Deploying Vaccine Distribution Sites for Improved Accessibility and Equity to Support Pandemic Response

CS4991 Capstone Report, 2023

Ann Li
Computer Science
The University of Virginia
School of Engineering and Applied Science
Charlottesville, Virginia, USA
al7gc@virginia.edu

ABSTRACT

In response to COVID-19, states such as Virginia have deployed mobile vaccination centers to distribute vaccines targeting two key values for improved pandemic response: accessibility and equity. We formulate a combinatorial problem that captures these factors, study the inherent hardness of the problem through strong impossibility results, and then develop efficient computational algorithms with theoretical guarantees on both of these factors. Finally, we run computational experiments on synthetic real-world population movement data from two different Virginia municipalities to show the efficacy of our algorithms. From our experiments, we conclude that our proposed algorithms generally yield a significant improvement over natural baselines. Additionally, we demonstrate that our algorithms compute a tradeoff between the maximum distance to a vaccination clinic and the number of clinics; this naturally enables us to give a recommendation to the government on the most cost-effective budget policy with important implications for the future of public health decision-making. Having demonstrated the significance of our problem formulation, we suggest that a natural next step for future work is to extend other variants of the facility location problem to this setting as well.

1. INTRODUCTION

As of February 2022, only 64% of the eligible population is fully vaccinated in the United States [2]. Furthermore, there is a significant disparity in vaccination rates between demographics—the rate among Whites was 1.2 times that of African Americans and 1.1 times that of Hispanic people. The reasons why some people have not been vaccinated include distrust and skepticism regarding COVID-19, accessibility issues, and concerns about the cost [1]. Lottery schemes, mandates, vaccine clinics, and other strategies have been implemented to increase the vaccination rate with varying levels of success. Since cost and accessibility remain a challenge for a fraction of the population, especially minorities and people in poorer neighborhoods, mobile vaccine clinics have been an important part of the public health response strategy of government agencies. In this paper, we study the problem of deploying mobile vaccine administration sites with the goal of improving the accessibility of vaccines to individuals. Deploying vaccination clinics is a form of a facility location problem [5, 9], referred to as the k-supplier problem, in which a limited set of k facilities needs to be placed so that every person (i.e., a client) is “close” to a facility; a common metric to measure closeness is the maximum distance between a client and their closest facility, though many other notions have been studied. Facility location problems are well understood, and efficient approximation algorithms and practical heuristics exist. However, deploying vaccine clinics leads to a novel facility location problem (referred to as the MobileVaccClinic problem) since people (clients) are mobile rather than stationary. Suppose each person p visits a set Sp of locations during the day; then it suffices to deploy a clinic close to at least one location in Sp. Our contributions are the following:

- We formalize the MobileVaccClinic problem for modeling the deployment of mobile vaccine clinics in a way that takes into account human mobility patterns (by considering the distance to a facility from any of the locations visited by a person), fairness (by requiring that at least a fraction of people in each demographic group have a nearby clinic), outliers (by allowing partial coverage), and capacity constraints (by restricting the number of people assigned to each clinic). We show that this problem is much harder than the standard k-supplier problem and getting any bounded polynomial-time approximation to the minimum distance is not possible, thus motivating bicriteria algorithms.

• We design two approximation algorithms and extend our algorithms to have fairness guarantees in both the original and the outliers formulation of the problem.

• We evaluate our algorithms for a realistic population of two municipalities in Virginia. The shortcomings of only considering a client's home (rather than their entire traveling route) emphasize the importance of our problem formulation. Additionally, our algorithms allow us to compute a tradeoff between the maximum distance to a clinic and the number of clinics; this naturally enables us to give a recommendation to the government on the most cost-effective budget policy.

2. RELATED WORKS

Due to its applications in a large number of domains, facility location and broader location theory is a very well-studied area; see, e.g., the surveys by [3, 5, 9]. The general goal in this family of problems is to deploy facilities to provide the best possible service to a set of clients. A huge number of objectives have been considered, along with a plethora of variations such as fairness variants and online or stochastic versions. The MobileVaccClinic problem we study here is a generalization of the well-known k-center problem, where the goal is to open at most k centers while minimizing the maximum distance of a point to its closest center. For this simple clustering setting, there exist efficient 2-approximation algorithms [10, 12]. Furthermore, it is shown that unless P=NP this is the best achievable approximation ratio [13]. Location theory problems have also been considered in the area of healthcare, e.g., [3, 8, 17, 18]. A lot of this work has been focused on placing mobile clinics or temporary facilities to ensure good service, especially in resource-poor countries. As mentioned in [3], the healthcare domain poses new challenges for location theory, such as uncertainty, reliability, operation efficiency, patient safety, and cost-effectiveness. Prior work has generally not considered the mobility of clients at a detailed scale, which provides more flexibility in deploying facilities. Our formulation of MobileVaccClinic explicitly models human mobility, thus providing a realistic framework for public health agencies in their response efforts.

3. PROJECT DESIGN

In our paper, we introduce a new variant of the facility location problem that follows a recent line of work on integrating the mobility patterns of the population into disease models [6, 20]. We will use the distance from a vaccination center as the metric for defining accessibility. The key change, however, is that clients will be represented by a set of locations that they visit (within a time period) instead of just one point. Though this will make the problem much harder to solve efficiently, it will more strongly correlate with the likelihood of a person going to a vaccine center.

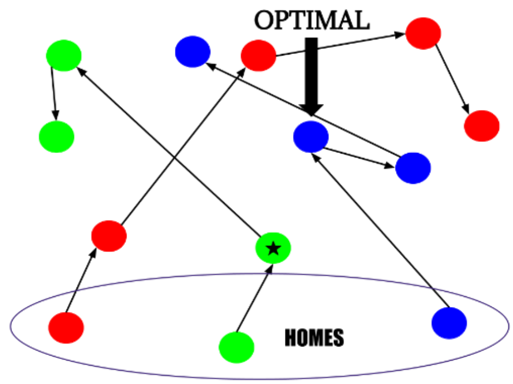

Figure 1. An example of MobileVaccClinic. The different colors represent different people and the circles represent the locations they visit (with the bottom three being their homes). Only considering homes in the problem formulation would result in the vaccination center being placed on the star. This would require people to deviate from their normal travels much more when getting a vaccine.

<u>Problem Statement:</u> We are given a set of locations C in a metric space characterized by the distance function d: C×C $\mapsto$ $R_{\geq 0}$. We additionally have a set of n clients P. Each individual $p \in$ P is associated with a set $Sp \subseteq$ C, which we can interpret as the set of locations p visits throughout the day. Finally, the input also includes a positive integer k constraining the number of facilities we can place, and a set S $\subseteq$ C containing the locations where we are allowed to place facilities. The goal of MobileVaccClinic is to choose a set $F \subseteq$ S with $|F| \leq k$ to place facilities, such that for every $p \in$ P we have d (, F) $\leq R$, for the minimum R possible. Here, we use the standard notation where (S, F) $=\min_{j\in S,'\in F}$ d(j, j'). Intuitively, this objective tries to minimize the maximum distance between the set of facilities placed and the locations visited by any client. We also consider three natural extensions:

• Outliers: To achieve herd immunity, we only need to vaccinate a large portion of the population (rather than every single person). To model this, we can take as input a parameter q, and seek to provide for only $\lfloor qn \rfloor$ of the clients. Formally, the new objective is to minimize R such that $|\{p \in \mathrm{P}: d(S_{p,F}) \leq R\}| \geq \lfloor qn \rfloor$.

• Fairness: Many studies have shown that COVID-19 disproportionately affects some demographic groups [19]. To counteract this, we seek to guarantee that different demographic groups have similar access to vaccines. As an example, when we solve the outliers formulation, we can guarantee that we are covering the same proportion of each demographic group when deciding the facility placements.

• Capacity: It is natural to assume that the number of vaccines that can be stored in each mobile facility is limited, say at most L. In this setting, we need to guarantee that every chosen facility will be assigned at most L people.

3.1. Hardness Result

For our hardness result, we use the following problem studied in [4], called γ-Colorful k-Center or γCkC for short. This problem is a generalization of the outliers version of k-center—in addition to the classical constraints, colors (representing demographic groups) are assigned to each client and the problem requires that a sufficient number of points of each color is covered. We use their work as a basis to prove a hardness result for our problem. Please refer to the proof in the complete paper [15] for more details.

Corollary 3.1. Even when the metric space is the Euclidean line, we have the following for MobileVaccClinic (unless P=NP):

(1) No approximation algorithm exists.

(2) Any bicriteria approximation algorithm must use at least $k \ln n$ facilities.

3.2. Algorithms

In this section, we introduce efficient methods which give (approximately) optimal facility placements, despite the hardness results. We also show how to extend each of our algorithms to ignore outliers, incorporate fairness constraints, and restrict the capacity of each facility. Please refer to the complete paper [15] for more details.

3.2.1. Fixed-Parameter Tractability

Let $U = \cup_{p \in \mathrm{P}} S_p$ denote the set of all the locations visited by the set of clients and $u = |U|$ be the number of locations in this set. Due to potential privacy concerns, we can assume that the client locations we have access to only include large public areas in the county such as malls, shopping centers, etc. Hence, it is reasonable to conclude that u is a fixed parameter, which we assume ranges from $15 - 30$. Given this fixed parameter, we develop an efficient algorithm for our problem.

Theorem 3.2. Algorithm 1 yields a 3-approximation algorithm for MobileVaccClinic and runs in time $2u$ poly(n, |C|).

Algorithm 1: FPT
1: for $A \in 2U$: $|A \cap Sp\ | \neq 0$, $\forall p \in$ P do
2: Obtain locations F_A by running the k-supplier algorithm on the appropriate instance discussed above.
3: Calculate the objective value for F_A.
4: end for
5: Pick the F_A with the smallest objective value.

Moving forward, we see that the same approach of guessing the correct set of client locations A will also apply in different settings. In fact, the only thing that may differ is the need for an alternative k-supplier algorithm that can incorporate the specific constraints of each unique setting; we survey some of these settings below.

Outliers: To modify our algorithm so that it only considers some fraction q of the population, we only need to change the objective value evaluated in line 3 of Algorithm 1. If we then feed A to the revised k-supplier algorithm, we will get a 3-approximation.

Fairness: Although our algorithm provides an upper bound guarantee for the maximum distance to a facility, the facility placement may significantly differ between individuals, with some having a facility right next to them, while others need to travel the whole $3R*$ guarantee. Luckily, the vaccine centers can vary from week to week. Thus, we can use a randomized algorithm such as the one given in [11], to guarantee that the reprovisioning of facilities over time will provide an improved per-point guarantee on expectation. Hence, we treat the clients stochastically fairly.

Capacity: In this case, we assume that each facility we use has a capacity L, i.e., at most L clients can be assigned to it in any solution. As we did for the regular case, we can create an instance of k-supplier where each location of Y for the k-supplier instance will have a capacity L. In other words, this will be an instance of capacitated k-supplier.

3.2.2. Covering Algorithm

In Corollary 3.1, we show that any bicriteria algorithm needs to open at least $k\ln(n)$ facilities in order to give a bounded approximation guarantee. Here, we show that this is essentially tight: we give an algorithm influenced by the standard Set Cover problem that outputs a set of locations of size at most ($\ln n + 1$), while guaranteeing that the objective value is at most that of an optimal solution.

Algorithm 2: ClientCover Search

1: Binary search on the sorted list $\{(i, j): j \in C, i \in S\}$, and let the current guess be R:
2: Use R to create the proper instance of ClientCover.
3: Obtain α-approximate solution FR for that instance.
4: If $|F_R| > \alpha \cdot k$, increase R; else, decrease R.
5: Output F_R for the minimum R such that $F_R \leq \alpha \cdot k$.

Theorem 3.3. Given an α-approximation algorithm for set cover, Algorithm 2 gives an $(1, \alpha)$-bicriteria algorithm for MobileVaccClinic.

As in the case of our FPT algorithm, we can easily extend Algorithm 2 to accommodate different settings. The only difference here lies in step 3, where instead of a classic Set Cover algorithm we can run a different algorithm.

Outliers: To modify our algorithm to only consider some fraction $q \in (0, 1)$ of the population, we can use some α-approximation algorithm for the Partial Set Cover problem, where the goal is to cover at least a q-fraction of the universe elements. Hence, we consider a variant of ClientCover, which we call Partial ClientCover, that requires only $\lfloor qn \rfloor$ points to be covered by balls of radius R^*.

Fairness: When solving MobileVaccClinic with outliers, the algorithm may view some demographic groups as outliers more often than others. To mitigate such possibilities, we can use an algorithm for the Partition Set Cover problem [14] to guarantee that a large proportion of each demographic group gets coverage. For example, we can guarantee that the algorithm considers a proportional number of people from each (demographic) group when choosing the vaccine center locations.

Capacity: As before, we assume that each facility we use has capacity L. We see that our general framework is still applicable: we can modify our algorithm to satisfy these capacity constraints by replacing the Set Cover algorithm with a Capacitated Set Cover algorithm when solving the ClientCover problem.

4. RESULTS

We run our experiments using synthetic data constructed from the 2019 U.S. population pipeline [7, 16] for Charlottesville City and Albemarle County in Virginia (see Table 1).

Table 1. Network Information

	Clients	Activity Locations	Residential Locations	Maximum Activity	Measured Diameter (km)
Charlottesville City	33156	5660	10038	9952	8.12
Albemarle County	74253	9619	32981	24506	61.62

We compare our algorithms with two heuristics: HomeCenters and MostActive. We set MostActive (opening vaccination centers at the k most visited locations) as the baseline because it is related to the current heuristic used by the Virginia Department of Health. In HomeCenters, we run k-supplier to place facilities at locations that minimize the maximum distance from client homes. We compare with this baseline to show the importance of considering mobility when placing the vaccination centers. For the complete experimental section, see [15] and our GitHub.

4.1. Tradeoff between Radius and Budget

It is important to evaluate the sensitivity of our algorithms to an increase in budget. We want to know how much the objective value would decrease if the county allocated more resources to deploy a greater number of mobile facilities. This knowledge can influence policy decisions: when an increase in budget yields a sharp decrease in objective, the government has more incentive to fund another vaccination center.

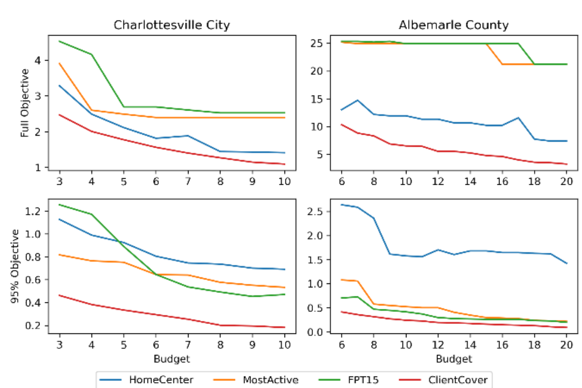

Figure 2. Tradeoff between vaccine accessibility and the number of vaccine centers placed.

As seen in Figure 2, there is generally a sharp decrease in the objective value when the budget is less than 6 for Charlottesville and 9 for Albemarle. As the budget increases past those thresholds, the marginal returns become so diminished that increasing the budget hardly changes the objective value. This is especially prominent in the full objective performance of FPT and MostActive. Hence, it is natural to recommend budgets of 6 and 9 clinics to the Charlottesville and Albemarle governments, respectively.

Additionally, we wish to bring attention to the overall poor performance of HomeCenters in these experiments. Though there is a general downward trend in the objective value for HomeCenters as the budget increases, there are cases in each county where increases in budget result in an increase in the objective value. This contradictory phenomenon is caused by the limited correlation between the distance to homes and our objective; as a result, noise/luck has a considerable effect. The noisiness of HomeCenters emphasizes the importance of modeling mobile populations.

5. FUTURE WORK

Since we have demonstrated the importance of modeling mobile populations, a natural next step is to extend other variants of the facility location problem to this setting as well. Furthermore, it would be valuable to see experimental results for the extended algorithmic versions for outliers, fairness, and capacity. Motivated by the rise of mobile vaccine distribution sites in rural Virginia, another line of work is to factor in not only the mobility of the population but also consider the mobility of these vaccine sites when targeting accessibility and equity.

6. CONCLUSION

In conclusion, we introduce a generalization of the classical k-supplier problem where we consider the mobility of populations when placing facilities. We show that designing an approximation algorithm for this variant is NP-Hard, so we turn to fixed-parameter tractability and bicriteria approximation algorithms to get around our hardness result. Finally, we experimentally show the efficacy of our algorithms in comparison to natural baselines.

ACKNOWLEDGMENTS

This paper is a summarized version of the work presented in [15]. The team thanks members of the Biocomplexity COVID-19 Response Team and the Network Systems Science and Advanced Computing (NSSAC) Division for their thoughtful comments and suggestions related to epidemic modeling and response support. George Li, Aravind Srinivasan, and Leonidas Tsepenekas were supported in part by NSF award number CCF-1918749. Ann Li, Madhav Marathe, and Anil Vullikanti were supported by DTRA (Contract HDTRA1-19-D-0007), University of Virginia Strategic Investment Fund award number SIF160, National Institutes of Health (NIH) Grants 1R01GM109718, 2R01GM109718, OAC-1916805 (CINES), CCF-1918656 (Expeditions), CNS-2028004 (RAPID), OAC-2027541 (RAPID), IIS-1908530, IIS-1955797, and IIS-2027848.

REFERENCES

[1] [n.d.]. “Meet the Four Kinds of People Holding Us Back From Full Vaccination”, New York Times. https://www.nytimes.com/interactive/2021/05/18/opinion/covid-19-vaccine-hesitancy.html.

[2] [n.d.]. See How Vaccinations Are Going in Your County and State. https://www.nytimes.com/interactive/2020/us/covid-19-vaccine-doses.html. Accessed:2022-01-23.

[3] Hamid Afshari and Qingjin Peng. 2014. Challenges and Solutions for Location of Healthcare Facilities. Industrial Engineering and Management 2014 (2014), 1–12.

[4] Georg Anegg, Haris Angelidakis, Adam Kurpisz, and Rico Zenklusen. 2020. A Technique for Obtaining True Approximations for k-Center with Covering Constraints. arXiv:2007.03946 [cs.DS]

[5] Derya Celik Turkoglu and Mujde Erol Genevois. 2020. A comparative survey of service facility location problems. Annals of Operations Research 292 (09 2020). https://doi.org/10.1007/s10479-019-03385-x

[6] Serina Chang, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky, and Jure Leskovec. 2021. Mobility network models of COVID-19 explain inequities and inform reopening. Nature 589 (2021). Issue 7840.

[7] Jiangzhuo Chen, Stefan Hoops, Achla Marathe, Henning Mortveit, Bryan Lewis, Srinivasan Venkatramanan, Arash Haddadan, Parantapa Bhattacharya, Abhijin Adiga, Anil Vullikanti, Aravind Srinivasan, Mandy L Wilson, Gal Ehrlich, Maier Fenster, Stephen Eubank, Christopher Barrett, and Madhav Marathe. 2021. Prioritizing allocation of COVID-19 vaccines based on social contacts increases vaccination effectiveness.

[8] Harwin de Vries, Joris van de Klundert, and Albert P.M. Wagelmans. 2020. The Roadside Healthcare Facility Location Problem A Managerial Network Design Challenge. Production and Operations Management 29, 5 (2020), 1165–1187.

[9] Zvi Drezner. 1995. Facility Location: A Survey of Applications and Methods.
[10] Teofilo F Gonzalez. 1985. Clustering to minimize the maximum intercluster distance. Theoretical computer science 38 (1985), 293–306.
[11] David G. Harris, Shi Li, Thomas Pensyl, Aravind Srinivasan, and Khoa Trinh. 2020. Approximation algorithms for stochastic clustering. arXiv:1809.02271 [cs.DS]
[12] Dorit S. Hochbaum and David B. Shmoys. 1985. A Best Possible Heuristic for the k-Center Problem. Math. Oper. Res. 10, 2 (May 1985).
[13] Dorit S. Hochbaum and David B. Shmoys. 1986. A Unified Approach to Approximation Algorithms for Bottleneck Problems. J. ACM 33, 3 (May 1986).
[14] Tanmay Inamdar and Kasturi Varadarajan. 2018. On the Partition Set Cover Problem. arXiv:1809.06506 [cs.DS]
[15] George Z. Li, Ann Li, Madhav Marathe, Aravind Srinivasan, Leonidas Tsepenekas, and Anil Vullikanti. 2022. Deploying Vaccine Distribution Sites for Improved Accessibility and Equity to Support Pandemic Response. In Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (Virtual Event, New Zealand) (AAMAS '22). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 789–797.
[16] Dustin Machi, Parantapa Bhattacharya, Stefan Hoops, Jiangzhuo Chen, Henning Mortveit, Srinivasan Venkatramanan, Bryan Lewis Mandy Wilson, Arindam Fadikar, Tom Maiden, Christopher L. Barrett, and Madhav V. Marathe. 2021. Scalable Epidemiological Workflows to Support COVID-19 Planning and Response.
[17] R. Meskarian, M. L. Penn, S. Williams, and T. A Monks. 2017. Facility location model for analysis of current and future demand for sexual health services. PLoS One (2017).
[18] Arman Nedjati and Mohammad Valipour. 2012. Solving Health Care Facility Location Problems with New Heuristic Algorithm Method. International journal of modeling and optimization (12 2012). https://doi.org/10.7763/IJMO.2013.V3.225
[19] Don Bambino Geno Tai, Aditya Shah, Chyke A. Doubeni, Irene G. Sia, and Mark L. Wieland. 2020. The Disproportionate Impact of COVID-19 on Racial and Ethnic Minorities in the United States. Infectious Diseases Society of America (2020).
[20] Lijing Wang, Xue Ben, Anniruddha Adiga, Adam Sadilek, Ashish Tendulkar, Srinivasan Ventatramanan, Anil Vullikanti, Milind Tambe, and Madhav Marathe. 2020. Using Mobility Data to Understand and Forecast COVID 19 Dynamics. medRxiv (2020).