

A SCALABLE ALGORITHM FOR LOCATING DISTRIBUTION CENTERS ON REAL ROAD NETWORKS

Saeed Ghanbartehrani, Ph.D.

J. David Porter, Ph.D.

School of Mechanical, Industrial and Manufacturing Engineering
Oregon State University, USA, 97331-6001

Abstract

The median problem is a type of network location problem that aims at finding a node with the total minimum demand weighted distance to a set of demand nodes in a weighted graph. In this research, an algorithm for solving the median problem on real road networks is proposed. The proposed algorithm, referred to as the multi-threaded Dijkstra's (MTD) algorithm, is then used to optimally locate Wal-Mart distribution centers on the 28-million node road network of the United States with the objective of minimizing the total demand weighted transportation cost. The resulting optimal location configuration of Wal-Mart distribution centers improves the total transportation cost by 40%.

1 Introduction

Location theory is a well-established and active research area. A common challenge in location theory is the *location problem* whose main objective is to find the *best* location for facilities to serve a set of demand points. The best location for a facility depends on the nature of the problem being studied, the problem's constraints, and the optimality criteria [1].

Determining the location of distribution centers plays a significant role on the efficiency, service quality, and economical sustainability of a distribution network. Many models have been proposed to optimally locate distribution centers. Mathematical programming algorithms, multi-criteria decision-making, heuristics, and simulation are among the most applied solution approaches [2].

In this paper, an algorithm for locating distribution centers on real road networks is proposed. The proposed algorithm is used to locate 78 Wal-Mart distribution centers on the continental United States road network graph (composed of over 28 million nodes) so that the total demand weighted transportation cost between each distribution center and

the set of stores it serves is minimized.

2 Literature Review

The term *location problem* refers to the modeling, formulation, and solution of a class of problems in location theory that can best be described as locating facilities in some given spaces. Location problems can be classified into the categories of *location*, *allocation*, and *location-allocation* [3]. In the context of the location problem, allocation involves assigning demand to the facilities being located.

Network location models have been applied to problems in location theory where a facility must be located on a network composed of *nodes* (or locations) and *edges* (or routes) [4]. The *median problem* (also known as the *1-median* problem) is a type of network location model introduced by Hakimi [5] whose objective is to minimize the total demand weighted distance between a set of demand points and a facility in a network. Hakimi [6] also formulated a generalized version of the median problem known as the *p*-median problem for locating *p* facilities.

Kariv & Hakimi [7] proved that the *p*-median problem is NP-hard (on general graphs) and proposed an algorithm with time complexity $O(n^2p^2)$ for solving the *p*-median problem on a tree network (i.e., a connected graph with no cycles). Perez-Brito et al. [8] proved that the optimal solution for the *p*-median problem on a network (and at least one of its corresponding spanning trees) are equal given that the objective function is non-decreasing in distance. This finding allows algorithms suited for solving location problems on tree networks to be applied to corresponding spanning trees of any cyclic network. Perez-Brito et al. [8] developed a heuristic solution process that randomly generates spanning trees and solves the location problem using a tree network solution algorithm. The authors also suggested using an algorithm developed by Tamir [9] for solving *p*-median problems, which has a time complexity of $O(pn^2)$. By adopting such an approach, the solution process may involve (in the worst case scenario) solving the location problem on all possible spanning trees of a network, which is excessively complex.

Many heuristics have also been developed for solving *p*-median problems. Reese [10] conducted a survey on solution methods developed for *p*-median problems on graphs with a minimum objective function that did not consider fixed location costs or employed mathematical programming. Many of the proposed heuristic solutions focused on efficiently solving problems on networks with up to 1,000 nodes, which are significantly larger problems compared to the problems solved by the exact algorithms reviewed.

Geographical information systems (GIS) based approaches have also been developed and applied to a variety of location problems, including the location of distribution centers, to allow for the incorporation of realistic factors such as real travel distances based on road networks, traffic, geological conditions, and so on [11]. Other GIS-based facility location approaches can be found in [12], [13], [14], and [15]. GIS based models

are not popular since they are very specific to the problem being solved and rely on internal features offered by (usually proprietary) GIS software packages.

The literature review conducted in this research indicates that none of the proposed exact algorithms for median problems are scalable to real road networks. Small real road networks are typically composed of thousands of nodes, whereas medium sized road networks may contain hundreds of thousands of nodes. Large road networks have millions of nodes (e.g., the United States road network is composed of approximately 28 million nodes). It is important to note that road networks are cyclic and, therefore, algorithms that work on trees cannot be applied to them. In contrast, approaches that involve finding spanning trees or network pre-processing are not practical for solving every single location problem due to the excessively large size of the road networks. Although proposed heuristic solutions for the median problem have significantly improved the applicability of the solution methods, they are still impractical for real road networks that contain several thousands of nodes or more. Also, the time complexity of the majority of the approaches just discussed is still an issue if optimal or close-to-optimal solutions are required.

In this research, a scalable algorithm for solving the median problem on real road networks is proposed. The proposed algorithm is not specific to the distribution center location problem and can be applied to any facility location problem based on the median model. The algorithm is used to locate 78 Wal-Mart distribution centers in the continental United States so that the total demand weighted transportation cost between each distribution center and the set of stores it serves is minimized.

3 Methodology

The main phases of the methodology aimed at optimizing the locations of Wal-Mart distribution centers are depicted in Figure 1. The boxes in Figure 1 depicted with a gray frame indicate the outputs of each of the phases of the methodology.

In the data preparation phase, data acquired from different sources are analyzed and processed in order to determine:

- 1) The locations of existing Wal-Mart distribution centers in the continental United States.
- 2) The locations of existing Wal-Mart stores in the continental United States.
- 3) The population of the urban area each Wal-Mart store is located in.
- 4) Allocation of stores to distribution centers.

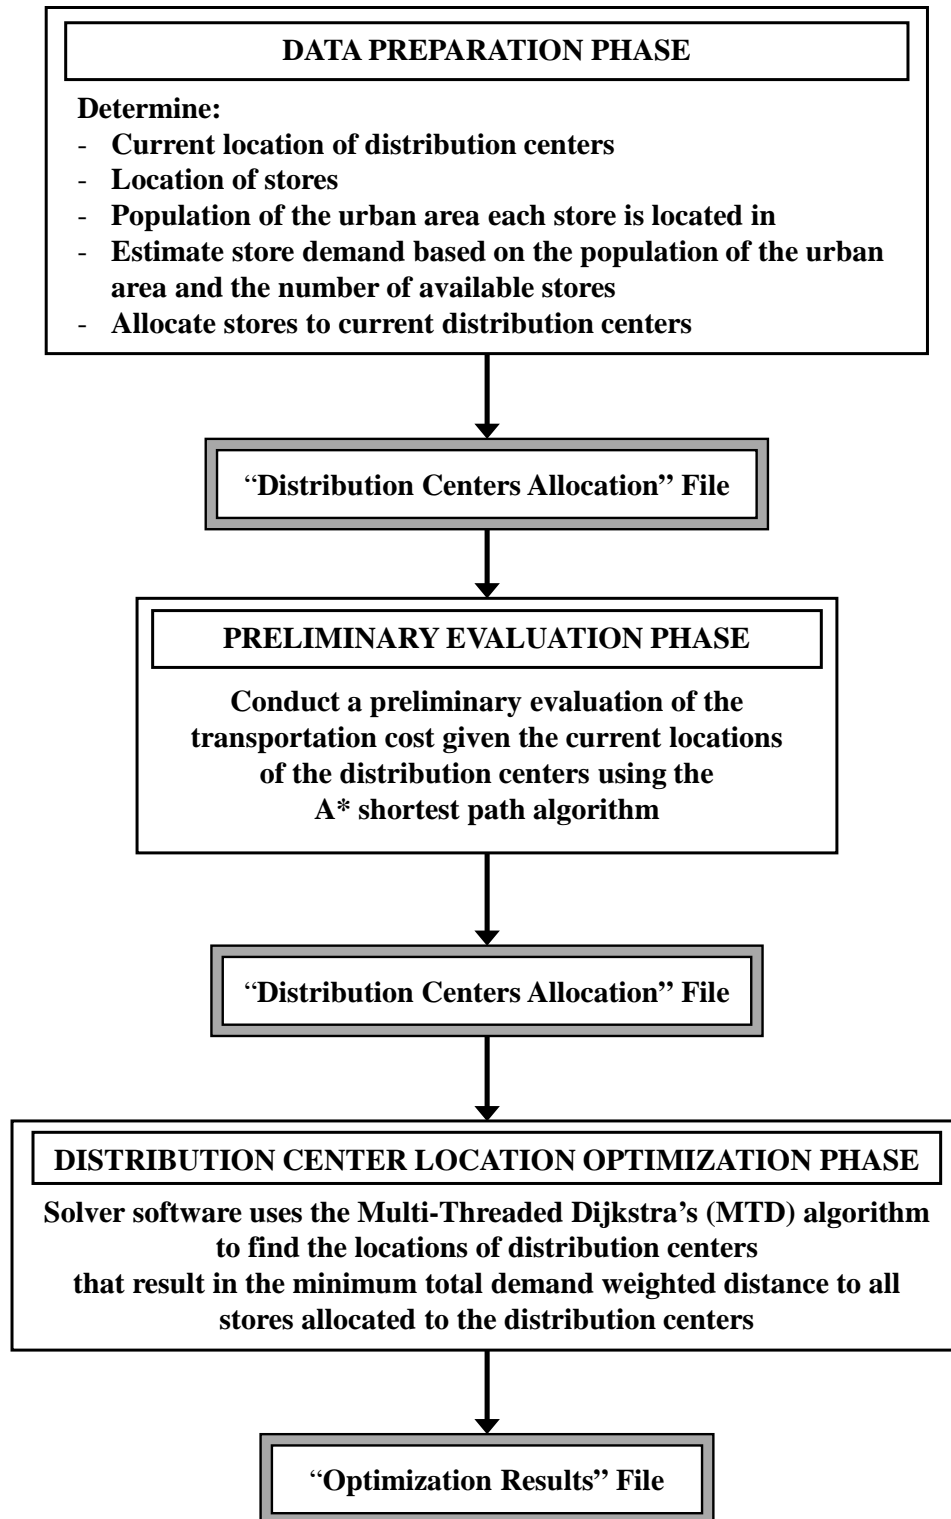


Figure 1: Main Phases of the Proposed Methodology

The output of the data preparation phase is a comma separated values (CSV) file referred to as the “Distribution Centers Allocation” file. The “Distribution Centers Allocation” file includes the following:

- 1) A list of Wal-Mart distribution centers and the stores allocated to each of these distribution centers.
- 2) The coordinates (i.e., latitude and longitude) of each Wal-Mart distribution center.
- 3) The coordinates (i.e., latitude and longitude) of Wal-Mart stores allocated to each distribution center.
- 4) The estimated demand for each Wal-Mart store based on the population of the urban area in which the store is located in.

The data included in the “Distribution Centers Allocation” file is then used to conduct a preliminary evaluation of the total transportation cost based on the current location of the Wal-Mart distribution centers and stores. The A* shortest path algorithm is used to find the shortest path between each Wal-Mart distribution center and the stores allocated to it.

Finally, the optimal location for each of the Wal-Mart distribution centers is found using a newly developed facility location algorithm whose main objective is to minimize the total transportation cost. The proposed facility location algorithm, referred to as the *multi-threaded Dijkstra’s* algorithm, is based on Dijkstra’s shortest path algorithm and it can find a point that has the minimum total weighted distance to a set of demand points.

3.1 Multi-Threaded Dijkstra’s Algorithm for Facility Location

The *multi-threaded Dijkstra’s* (MTD) algorithm is based on the bidirectional Dijkstra’s shortest path algorithm. Dijkstra’s and bidirectional Dijkstra’s algorithms find the shortest path between a *single* source and a *single* destination node in a graph. The bidirectional Dijkstra’s algorithm improves the search efficiency (when compared to Dijkstra’s algorithm) by performing a simultaneous search from both the source and destination nodes.

The MTD algorithm starts the search from *all* input nodes (demand points) and finds the location where the total demand weighted distance is minimized in the graph. The MTD algorithm is similar to Dijkstra’s algorithm and bidirectional Dijkstra’s algorithm in that all of them search the graph to find the shortest path between some nodes. However, the unique feature of the MTD algorithm is that it finds a node that has the minimum total weighted distance to a *set* of demand points. In contrast, Dijkstra’s algorithm and bidirectional Dijkstra’s algorithm find the shortest path only between two *predefined* and *fixed* source and destination nodes.

Given a road network graph G with k demand points (v_i), edge weights (c_{ij}), and demand w_i associated with a demand point v_i , the MTD algorithm finds the node in G

with the minimum total weighted distance to all demand points. The steps taken by the MTD algorithm are as follows:

1. Set the distance property for all v_i nodes to $d_i(v_i) = 0$. The distance property for all other nodes v is set to $d_i(v) = \infty$.
2. Start from node v_i and add v_i to the i^{th} open list. The open list is a priority queue. Do this step for $i = 1$ to k .
3. Select the non-empty open list that contains the smallest top element. The top element of an open list is the node with the smallest d_i property. Expand the top element of the selected open list by calculating $d_i(v)$ for all unvisited adjacent nodes v using Equation (1). Mark the expanded node as visited from demand point i .

$$d_i(v) = d_i(s) + w_i \cdot c_{sv} \quad (1)$$

4. Add the unvisited adjacent nodes v (from demand point i) and their corresponding $d_i(v)$ to the same open list as children of the expanded node. If the same node already exists in the open list, keep only the instance with the smaller d_i property.
5. Remove the expanded node from the open list. In case all d_i properties of a node v are found, set the value of the objective function for the best solution found so far (i.e., μ) to $\sum_{i=1}^k d_i(v)$. Update μ when a smaller value is found.
6. Repeat steps 3 through 5 until μ becomes smaller than all d_i values for the top elements in all open lists (referred to as the optimality criteria), or all open lists become empty. The optimality criterion is presented in Equation (2) where d_i^{top} is the top element in the i^{th} open list.

$$\mu \leq \text{Min}_{i=1, \dots, k} d_i^{\text{top}} \quad (2)$$

3.2 Data Preparation

The main data utilized in this study included Wal-Mart's store and distribution center openings from 1962 to 2006 [16]. The list of Wal-Mart distribution centers and stores was built using several data sources including Wal-Mart's website, Wal-Mart's Environmental Protection Agency (EPA) reports, and Wal-Mart's annual reports [17].

The store openings data include the store opening date and the street address of 3,176 Wal-Mart stores in the continental United States. The street addresses were converted to geocoded locations (i.e., latitude and longitude) using an online geocoding service [18]. The data for the distribution centers include the street address and the GPS coordinates for 78 Wal-Mart distribution centers in the continental United States and, therefore, did

not require geocoding. A shape file containing the continental United States national urban areas was downloaded from the US Census Bureau website [19].

The geocoded locations for all Wal-Mart stores and distribution centers, as well as the urban areas shape file, were imported into the open source, relational spatial database PostgreSQL [20]. PostgreSQL offers a spatial extension known as PostGIS used for spatial analysis [21]. The population of the urban areas associated with each Wal-Mart store was found using a query developed in PostGIS.

An important step in the data preparation phase was the allocation of Wal-Mart stores to distribution centers, which is depicted graphically in Figure 2. It is important to note that since real Wal-Mart store allocation data (i.e., which stores are served by each distribution center) was not readily available, a decision was made to allocate only the 15 closest stores to each distribution center. If real allocation data were available, a more precise allocation process could be implemented. Out of the total of 3,176 Wal-Mart stores available, only 1,064 were allocated to distribution centers and used in the analysis. This is due to the fact that some stores were allocated to more than one distribution center.

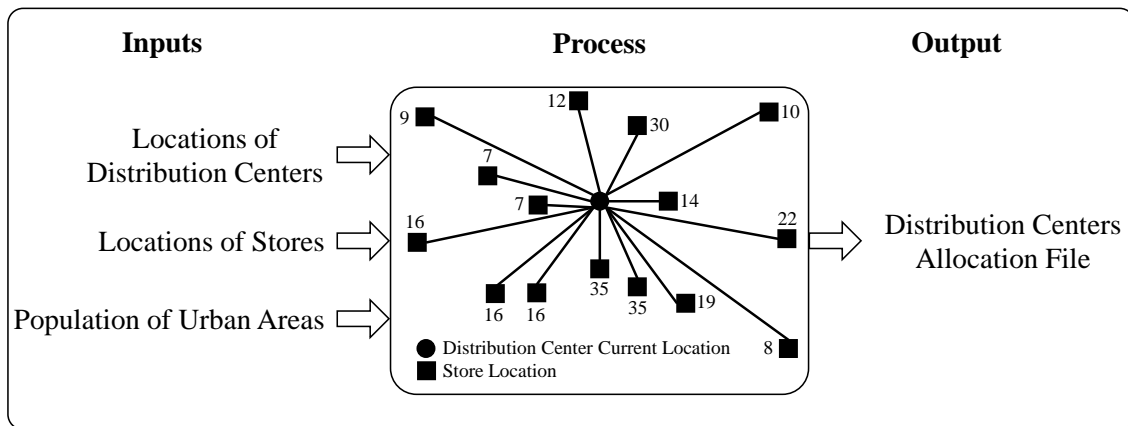


Figure 2: Inputs, Process, and Outputs of the Data preparation Phase

The demand for each Wal-Mart store (i.e., the numbers next to the store location icons in Figure 2) was estimated using Equation (3):

$$\text{Store estimated demand} = \frac{\text{Total urban area population}}{1000 \times \text{Number of stores in the urban area}} \quad (3)$$

The output of the data preparation phase is a file exported from PostgreSQL in comma separated values (CSV) format referred to as the “Distribution Centers Allocation” file. The “Distribution Centers Allocation” includes a list of all the Wal-Mart distribution centers, the 15 closest stores allocated to each distribution center, and the estimated demand for each store.

3.3 Preliminary Evaluation of the Transportation Cost

Once all the required data had been collected and organized, a preliminary evaluation of the current allocation of stores to the 78 Wal-Mart distribution centers was performed by calculating the *total transportation cost (TTC)* using Equation (4):

$$\text{Total Transportation Cost (TTC)} = \sum_{i=1}^{78} \sum_{\forall j \in S_i} d_{ij} p_j \quad (4)$$

Where:

- d_{ij} : Shortest road distance (in kilometers) from distribution center i to its allocated store j
- p_j : Estimated demand for store j
- S_i : Set of stores allocated to distribution center i

A solver software was developed in Java/JavaScript based on the open source project GraphHopper to compute the shortest road distances (i.e., d_{ij}). The solver software was also used as a basis to implement the MTD algorithm. GraphHopper is an open source, web-based routing engine developed in Java. Some of the most popular path finding algorithms such as A* and Dijkstra's are already built into GraphHopper's routing engine [22]. GraphHopper is released under the *Apache License* which allows developers to "use the software for any purpose, distribute, modify, or distribute the modified version of software without the concern of royalties" [23].

As Figure 3 illustrates, the "Distribution Centers Allocation" file was the main input to the solver software in the preliminary evaluation phase. The data in the "Distribution Centers Allocation" file were used to compute the shortest road distance from every Wal-Mart distribution center to its allocated stores using the A* shortest path algorithm.

Open source map data of the five main regions of the continental United States (i.e., Midwest, Northeast, Pacific, South, and West) were downloaded from Geofabrik & OpenStreetMap [24]. Map data files were then merged using OsmConvert [25] to produce a single map file of the continental United States. The map data of the continental United States (in OSM format) was also provided as input to the solver software.

The preliminary TTC based on the current allocation of stores to distribution centers was calculated as 4,730,759 kilometers-people.

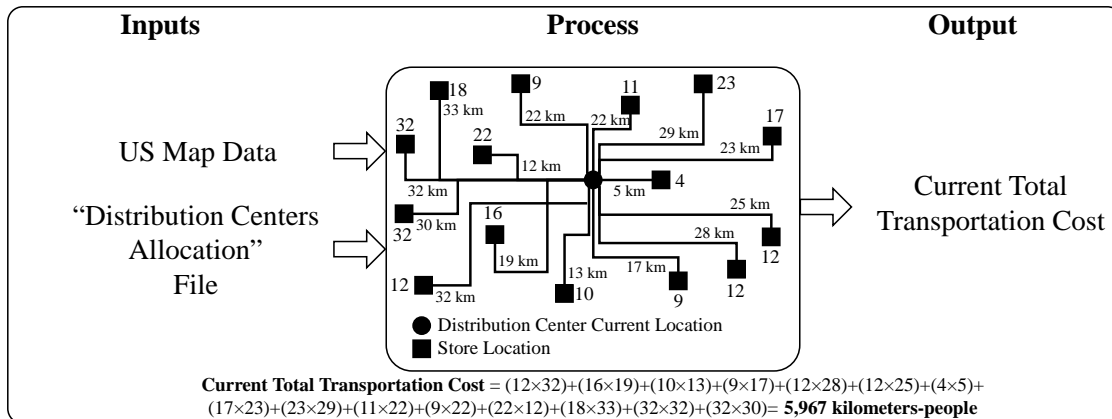


Figure 3: Inputs, Process, and Outputs of the Preliminary Evaluation Phase

3.3 Optimizing the Locations of Distribution Centers

The “Distribution Centers Allocation” file was again used as the main input to the solver software to determine the optimal location of each individual Wal-Mart distribution center. The solver software uses the MTD algorithm to determine the optimal location for each distributions center that results in the minimum total weighted distance to its allocated stores.

It is important to mention that for every input point (i.e., the location of a Wal-Mart distribution center or a store), the solver software finds the closest point on the road network graph and does all the required computations based on this point. Also, all output points (i.e., improved distribution center locations) are points on the road network graph which makes them practical distribution center location options in terms of accessibility to the road network.

The map data file of the continental United States is approximately 110 Gigabytes in size and its corresponding graph (i.e., the graph the solver software uses to solve location problems) has over 28 million nodes. To perform this analysis, an HP Compaq 8100 Elite desktop computer with a Core i7 860 processor and 16 GB of RAM was used. The solver software and solution methodology can be scaled to larger geographic areas (e.g., individual countries or a whole continent) to locate facilities on a road network graph to minimize the total demand weighted distance. Figure 4 depicts an example of optimizing the location of a distribution and calculating the improved TTC.

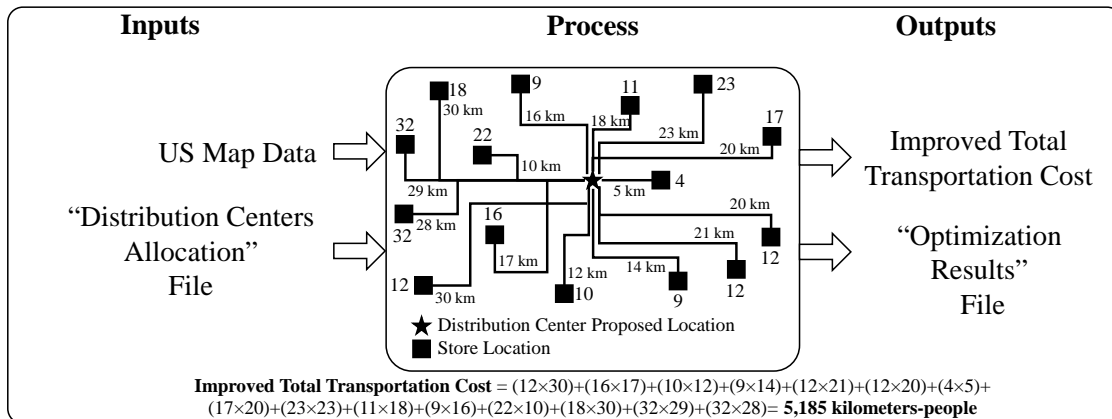


Figure 4: Inputs, Process, and Outputs of the Distribution Center Location Optimization Phase

4 Analyses and Results

Table 1 summarizes the results of analyzing the total transportation cost and transportation distance of 78 Wal-Mart distribution centers before and after relocating the distribution centers using the software implementation of the MTD algorithm. As stated earlier, the preliminary evaluation resulted in a TTC of 4,730,359 kilometers-people. After relocating the Wal-Mart distribution centers using the MTD algorithm, the TTC reduces to 2,859,738 kilometers-people. This represents an improvement in the cost of 12%.

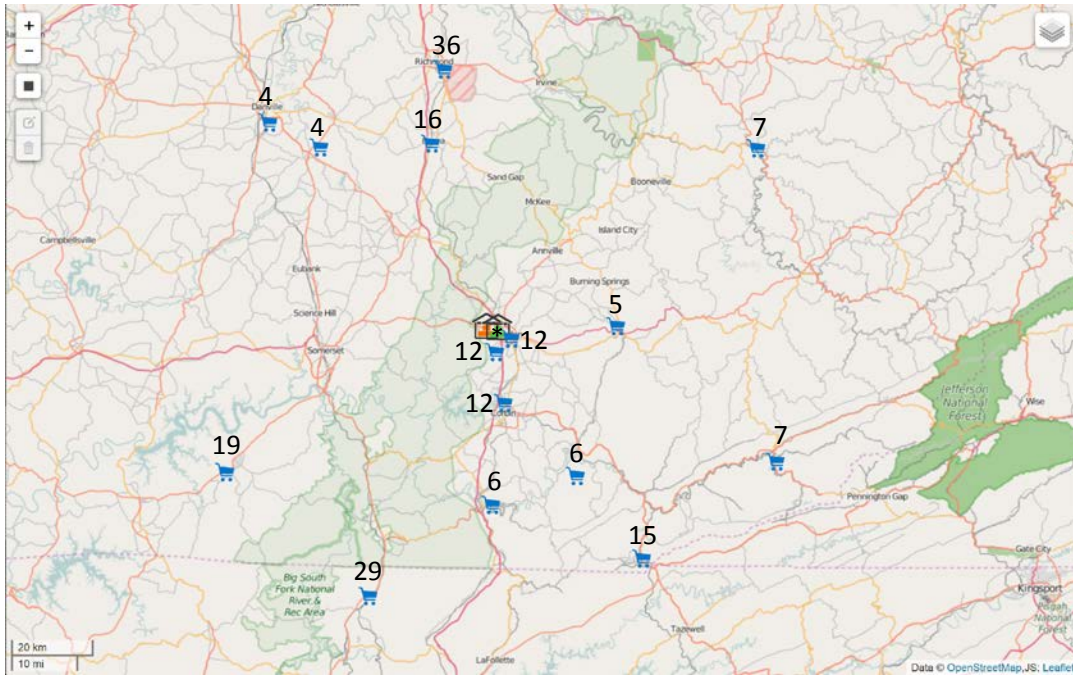
The maximum transportation distance shows an increase of 18%, which means that some Wal-Mart distribution centers have been relocated farther away from some stores. However, the average transportation distance shows an improvement of 40%. The analysis of transportation distances was performed to ensure that the proposed relocations did not result in extremely long transportation distances, which may have adverse effects on lead times and the quality of service provided to stores by their corresponding distribution center. The decrease in minimum transportation distance suggests that some of the Wal-Mart distribution centers have been relocated to serve stores with high values of demand.

Table 1: Summary of Results

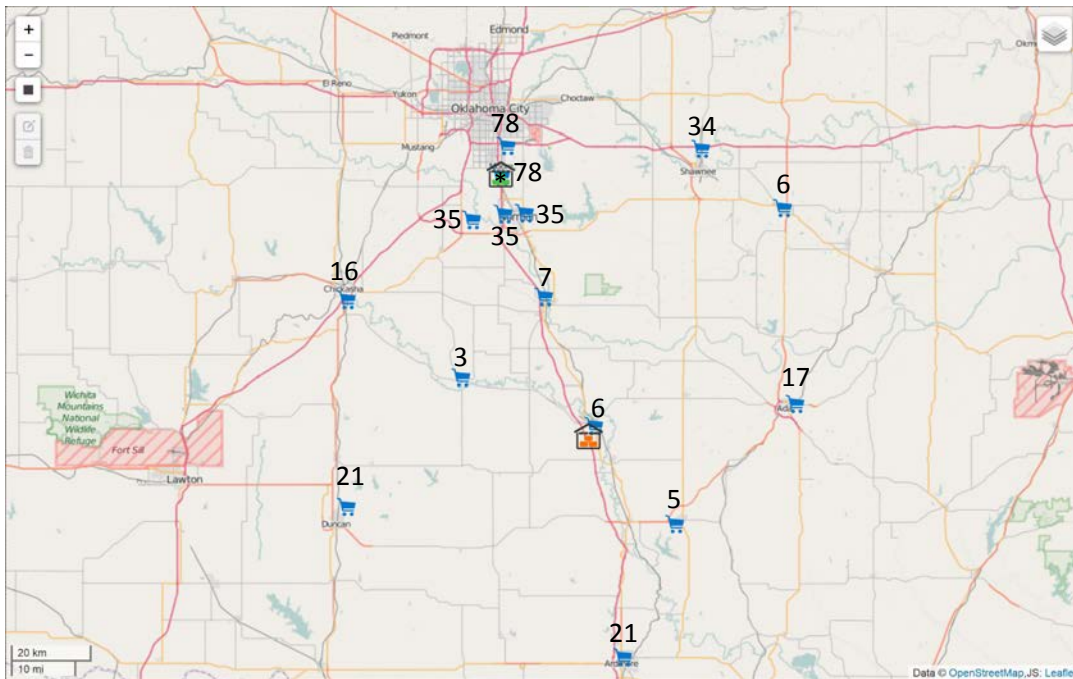
| Metric | Current | Proposed | % Change |
|--------------------------------------|-----------|-----------|----------|
| Total Transportation Cost | 4,730,759 | 2,859,738 | 40% |
| Average Transportation Cost | 4,061 | 2,455 | 40% |
| Minimum Transportation Cost | 3 | 0 | 100% |
| Maximum Transportation Cost | 90,305 | 73,726 | 18% |
| Average Transportation Distance (km) | 69 | 61 | 12% |
| Minimum Transportation Distance (km) | 0.1 | 0 | 100% |
| Maximum Transportation Distance (km) | 351 | 486 | -38% |

Figure 5 depicts two examples of recommended distribution center relocations generated by the software implementation of the MTD algorithm. The shopping cart icons represent Wal-Mart stores, whereas the house icon represents a Wal-Mart distribution center. A house icon with an asterisk (*) inside represents the proposed (optimal) location of the distribution center, whereas a house icon without an asterisk represents the current location of a distribution center. The numbers displayed above the shopping cart icons represent the demand associated with each store.

Figure 5a depicts a case in which the current location of the Wal-Mart distribution center serves its allocated stores quite well since the optimal distribution center location found by the software implementation of the MTD algorithm is only 3 km (1.8 mi.) away. In this case, the improvement in the transportation cost that resulted from relocating the distribution center is 7%. In contrast, Figure 5b depicts a case in which the current location of the Wal-Mart distribution center is too far from the proposed (optimal) location, which indicates that the current location of the Wal-Mart distribution center is too far from stores with high demand levels. This situation can be improved by relocating the distribution center to or near the recommended location. In this case, the optimal distribution center location is 71 kilometers (44 mi.) from the current location. If the Wal-Mart distribution center were relocated, it would translate into a 53% improvement in the transportation cost.



(a)



(b)

Figure 5: Examples of Distribution Center Relocations Recommended by the Software Implementation of the MTD Algorithm

5 Conclusions

In this research, a methodology to optimally locate distribution centers based on the newly developed multi-threaded Dijkstra's (MTD) algorithm is proposed.

A case study involving 78 Wal-Mart distribution centers and 1,064 stores was used to evaluate the performance of the proposed methodology. The demand for each store was estimated based on the population of urban area the store is located in. The real road network graph of the United States, composed of over 28 million nodes, was used as a basis for locating distribution centers.

The results showed that the total transportation cost improved by 40%, while the average transportation distance decreased by 12%. The software implementation of the MTD algorithm was able to optimally locate all 78 distribution centers within reasonable time (~2 hours). Several opportunities to improve the proposed methodology have already been identified:

- The Wal-Mart store and distribution center location data used in this research dates back to 2006. Newer data could be obtained to demonstrate the effectiveness of the proposed methodology on a more realistic case study.
- Demands of stores were estimated based on the population of urban areas. Real store demand data categorized by the products sold in the stores can replace the estimated values to improve the accuracy of the results.
- It was assumed that each distribution center serves the 15 closest stores, which may not be an accurate assumption in reality. This can be replaced by real allocation data based on real Wal-Mart logistics historical data.
- Finally, more complex factors such as traffic, road type, and elevation could be incorporated (along with distance) to calculate the transportation cost.

References

- [1] Laporte, G., Nickel, S., & Gama, F. S. da. (2015). *Location Science*. Springer.
- [2] Xu, X., & Liu, W. (2013). Research on Logistics Distribution Center Location Model Based On GIS. *International Journal of Digital Content Technology and Its Applications*, 7(7), 538–545. <http://doi.org/10.4156/jdcta.vol7.issue7.63>
- [3] Zanjirani Farahani, R., & Hekmatfar (Eds.). (2009). *Facility Location Concepts, Models, Algorithms and Case Studies*. Springer Verlag.
- [4] ReVelle, C. S., Eiselt, H. A., & Daskin, M. S. (2008). A bibliography for some fundamental problem categories in discrete location science. *European Journal of Operational Research*, 184(3), 817–848. <http://doi.org/10.1016/j.ejor.2006.12.044>
- [5] Hakimi, S. L. (1964). Optimum Locations of Switching Centers and the Absolute Centers and Medians of a Graph. *Operations Research*, 12(3), 450–459.

- [6] Hakimi, S. L. (1965). Optimum Distribution of Switching Centers in a Communication Network and Some Related Graph Theoretic Problems. *Operations Research*, 13(3), 462–475.
- [7] Kariv, O., & Hakimi, S. L. (1979). An Algorithmic Approach to Network Location Problems. II: The p-Medians. *SIAM Journal on Applied Mathematics*, 37(3), 539–560. <http://doi.org/10.1137/0137041>
- [8] Perez-Brito, D., Mladenovic, N., & A. Moreno-Perez, J. (1998). A note on Spanning Tree for Network Location Problems. *Yugoslav Journal of Operations Research ISSN: 0354-0243 EISSN: 2334-6043*, 8(1). Retrieved from <http://yujor.fon.bg.ac.rs/index.php/journal/article/view/459>
- [9] Tamir, A. (1996). An $O(pn^2)$ algorithm for the p-median and related problems on tree graphs. *Operations Research Letters*, 19(2), 59–64. [http://doi.org/10.1016/0167-6377\(96\)00021-1](http://doi.org/10.1016/0167-6377(96)00021-1)
- [10] Reese, J. (2005). Methods for solving the p-median problem: An annotated bibliography. Retrieved from http://digitalcommons.trinity.edu/math_faculty/28/
- [11] Liang, Y., Zhang, X., & Sun, Q. (2008). Study on logistics distribution center location based on GIS environment. In *IEEE International Conference on Service Operations and Logistics, and Informatics, 2008. IEEE/SOLI 2008* (Vol. 2, pp. 3007–3010). <http://doi.org/10.1109/SOLI.2008.4683050>.
- [12] Yeh, A. G., & Chow, M. H. (1996). An integrated GIS and location-allocation approach to public facilities planning—An example of open space planning. *Computers, Environment and Urban Systems*, 20(4–5), 339–350.
- [13] Li, X., & Yeh, A. G. (2005). Integration of genetic algorithms and GIS for optimal location search. *International Journal of Geographical Information Science*, 19(5), 581–601. <http://doi.org/10.1080/13658810500032388>.
- [14] Panichelli, L., & Gnansounou, E. (2008). GIS-based approach for defining bioenergy facilities location: A case study in Northern Spain based on marginal delivery costs and resources competition between facilities. *Biomass and Bioenergy*, 32(4), 289–300.
- [15] Fiorese, G., & Guariso, G. (2010). A GIS-based approach to evaluate biomass potential from energy crops at regional scale. *Environmental Modelling & Software*, 25(6), 702–711.
- [16] Holmes, T. J. (2010, March). Wal-Mart Data for “Diffusion of Wal-Mart...” Retrieved May 16, 2016, from <http://users.econ.umn.edu/~holmes/data/WalMart/>
- [17] Holmes, T. J. (2011). The Diffusion of Wal-Mart and Economies of Density. *Econometrica*, 79(1), 253–302. <http://doi.org/10.3982/ECTA7699>.
- [18] Dotsquare LLC. (2014). Geocodio. Retrieved May 18, 2016, from <https://geocod.io>
- [19] US Census Bureau, G. D. (2010). 2010 TIGER/Line® Shapefiles. Retrieved May 18, 2016, from <https://www.census.gov/cgi-bin/geo/shapefiles/index.php?year=2010&layergroup=Urban+Areas>
- [20] PostgreSQL. (2016). PostgreSQL: The world’s most advanced open source

- database. Retrieved May 18, 2016, from <http://www.postgresql.org/>
- [21] PSC. (2016). PostGIS — Spatial and Geographic Objects for PostgreSQL. Retrieved May 18, 2016, from <http://postgis.net/>
- [22] Karich, P. (2015, August 12). GraphHopper - OpenStreetMap Wiki. Retrieved November 6, 2015, from <http://wiki.openstreetmap.org/wiki/GraphHopper>
- [23] Apache Software Foundation. (2004, January). Apache License. Retrieved March 7, 2014, from <http://www.apache.org/licenses/LICENSE-2.0.html>
- [24] Geofabrik, G., & OpenStreetMap. (2016). OpenStreetMap data for North America. Retrieved May 18, 2016, from <http://download.geofabrik.de/north-america.html>
- [25] OpenStreetMap. (2016, March 7). Osmconvert - OpenStreetMap Wiki. Retrieved May 18, 2016, from <http://wiki.openstreetmap.org/wiki/Osmconvert>