] *Tinbergen Instituut Rotterdam, Erasmus Universiteit Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, Netherlands*
[c] *Vakgroep Toegepaste Economie, Sectie Economische Geografie, Erasmus Universiteit Rotterdam, P.O. Box 1738,
3000 DR Rotterdam, Netherlands*

Abstract

Traditional location-allocation models aim to locate network facilities to optimally serve demand expressed as weights at nodes. For some types of facilities demand is not expressed at nodes, but as passing network traffic. The flow-capturing location-allocation model responds to this type of demand and seeks to maximize one-time exposure of such traffic to facilities. This new model has previously been investigated only with small and contrived problems. In this paper, we apply the flow-capturing location-allocation model to morning-peak traffic in Edmonton, Canada. We explore the effectiveness of exact, vertex substitution, and greedy solution procedures; the first two are computationally demanding, the greedy is very efficient and extremely robust. We hypothesize that the greedy algorithm's robustness is enhanced by the structured flow present in an authentic urban road network. The flow-capturing model was derived to overcome flow cannibalization, wasteful redundant flow-capturing; we demonstrate that this is an important consideration in an authentic network. We conclude that real-world testing is an important aspect of location model development.

Keywords: Flow-capturing; Heuristics; Network; Transportation

1. Introduction

Location-allocation (LA) models endeavour to optimally locate systems of facilities and allocate demand to them. Traditional distance minimizing and covering models view demand as being exerted by weights occurring at nodes. Demand in networks is however not always expressed at nodes; for some types of facilities, demand is exerted by traffic flowing between origins and destinations. Examples are convenience stores, gasoline stations, banking machines, and billboards. Location scientists have only recently recognized flow-based demands (Goodchild and Noronha, 1987), and even more recently formulated location-allocation models for serving them (Hodgson, 1990; Berman et al., 1992; Hodgson and Rosing, 1992).

The locations and service areas of optimal facility systems function as an interacting whole.

---

* Corresponding author. Van Polanenpark 23, 4791 BP Klundert, Netherlands.

Locating facilities piecemeal can lead to damaging spatial competition among facilities, colourfully termed *cannibalization*. In a traditional coverage situation, cannibalization would arise from locating facilities so close together that they cover some demands multiply at the cost of not covering other demands at all. Cannibalization of flows in networks can also occur; two facilities may cover the same flow at different locations in the network. Point-based LA models were designed specifically to avoid this source of system degradation; so too are the new flow-based LA models.

Hodgson (1990) and Berman et al. (1992) independently developed a location-allocation model aimed at flow-based demand, for which we use Hodgson's term *Flow-Capturing Location-Allocation Model* (FCLM). Both papers prove nodal locations to be optimal and observe the model's structural identity to the node-based maximal covering location problem (MCLP) (Church and ReVelle, 1974). Hodgson (1990) gives a detailed discussion and example of the flow cannibalization problem and how the model avoids it, as well as assessing the robustness of a greedy heuristic. Berman et al. (1992) prove the worst case performance characteristics of the same greedy heuristic, and provide a branch and bound solution technique. Hodgson and Rosing (1992) consider the trade-off between the traditional *p*-median and flow-capturing goals. These papers are introductory and exploratory, based on small contrived networks.

Berman et al. (1992) demonstrated and tested their model and algorithms with two small test problems designed to illustrate their proofs as well as with a number of randomly generated problems. Neither the network structures nor the flows were meant to be realistic representations of transportation situations. Hodgson (1990) used a 25 node test network with random link lengths and node populations. Flows on this network were generated with a gravity model to provide a degree of flow structure realism.

Real-world transportation systems differ from previously used test problems in two important ways. Urban transportation networks are much larger than those which have been used to study the FCLM. This fact, and the model's great com-

putational demands, raise the question of our ability to solve it effectively in realistic settings. Perhaps more importantly, real-world transportation systems are not randomly structured. Urban residential, workplace, and shopping distributions possess considerable order, and the corresponding transportation networks and travel patterns result from complex processes in reaction to that order. Travel patterns tend to focus on central business districts, and movement tends to converge from local streets onto larger arterial roads and expressway systems. Such order will alter the operation and results of the FCLM substantially, as flow cannibalization is bound to be more damaging in a highly rational travel network than in a random one.

In this paper, we explore the characteristics of the FCLM in a realistic transportation setting. We concentrate particularly on an examination of solution procedures for the FCLM and on the very rationale for the model: the potential for flow cannibalization if the FCLM's location-allocation capability is not exploited. The setting we use is the road network, with morning-peak traffic patterns, of the city of Edmonton, Canada.

2. The flow-capturing location-allocation problem

In a network with n nodes, the FCLM considers all n nodes to be potential facility locations, and internodal OD pairs to be the basic element of demand. Potential facility location nodes are subscripted $i = 1, \ldots, n$. Each node can in turn serve as the origin or destination of a traffic flow; there is a potential $n_q = (n)(n - 1)$ OD pairs, considering directionality. We identify their flows with a single subscript $q = 1, \ldots, n_q$. A flow is captured by locating a facility at a node through which that flow passes or where it terminates. An optimum will be obtained by locating at nodes (Berman et al., 1992; Hodgson, 1990). The model is formulated as:

Maximize:

$$Z = \sum_{q=1}^{n_q} f_q Y_q \tag{1}$$

subject to:

$$\sum_{i \in A_q} X_i \geq Y_q, \tag{2}$$

$$\sum_{i=1}^{n} X_i = p, \tag{3}$$

where:

f_q is the amount of unidirectional flow from O to D (OD_q);

p is the number of facilities to be located;

A_q is the set of nodes capable of capturing a particular OD pair q's flow, i.e.: the set of all nodes on the shortest path between that origin and destination;

Y_q = 1 if f_q is captured, 0 if not.

X_i = 1 if a facility is located at node i, 0 if not.

The objective (1) is to maximize the amount of flow captured. Flow cannot be captured unless there is at least one facility at a node capable of capturing it (2), and p facilities are located.

3. Data

Edmonton, Alberta, is a Canadian prairie city with an approximate area of 700 km^2 and a 1991 population of 616,714 (Statistics Canada, 1992). The North Saskatchewan River, spanned by nine bridges, bisects the city. The roadway system consists of 103 km of freeways and expressways, 729 km of arterials, 451 km of collectors and 1644 km of local roads, a total of 2,927 km. Many of the roadways, and two of the bridges, are one-way, which strongly influences shortest paths and makes the flow structure asymmetrical. The transportation infrastructure also embraces 12.3 km of Light Rail Transit, a truck route system, and a 143 km bicycle route system (Transportation Department, 1990).

The *City of Edmonton Forecasting and Assessment Branch* (FAB), which maintains accurate data bases for travel demand forecasting, provided data from a 1989 update. Data are aggregated into traffic zones delineated by arterial roadways, major transit corridors, neighbourhood and natural boundaries. The zones are designed to be functionally homogeneous and to conform

to federal census enumeration areas (Transportation Department, 1991).

FAB provided data for the 327 traffic zones of the full Edmonton Metropolitan Area, but we limit our analysis to the 177 municipal traffic zones. The full street network is too complex for FAB's traffic (and our) modelling efforts; it is aggregated to a system of zonal centroids and freeway and arterial road network intersections. The basic aggregation unit for local trip-making and street network data is the traffic zone; intra-zonal traffic, travel times and distances are estimated for each traffic zone's centroid. Zone centroids are connected to the transportation network by one or more feeder links which are treated as normal network links. FAB provided full two-way network topology and link lengths thereby accounting for such network characteristics as one-way streets, traffic barriers, and bridge crossings.

FAB also provided their 1989 estimate of the morning peak hour, all purpose, OD trip volume. We have no control over the accuracy of these data, but FAB claims that all fall within their 10% accuracy guideline and that most are far better than that (Transportation Department, 1989). The system of 177 zonal centroids produces a total of 31,152 ($n^2 - n$) OD pairs, but since paths with zero flows need not be considered by the FCLM and they can be removed, a total of $n_q = 23,350$ flow pairs remain.

The FCLM will capture flows on paths however they are defined. For this study, we assume away congestion and behavioral influences and consider that all travel is by the shortest path. We pared the network to remove all nodes and links which do not fall on the shortest paths between non-zero OD pairs. A total of 703 nodes and 2,198 links (Fig. 1) remained in the network from original totals of 1,722 nodes and 3,743 links. The problem remains very large for some of the solution procedures.

The network and flow data provide a realistic situation for demonstrating and testing the FCLM. The aggregation might pose a problem, but because of the nature of traffic to bundle onto major arteries, the FCLM would be unlikely to locate facilities on the lowest level (aggregated

away) streets. Shortest path flow assignment is not realistic. More sophisticated assignment procedures would require us to make assumptions about congestion, turn penalties, and other aspects of the transportation network which would be equally difficult to support and would distract from our intent. We are cognizant, however, of a fundamental difficulty with the assumption of shortest paths. In a congested network, traffic seeks alternative routes, which are beyond the

ability of normal traffic assignment methods to identify uniquely. Dealing with the probabilistic nature of traffic assignment in a flow-capturing framework remains a challenging research problem. Ignoring traffic from outside the 177 node system is also worrisome; this traffic would likely alter our numerical results. This should, however, not unduly compromise our results relating to demonstration of cannibalization problems and the testing of solution procedures.

Fig. 1. Transportation network, Edmonton, Canada.

4. Solution procedures

The FCLM is NP-hard (Berman et al., 1992) and, because n_q increases roughly as the square of n, real-world problems are very large. The municipality of Edmonton, a relatively small city, yields 703 nodes and 23,350 flows; the Berman et al. (1992) branch and bound routine required several hours of IBM PC (*sic*) time for a problem with 100 nodes and 100 flows. It would be very desirable if an efficient algorithm were shown also to be robust. We consider three solution procedures with this desire in mind.

4.1. An exact solution procedure

For smaller problems, $p = 1, \ldots, 15$; $n = 703$, $q = 23350$; we solved the FCLM optimally on a Convex 210 using a fully vector version of the LAMPS (1991) mathematical programming package. LAMPS relaxes the 0/1 condition, solves the problem as a continuous linear program, and then resolves non-binary solution variables through branch and bound. The continuous programs were solved quickly and easily, never requiring more than 29 CPU minutes, but the resolution of the fractions required a huge computational effort which increased as a function of p (Table 1). These times, and the potential need to solve still larger problems, beseech the use of a more efficient, yet robust, solution procedure. Non-the-less the optimal solutions are valuable, as far as they can be obtained, to provide a baseline for the consideration of the later heuristic solutions.

The FCLM, as is the MCLP, is a form of set-covering problem (where not all sets are covered). The linear program's complex fractional nature can be explained by reference to earlier work on the set-covering problem (SCP). Rosing et al. (1992) and Rosing and ReVelle (1993) demonstrated that, with the SCP, fractions can occur when the sets that compose the solution space overlap and that the greater the degree of overlap, the higher is the likelihood of fractions. In the MCLP, the covering distance limits the membership of each covering set to local nodes; in the FCLM, membership is a function of overall

Table 1
Computational characteristics, integer programming solution

p	Continuous solution time (s)	Branch and bound time (s)	Number of nodes resolved	Total solution time (s)
2	5353	15	0	5369
3	4202	15	0	4217
4	189	15	0	204
5	897	122	0	1019
6	725	16	0	741
7	1280	2329	4	3609
8	1173	3866	8	5039
9	865	8589	12	9454
10	1446	25137	48	26583
11	1738	57221	164	58959
12	1102	165811	342	166913
13	1700	93863	226	95563
14	1658	212134	376	213792
15	1516	469770	962	471286

network flow structure rather than of propinquity, is more prolific and spatially diffuse. In the FCLM, Λ_q is the set of nodes capable of capturing OD pair q's flow, i.e.: the set of all nodes on the shortest path between O and D. Except for a small set of terminal nodes, all nodes are on many shortest paths; one is on the empirical maximum of 3,829 shortest paths, and thus a member of 3,829 covering sets, the median number of sets per node is 359. Thus, each node's membership in the sets Λ_q is large and spatially widespread; the complex overlapping of these shortest path sets results in highly fractional solutions to the relaxed LP. The objective of the FCLM is to maximize the flows captured once. When p is small, the facilities are well spread and the chosen sets tend not to overlap. As p increases, the overlapping can no longer be avoided, it becomes intricate, and fractions proliferate.

4.2. A vertex substitution procedure

We solved larger problems, $p = 16, \ldots, 50$, with a vertex substitution heuristic (VSH) of the Teitz and Bart (1968) type, typified as an H-3 interchange by Love et al. (1988). This algorithm is known to perform well on median problems (Rosing et al., 1979; Cornuejols et al., 1977), but its

performance on covering problems in general and on the FCLM in particular is not known. As with any heuristic, the optimality of any solution is not guaranteed; our worst trial VSH solutions, discussed below, confirm that the algorithm does fail. We adopted the accepted procedure of running the model many times from random starting solutions, hoping to avoid poor local maxima. We used 50 random starts for all values of p except multiples of 5 ($p = 5, 10, \ldots, 50$), for which we used 200 starts.

To test this algorithm's robustness, we consider its performance against the known optima for $p = 2, \ldots, 15$. The VSH identified the optimal solution for all p on at least 24 percent of attempts. For $p = 2, \ldots, 6$ the VSH was optimal on all attempts, and its worst trial performance was always better than 97 percent of optimal. Worst trial performance showed no obvious relationship with p, but the proportion of attempts which found the optimum declined with increasing facility numbers (Table 2). Because this deterioration reduced our trust in the robustness of the VSH, we took two further approaches to assessing that robustness.

Because each run of the VSH begins from a random starting solution, it is totally independent of all other runs, thus each series of runs for a different value of p is also independent of each other series of runs. Thus a series with a particular p, in which one or more very good (perhaps even the optimal) solution has been found, could

well be followed by a series, with $p + 1$, in which the best solution is very bad. Based on this observation we can identify two characteristics of the FCLM VSH solutions which can signal gross departures of the heuristic from optimality. Firstly, until all flow has been captured, the addition of an additional facility (moving from p to $p + 1$) in the *optimal* solutions will result in an increase in the amount of flow captured. In the VSH, as p increases, the best found values of the objective function must also monotonically increase unless that of $p + 1$ deviates considerably further from optimality than does that of p. The values of the best objective functions found by the VSH do so increase indicating that there are no gross deviations from optimality. Second, in the optimal solutions, the addition of a $p + 1$st facility cannot provide a greater benefit than did the addition of the pth. Thus, as p increases, the marginal improvement in the objective function of optimal solutions must be equal or decreasing; capturing returns to the number of facilities is non-increasing. The objective functions produced by the VSH exhibit this pattern. Thus, although optimality is not, and cannot be, proven, the VSH results do not violate the numerical expectations of optimality.

Recent literature provides another way to evaluate the results of the VSH. Because the starts of each problem are independent and randomly generated, the set of objective functions can be treated statistically (Van Dijk, 1990). We would most like to be able to state the probability that the best solution obtained by the VSH is optimal; unfortunately we cannot make such statements. We can, however, by applying Generalized Extreme Value (GEV) theory (De Haan, 1976; Buysland, 1989), make a statement about the likelihood of obtaining a better result than any yet obtained. Using the GEV approach, one estimates the parameters of the probability density function of a random variable in the vicinity of an extreme value (maximum or minimum) from observed values. One uses these parameters in turn to make probability judgements about further observations of this variable. Here, we fit the PDF of objective function values obtained with the VSH, and use the parameters of this PDF to

Table 2
Performance of the vertex substitution heuristic

Number of facilities	Number of tries on which optimum found	Worst case as proportion of optimum
7	49/50 (0.98)	0.9849
8	44/50 (0.88)	0.9937
9	44/50 (0.88)	0.9941
10	114/200 (0.57)	0.9942
11	27/50 (0.54)	0.9946
12	25/50 (0.50)	0.9703
13	14/50 (0.28)	0.9888
14	12/50 (0.24)	0.9913
15	52/200 (0.26)	0.9782

estimate the probability that the next single run of the heuristic will give a better solution than the best obtained so far. More details of the theory and application of this process are presented by Rosing and Van Dijk (1993).

For the GEV based upon 50 trials all experimental probabilities are less than 0.01 with the exception of $p = 17$ (0.01304), 19 (0.01065), 20 (0.01152), 21 (0.01179), 34 (0.01457), and 49 (0.01007). For $p = 20$ we also ran 200 trials and in this case the probability of the next run improving the best so far is 0.00119; indeed for all values of p where 200 starts were made, the GEV is always less than 0.0023. Considering all evidence, we conclude that, with multiple starts, the VSH is very robust and that its results are excellent and probably approach the true optimal very closely.

The robustness of the VSH is obtained at great computational expense. Computation times on the Convex 210 rise linearly with p from an average 55.8 seconds for $p = 2$ to 2,895 seconds at $p = 50$ (Fig. 2). The need for multiple runs can exert very high computational demands; our 200 runs at $p = 50$ required 579,100 seconds (6.7 days) of CPU time. The number of runs required to ensure desired levels of robustness is a topic which demands further research. In this paper, we turn to consideration of the robustness characteristics of a greedy algorithm.
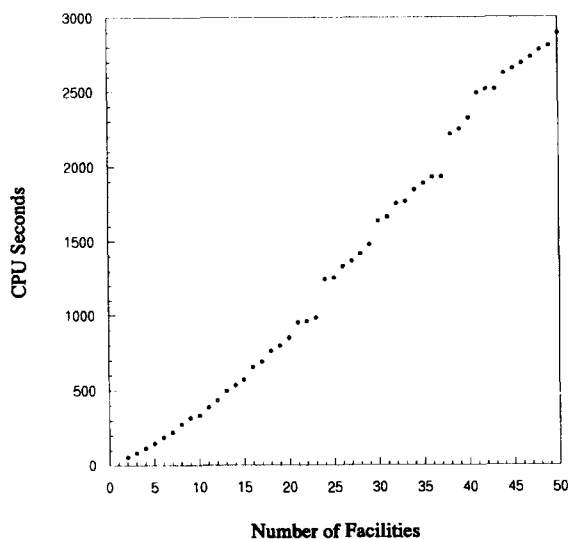

Fig. 2. Computation times, vertex substitution heuristic.

4.3. A greedy solution procedure

Computation times were high for the VSH, and inordinately so for our exact solution procedure, beyond $p = 15$. Hodgson (1990) used a greedy algorithm based on that of Kuehn and Hamburger (1963) which locates facilities one by one, selecting the location at any stage which will capture the most as yet uncaptured flow. It is implemented by, at each stage, locating a facility at the node with the most traffic, and removing the flows that location captures from further consideration. This heuristic involves no combinatorial manipulation, simply p applications of the locate and remove operation, and is thus computationally very efficient. The entire run, for $p = 1, \ldots, 50$ required only 352 CPU seconds on the Convex, and 912 elapsed seconds on a 486/33. Hodgson (1990) demonstrated that on a small test problem this heuristic was quite robust; here, we consider the greedy algorithm's robustness in dealing with the Edmonton problem.

We test the greedy algorithm's performance by comparison with the best available solution, henceforth termed the *best solution*. For $p = 1, \ldots, 15$ this is the known optimum; for $p > 15$, it is the best VSH solution we discovered. The greedy heuristic performs very well. For $p = 1, \ldots, 12$ it generates the optimal solution, for $p = 13, \ldots, 15$ a solution within 0.44 percent of optimum, and beyond that a solution which is never more than 0.9 percent less than the best VSH solution (Table 3, Fig. 3). In all cases, the greedy algorithm produces a solution which is as good as, or better than, the worst case VSH solution.

Berman et al. (1992) derived a worst case performance indicator for the greedy heuristic by which:

$$\frac{Z^* - Z^g}{Z^*} \leq 1, \tag{4}$$

where Z^* is the optimal flow captured and Z^g is the flow captured by the greedy heuristic.

We compare the greedy heuristic's performance on the Edmonton data to this indicator, using the best value obtained with the VSH as Z^* for $p > 15$. The effect of a sub-optimal VSH

value would be to underestimate the worst case performance of the greedy heuristic and thus overestimate its performance on our test problem; this effect is minor here, where the VSH solution appears to be an excellent approximation of the optimal.

Table 3
Robustness of greedy heuristic: percentage of total flows captured

p	Solution				Percentage of best solution obtained by	
	Greedy	Best [a]	Worst trial VSH	Greedy worst case [b]	Worst trial VSH	Greedy
1	9.6	9.6	-	-	-	100.0
2	18.3	18.3	18.3	9.1	100.0	100.0
3	26.8	26.8	26.8	17.8	100.0	100.0
4	33.0	33.0	33.0	24.6	100.0	100.0
5	38.1	38.1	38.1	30.3	100.0	100.0
6	43.1	43.1	43.1	35.6	100.0	100.0
7	47.1	47.1	46.3	40.0	98.3	100.0
8	50.7	50.7	50.4	43.9	99.4	100.0
9	54.2	54.2	53.9	47.7	99.4	100.0
10	57.1	57.1	56.8	50.8	99.5	100.0
11	59.9	59.9	59.6	53.7	99.5	100.0
12	62.6	62.6	60.7	56.5	97.0	100.0
13	64.7	65.0	64.2	58.9	98.8	99.7
14	66.8	67.0	66.5	61.1	99.3	99.6
15	68.6	68.9	67.3	63.0	97.7	99.6
16	70.3	70.7	69.8	64.9	98.7	99.4
17	71.9	72.4	71.2	66.6	98.3	99.3
18	73.3	74.0	72.3	68.2	97.7	99.1
19	74.7	75.4	73.6	69.6	97.6	99.1
20	76.1	76.8	75.4	71.0	98.2	99.1
21	77.4	78.1	76.7	72.3	98.2	99.1
22	78.6	79.3	77.9	73.4	98.2	99.1
23	79.8	80.2	79.1	74.3	98.6	99.4
24	80.9	81.4	80.4	75.5	98.8	99.3
25	81.8	82.3	81.0	76.3	98.4	99.4
26	82.8	83.3	82.2	77.2	98.7	99.5
27	83.7	84.1	82.8	78.0	98.5	99.6
28	84.6	85.0	83.9	78.8	98.7	99.5
29	85.4	85.8	84.6	79.5	98.6	99.6
30	86.2	86.6	85.6	80.2	98.8	99.5
31	86.9	87.3	86.4	80.9	99.0	99.5
32	87.6	88.0	86.9	81.5	98.8	99.5
33	88.2	88.6	87.6	82.0	98.9	99.5
34	88.8	89.2	88.5	82.5	99.2	99.5
35	89.3	89.8	88.9	83.0	99.0	99.5
36	89.8	90.3	89.6	83.4	99.2	99.5
37	90.4	90.8	90.1	83.9	99.2	99.5
38	90.8	91.3	90.6	84.2	99.2	99.5
39	91.2	91.8	91.1	84.6	99.2	99.4
40	91.6	92.3	91.1	85.0	98.7	99.3
41	92.0	92.7	91.9	85.3	99.1	99.3
42	92.4	93.0	92.4	85.5	99.4	99.3
43	92.8	93.4	92.7	85.8	99.3	99.3
44	93.1	93.8	93.1	86.0	99.3	99.3

Table 3 (continued)

p	Solution				Percentage of best solution obtained by	
	Greedy	Best [a]	Worst trial VSH	Greedy worst case [b]	Worst trial VSH	Greedy
45	93.5	94.1	93.4	86.3	99.3	99.3
46	93.8	94.5	93.8	86.5	99.3	99.3
47	94.1	94.8	93.9	86.7	99.8	99.3
48	94.4	95.0	94.4	86.8	99.4	99.3
49	94.6	95.3	94.6	87.0	99.3	99.3
50	94.9	95.6	94.8	87.1	99.2	99.3

[a] Optimal for $p = 1,...,15$; VSH for $p > 15$.
[b] Berman et al. (1992).

The greedy algorithm's performance clearly exceeds the worst case value derived from Berman et al. (1992) (Fig. 3). We suggest that the algorithm's excellent performance is at least in part due to the structure of flows within the network. Hodgson's (1990) conclusion that the greedy heuristic is very robust for the FCLM is borne out in the Edmonton situation. We tentatively advance the hypothesis that in real-world networks, with actual traffic flows, the greedy heuristic's robustness will far exceed its theoretical limits. Further research in other real-world settings should test this assertion.

5. Flow cannibalization

A unit of flow passing along a link or through a node cannot be considered in isolation at that location. Rather, the flow has a spatial extent, coming from some origin and going to some destination, perhaps close together, perhaps far apart. It usually has passed, and will pass, along other links and through other nodes. As, in the MCLP, two facilities located close together may cover the same demand nodes, two facilities in a network may capture the same flows. Where flows are involved, multiple coverage (capturing) is not so easily detected, for it, like capturing set membership, is a function of overall network flow structure rather than of propinquity.

The traffic passing through a node is the sum of all individual flows for which that node lies on

the path between origin and destination. Some of these flows may share common paths of some length, the routes of others may have only this node in common. Because of this complexity, it is difficult to assess the potential for flow cannibalization visually. Hodgson (1990) steps the reader through a facility by facility flow-capturing location process, indicating at each stage which flows are covered and which remain uncovered. That demonstration's graphical devices would overly clutter Edmonton's complex network topology map; the reader who is unclear about the flow cannibalization mechanism is advised to consult Hodgson (1990). Here, we restrict cartographic
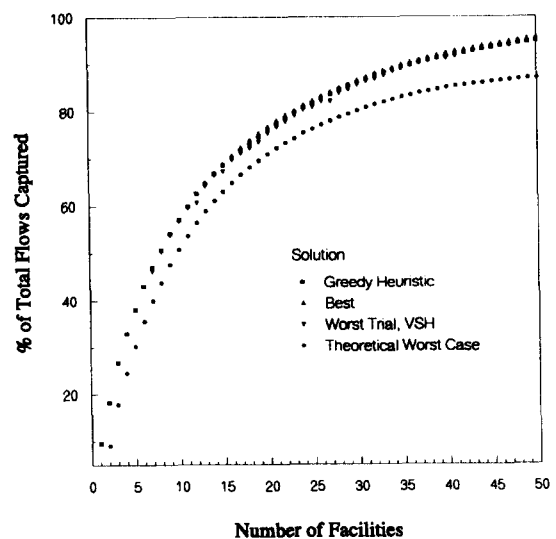

Fig. 3. Greedy heuristic robustness.

comparison to 15-facility maps derived with and without cannibalization and otherwise use numerical comparisons to indicate, for different values of p, the degree of flow cannibalization.

We adopt the approach of Hodgson (1990), who considered the consequences ignoring the cannibalization phenomenon altogether by implementing a *naive* strategy which simply locates facilities at the p nodes with the highest traffic. The degree of personification of this strategy in

the following narrative is for ease of discussion. Moreover, we do not view the naive strategy as one that a decision maker would pursue; it is a baseline of potential cannibalization effects, the hypothetical result of attempting to maximize benefit with no acknowledgment of self competition.

We demonstrate the cannibalization effect by comparing three maps of $p = 15$ solutions. The circles are proportional in area to the amount of
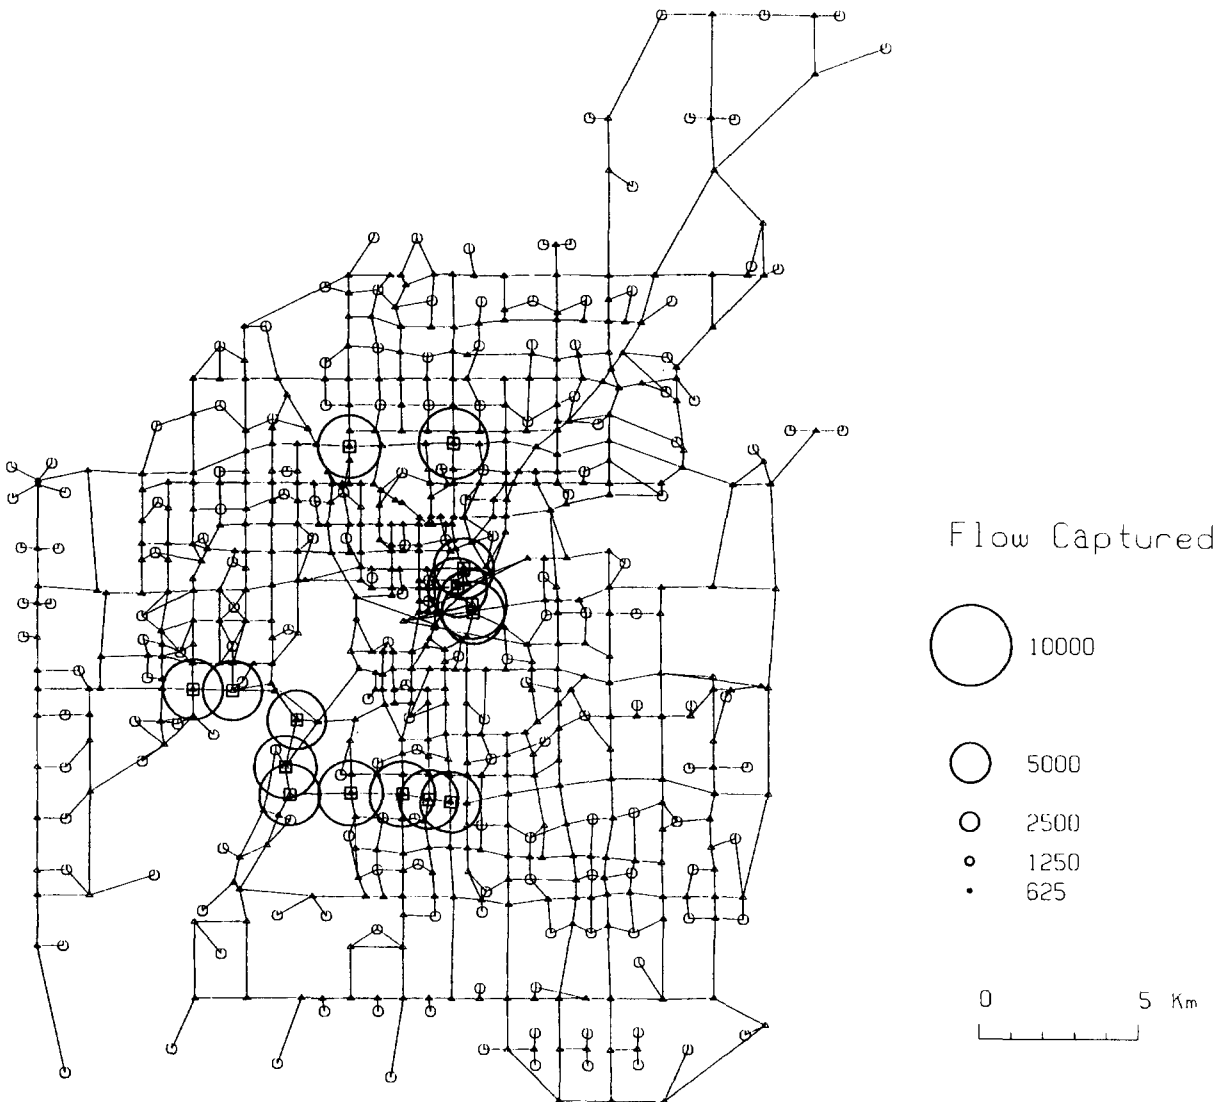
Fig. 4. Naive strategy solution, $p = 15$, expected flows captured.

flow captured at particular nodes. The naive approach clearly expects to capture much flow with 15 facilities (Fig. 4). The spatial configuration of the facilities foreshadows the cannibalization problem, however. The central cluster of facilities is in the central business district, a major destination, and thus a major area of bundling, of early morning traffic. The southwest linear arrangement of facilities is laid out along the links representing parts of the Whitemud Freeway, part of a

yet uncompleted inner ring road system, and a natural corridor for traffic bundling. In both areas, the potential for cannibalization is clear.

We can evaluate the degree of cannibalization in the naive solution by applying the greedy algorithm to that solution's facility set. Of the 15 facilities located by the naive strategy, we locate the first, remove the flows it captures, locate the next, and so on. This allows an assessment of the flows that would actually be captured by a partic-
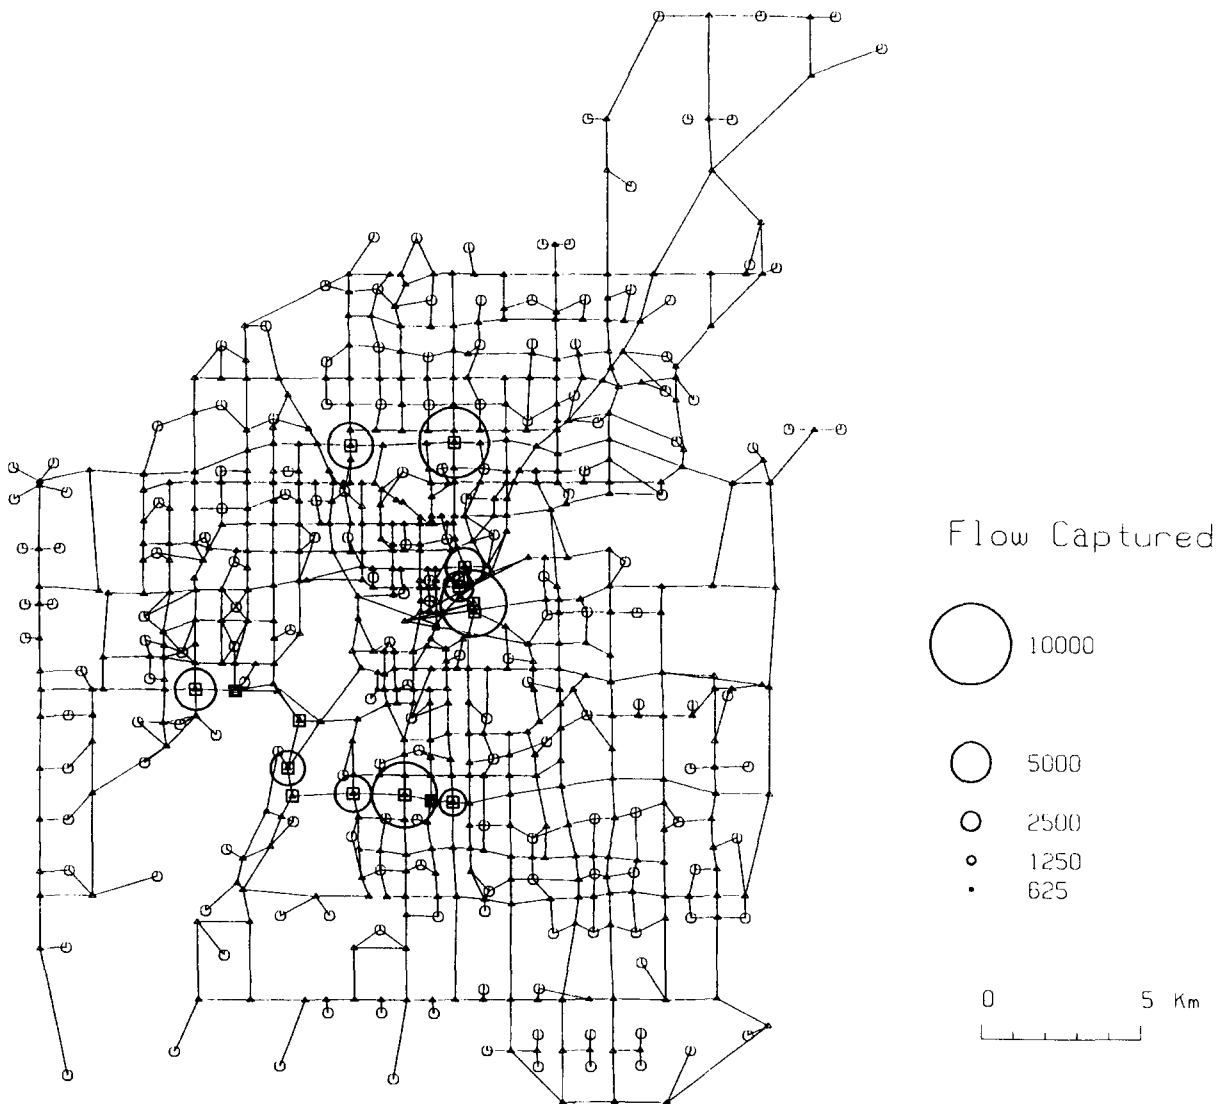

Fig. 5. Naive strategy solution, $p = 15$, actual flows captured.

ular solution. In the $p = 15$ situation, far fewer flows than expected by the naive strategy are captured (Fig. 5), with major cannibalization occurring in the two tight facility clusters (Fig. 6).

The 15-facility situation is by no means an isolated example. Total system flows are 77328. We compare the percentages of this total that the naive model expects to capture, that it actually does capture, and that the best solution captures (Fig. 7). The term *naive* is apt, this strategy expects to capture over 300 percent of the available flows with 50 facilities. It would, in fact, capture only 78.5 percent of total flows, in contrast with a VSH best of 95.6 percent. For $p = 1, \ldots, 3$ the naive strategy is optimal, beyond that, the gap between the naive strategy's expectation and attainment widens to a maximum at $p = 15$ where over 32 percent of the flows would be cannibalized under the naive strategy (Table 4). For $p = 6, 8, 15, 39, 42, 47$ and 49, the naive strategy's additional facility captures no new flows; all are instances of cannibalization. As p
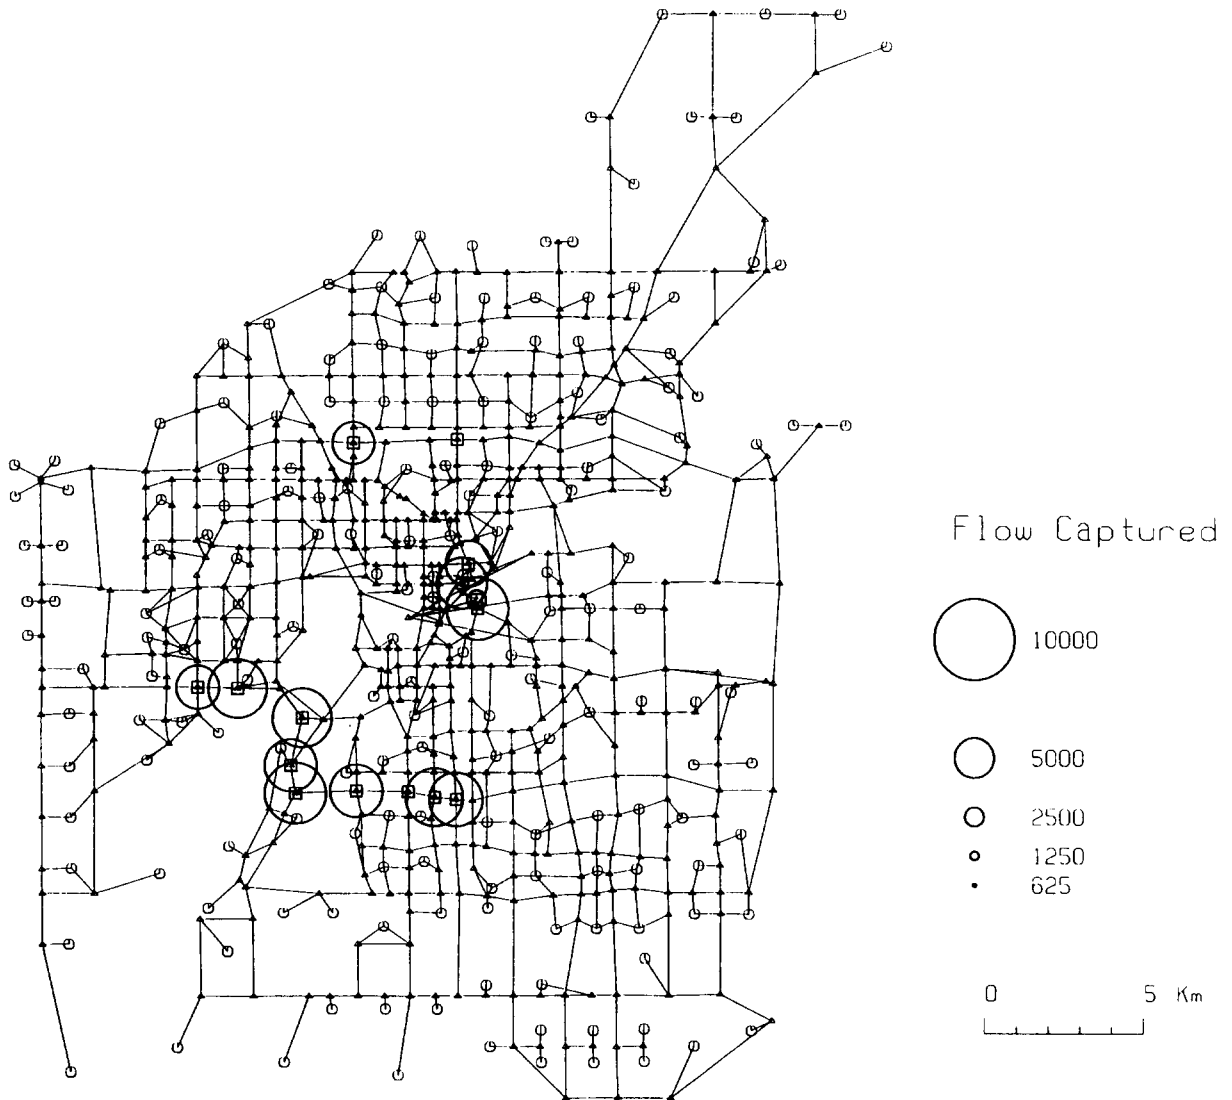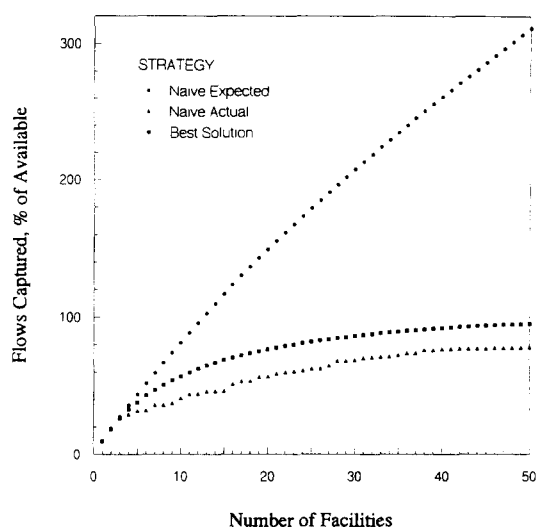

Fig. 6. Naive strategy solution, $p = 15$, flows cannibalized.

Fig. 7. Flow cannibalization, $p = 1,\dots,50$.

Table 4
Flow cannibalization

Number of facilities	Naive strategy	Best solution	% Cannibali- zation
1	9.6	9.6	0.0
2	18.3	18.3	0.0
3	26.8	26.8	0.0
4	29.5	33.0	10.7
5	31.9	38.1	16.2
6	31.9	43.1	25.9
7	36.1	47.1	23.3
8	36.1	50.7	28.8
9	37.6	54.2	30.7
10	41.0	57.1	28.2
11	44.0	59.9	26.6
12	44.2	62.6	29.3
13	46.1	65.0	28.9
14	46.2	67.0	30.7
15	46.2	68.9	32.6
16	51.4	70.7	27.0
17	53.4	72.4	25.8
18	53.5	74.0	27.1
19	56.4	75.4	24.5
20	56.8	76.8	25.4
21	59.1	78.1	23.7
22	59.9	79.3	23.7
23	60.7	80.2	23.8
24	61.2	81.4	24.3
25	62.5	82.3	23.6
26	62.8	83.3	24.2
27	64.8	84.1	22.6
28	68.2	85.0	19.4
29	68.2	85.8	20.1
30	68.9	86.6	20.1
31	70.3	87.3	19.1
32	71.0	88.0	18.9
33	71.2	88.6	19.2
34	71.3	89.2	19.7
35	72.5	89.8	18.9
36	73.9	90.3	17.8
37	73.9	90.8	18.2
38	76.3	91.3	16.0
39	76.3	91.8	16.4
40	76.9	92.3	16.1
41	77.2	92.7	16.1
42	77.2	93.0	16.5
43	77.3	93.4	16.6
44	77.5	93.8	16.8
45	77.5	94.1	17.1
46	77.9	94.5	16.9
47	77.9	94.8	17.2
48	78.0	95.0	17.3
49	78.0	95.3	17.5
50	78.5	95.6	17.2

increases and the marginal flows available decline, the naive strategy closes the gap somewhat.

In contrast to the naive strategy's facility bunching, the 15 optimally located facilities are much more dispersed (Fig. 8). Thus, well placed facilities at only 15 of the network's 703 nodes can capture 68.9 percent of all flows in the network. The marginally sub-optimal greedy solution captures 68.6 percent of total flows.

We have demonstrated that the greedy, the VSH and the optimal procedures, when compared to the simplistic naive strategy, are of a much higher quality and have argued that this is because of flow cannibalization. This contention is supported by a cartographic demonstration of the naive strategy's bunching of facilities in the central business district and on a single major transportation route. We have not, however, specifically identified the flow cannibalization process, "capturing flow more than once at the expense of flows elsewhere in the network not being captured at all" (Hodgson and Rosing, 1992). Other processes might be responsible for the naive strategy's poor performance.

We now exhibit the cannibalization process through the concept of flow capturing *redundancy*. We contend that the naive strategy captures some flows more than once, and through

this redundancy, leaves more flows uncovered than those approaches which avoid cannibalization. In the $p = 15$ solution, the naive model captures almost two percent of flows nine times, yet fails to capture 53.8 percent of flows at all (Table 5). Within a complex network flow situation, no solution can capture flows only once, but the greedy and optimal procedures capture 45.8 percent of network flows only once, and none

more than four times. The redundancy data are strongly consistent with the postulated flow cannibalization process (Fig. 9).

6. Conclusions and future research

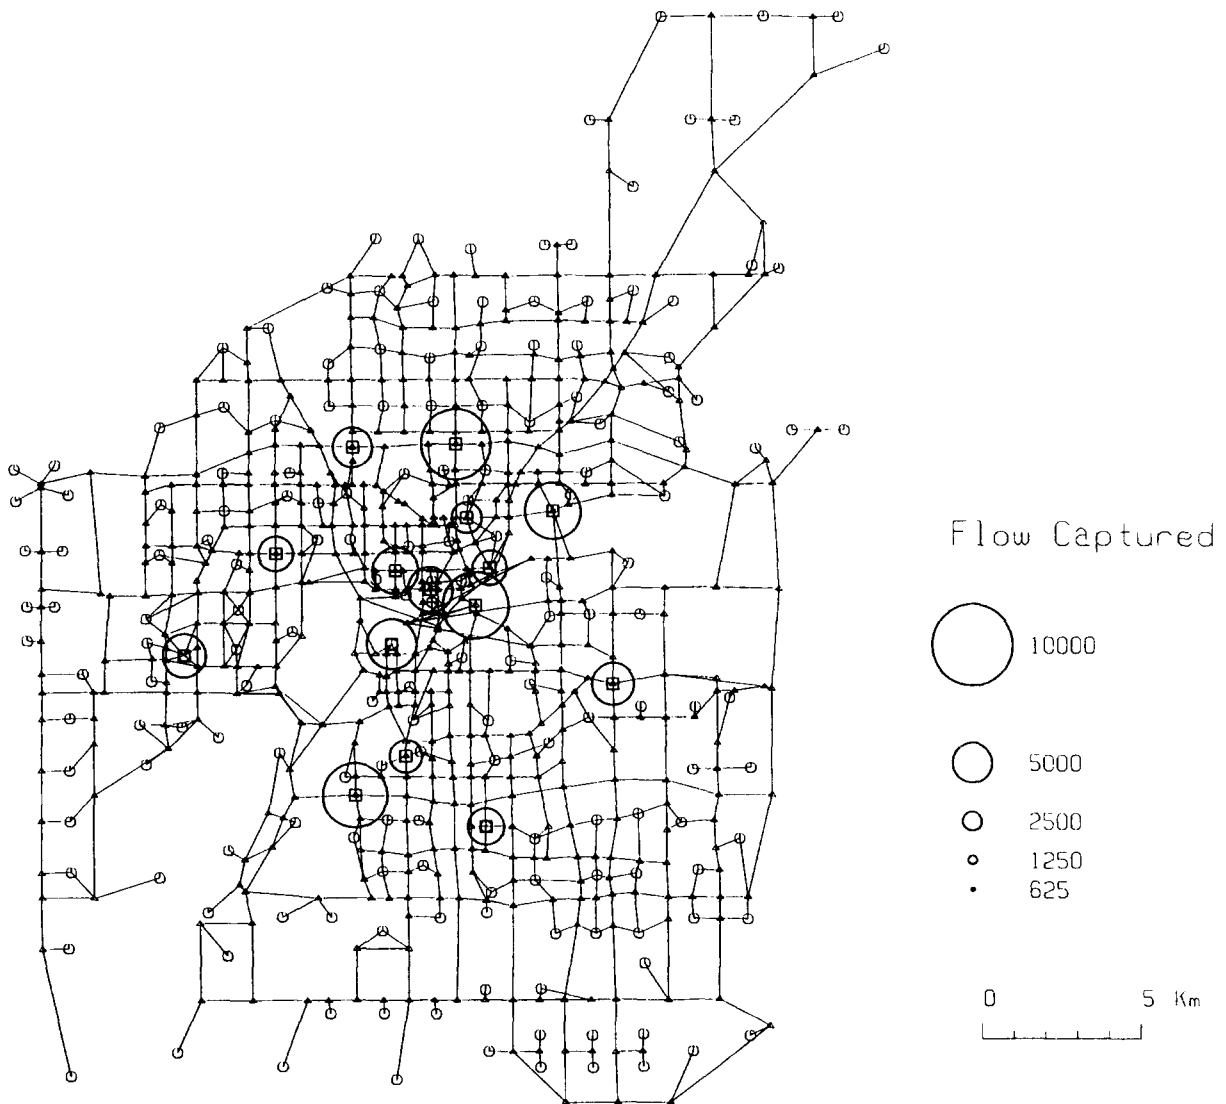The FCLM is combinatorially complex and dimensionally large; real-world problems demand

Fig. 8. Optimal solution, $p = 15$.

Table 5
Redundancy: the percentage of flows not captured, captured once, or more in $p = 15$ solution

Number of times captured	Naive	Greedy	Optimal
0	53.8	31.4	31.1
1	19.9	45.8	45.8
2	13.0	19.2	20.4
3	3.2	3.5	2.5
4	2.9	0.1	0.1
5	2.5	-	-
6	0.8	-	-
7	1.3	-	-
8	0.8	-	-
9	1.9	-	-

a robust and efficient heuristic solution procedure. Hodgson (1990) and Berman et al. (1992) suggested that a simple greedy algorithm was both very efficient and robust. We have confirmed the robustness of the greedy approach; it provides excellent solutions for real-world data, far surpassing its theoretical worst case expectations. We suspect that the specific structure of the problem, the spatial concentration of morning traffic flows, enhances the algorithm's robustness; further work could explore the mechanism

Fig. 9. Flow-capturing redundancy, $p = 15$.

through which it occurs. The FCLM provides an important additional approach to facility location. We suggest that further work be done to identify potentially better effective solution methods. For example, specific methods which would exploit the problem's structure, and, perhaps, Lagrangian relaxation remain to be explored.

We have substantiated, within a real-world situation, Hodgson's (1990) contention that flow cannibalization can badly impair efforts to locate facilities to capture flows in a transportation network. Urban traffic tends to focus upon the central business district and to become bundled on major traffic arteries, creating a situation in which multiple high-traffic nodes have many flows in common. We have demonstrated that piecemeal facility location aimed at major flow sites results in some flows being captured several times at the expense of others not being captured at all. The importance of the location-allocation approach introduced in the FCLM, which avoids such cannibalization, clearly extends to real-world problems.

For many types of service, flows are a better representation of demands than are nodal populations; the flow-capturing location-allocation model fills a need in facility location theory. We have shown that the model can provide great benefits in avoiding flow cannibalization within a real-world traffic network, and that a simple and efficient feasible solution procedure provides excellent solutions. Our findings reinforce our earlier belief that the FCLM is a worthy subject for further research and development.

Although traffic flows exert demand for many types of facility, it is likely that they would occur in tandem with point-based demands; that such facilities would serve persons who travel specifically to the facility from their homes as well as passing traffic. Hodgson and Rosing (1992) argued that a dual objective approach trading off the FCLM and the p-median model is suited to locating facilities which draw their patronage both from neighbourhood shopping and passing traffic. Their results, obtained with a small test problem, suggested that p-median solutions are more compatible with flow-capturing needs than are flow-capturing solutions with p-median needs. This
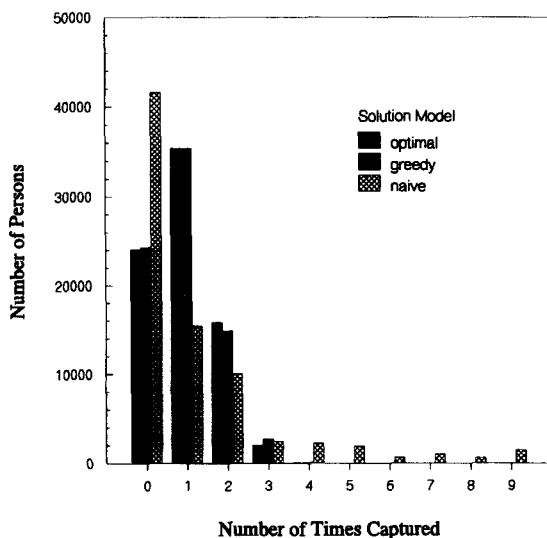
counterintuitive finding remains to be tested with a real-world network.

Working with data which purport real-world authenticity raises two important real-world questions. Traffic flows in congested networks do not follow unique paths. Current traffic assignment procedures provide traffic totals through probabilistic flow assignment which the FCLM cannot deal with. Development of probabilistic flow-capturing methods is another challenging and pressing research problem. Location-allocation modellers often ignore site suitability in the interests of providing general systemwide solutions, but in dealing with flow-capturing, site characteristics may be too important to ignore. In some situations, traffic safety considerations may prohibit location at intersections. Simple observation shows that traffic intersection locations are only a small proportion of the sites actually developed for retail activity. Nodal locations capture flows optimally; the likelihood and implications of these locations being infeasible sites should be investigated.

Flow-capturing, as it has been posed, is based on the assumption that direct exposure is required for capturing. This is certainly the case for billboards and impulse purchases, where the visibility of a facility is of utmost importance. For routine service acquisition such as dry cleaning or use of daycare facilities, some deviation from the network might be permitted (Hodgson, 1981). The flow-capturing model is dedicated to the complete eradication of flow cannibalization. For certain types of service, billboard advertising, for instance, multiple coverage may be advantageous. Hogan and ReVelle (1983)'s work on backup coverage in emergency facility location offers a route to development of these ideas. Further work should also involve capacitated facilities. Present work by Berman et al. (1995) casts the FCLM in a probabilistic framework.

The flow-capturing model is in its infancy; we expect that it will be found to have applications in other areas than those retail purposes which we have heretofore suggested. Methods of locating water treatment and highway weigh stations, for instance, may benefit from applications and extensions of the flow-capturing model. The model also seems particularly well suited to become the foundation of a system for the location of hazardous waste inspection stations.

A final important conclusion bears, not directly on the particular model we have examined, but on our use of a real-world data base to examine it. The Edmonton morning traffic data do more than provide a larger data set; they represent a system which exhibits the structure of the real-world problem that the model is designed to solve. Operations research is long on new models and algorithms, but short on real-world tests. In some instances, testing with real-world data may weaken the case for a location model, in other instances, such as this, the model's importance is enhanced. Earlier presentations of the FCLM followed the norm, demonstrating and testing the models on randomly contrived test problems. The FCLM's rationale and performance are strongly attached to the nature of the phenomenon it deals with: flow-capturing. Its justification arises from its ability to maximize flow-capturing by overcoming flow cannibalization, a very real real-world possibility, but a phenomenon which is very unlikely to occur in random data. Operations research in general and location theory in particular purport to solve real problems in the real-world; their models must be shown to function in that arena.

References

Berman, O., Larson, R.C., and Fouska, N. (1992), "Optimal location of discretionary service facilities", *Transportation Science* 26, 201–211.

Berman, O., Krass, D., and Xu, C.W. (1995), "Locating discretionary service facilities based on probabilistic customer flows", *Transportation Science* 29, 276–290.

Buysland, T.A. (1989), "Statistics of extremes in climatology", *Statistica Neerlandica* 43, 1–30.

Church, R.L., and ReVelle, C.S. (1974), "The maximum covering location problem", *Papers of Regional Science* 32, 101–118.

Cornuejols, G., Fisher, M.L., and Nemhauser, G. (1977), "Location of bank accounts to optimize floats: An analytic study of exact and approximate algorithms", *Management Science* 23, 789–810.

De Haan, L. (1976), "Sample extremes: An elementary introduction", *Statistica Neerlandica* 30, 161–172.

Goodchild, M.F., and Noronha, V.T. (1987), "Location-allocation and impulsive shopping: The case of gasoline retailing", in: A. Ghosh and G. Rushton (eds.), *Spatial Analysis and Location-Allocation Models*, Van Nostrand Reinhold, New York, 121–136.

Hodgson, M.J. (1990), "A flow-capturing location-allocation model", *Geographical Analysis* 22, 270–279.

Hodgson, M.J. (1981), "The location of public facilities intermediate to the journey to work", *European Journal of Operational Research* 6, 199–204.

Hodgson, M.J., and Rosing, K.E. (1992), "A network location-allocation model trading off flow capturing and p-median objectives", *Annals of Operations Research* 40, 247–260.

Hogan, K., and ReVelle, C.S. (1983), "Backup coverage concepts in the location of emergency services", *Modelling and Simulation* 14, 1423–1428.

Kuehn, A.A., and Hamburger (1963), "A heuristic program for locating warehouses", *Management Science* 9, 643–666.

LAMPS (1991), *Users' Guide 1.66*, Advanced Mathematical Software, London.

Love, R.F., Morris, J.G., and Wesolowsky, G.O. (1988), *Facilities Location: Models and Methods*, North-Holland, New York.

Rosing, K.E., Hillsman, E.L., and Rosing-Vogelaar, H. (1979), "A note comparing optimal and heuristic solutions to the p-median problem", *Environment and Planning, A* 11, 86–89.

Rosing, K.E., ReVelle, C.S., and Rosing-Vogelaar, H. (1992),

"Fractions in the location set-covering problem", *Environment and Planning, B* 19, 125–130.

Rosing, K.E., and ReVelle, C.S. (1993), "Integers in the location set-covering problem", *Environment and Planning, B* 20, 481–482.

Rosing, K.E., and Van Dijk, J.J. (1993), "Estimating the probability of heuristic improvement: A large scale application of extreme value theory", *Studies in Locational Analysis* 4, 301–304.

Statistics Canada (1992), *Census of Canada: Population and Dwelling Counts for Federal Electoral Districts and Enumeration Areas, 1991*, Ottawa.

Teitz, M.B., and Bart, P. (1968), "Heuristic methods for estimating the generalized vertex median of a weighted graph", *Operations Research* 16, 955–961.

Transportation Department (1991), *1990 Transportation System Status Report*, City of Edmonton, Edmonton.

Transportation Department (1990), *Working Paper: Recalibration of the City of Edmonton Regional Travel Model to 1989 Base Conditions and Development of Travel Demand Forecasts 1989 – Long Term*, City of Edmonton, Edmonton.

Transportation Department (1989), *1989 Travel Survey Working Paper: Survey Design, Organization and Analysis*, City of Edmonton, Edmonton.

Van Dijk, J.J., "Heuristics and statistics in cluster analysis", in: Fr. Orban-Ferauge and J.P. Rasson (eds.), *Proceedings of Meeting V of the Euro Working Group on Locational Analysis*, Facultés Universitaires Notre-Dame de la Paix, Namur, Belgium, 189–203.