A Robust Optimization Approach for Selecting Urban Fire Engine Company Locations

by Vivek Gopalan

Pennsylvania State University | email: Vivek.Gopalan21@gmail.com

Sponsor: Dr. Murali Kodialam[1]

Abstract

When a fire breaks out in a city, an emergency call to 911 is made and a fire engine company responds to the incident. This average response time is an important system metric that must be kept under a suitable threshold. The goal of this project is to study the relationship between the number and locations of fire engine companies and their response times to help cities optimize resources. Since the number and locations of fires are not known apriori, any selected set of engine company locations must be robust across a wide spectrum of fire incidents. Furthermore, if too much emphasis is placed on optimizing resources and system costs, some engine companies could bear a disproportionate fraction of the workload, leading to dissatisfaction or fatigue among firefighters. This project therefore also considers the incremental cost of ensuring equitable engine company workloads. A methodology combining integer programming techniques, ensemble learning, and local search using genetic algorithms is proposed to determine robust locations for the fire engine companies. Tested with a data set for the city of Philadelphia, the results support the hypothesis that optimizing fire engine company locations can result in significant savings for resource-strapped cities.

1. Introduction

On July 5, 2014, four children were killed in a 3-alarm fire that broke out in Philadelphia [14] and a similar fire entailed evacuating a hundred residents from an apartment building more recently [17]. When a fire or medical emergency occurs, a call is placed to 911 and even a minute reduction in the response time (i.e., the time between the call and arrival of Emergency Medical Services (EMS) personnel at the site) could save several lives [18]. A recent audit report published by the Office of the Controller for the City of Philadelphia, reports that the National Fire Prevention Association (NFPA) has set a standard where the response time for the fire engines should be within 5 minutes and 20 seconds after the call is dispatched for 90 percent of their runs. The Philadelphia Fire Department responds to approximately 54,000 calls a year [18]. While response times can be reduced by adding more ambulance shelters or fire engine company locations, cities are already budget-constrained and must do their best with very limited resources.

The fields of operations research/management science have contributed immensely to improving urban public services [8]. While EMS encompasses many domains, such as locating ambulance shelters [1, 5], and optimizing police patrols [12], this research will focus on studying the relationship between number and locations of fire engine companies and their response times to emergency incidents.

[1] Address: Nokia Bell Labs, 791 Holmdel Road, Holmdel, NJ 07733.
Email: Murali.Kodialam@nokia-bell-labs.com

Previous work has studied the relationship between travel distances and travel times for fire engines, and has also outlined an algorithm for engine company relocation(i.e., when an engine company is called away to respond to a fire, another engine company may be relocated to its base as a surrogate, in case a fire breaks out), but has ignored the development of robust solutions for fire engine company location problems [10, 11]. We focus on obtaining robust solutions for fire engine company location problems. A solution to a facility location problem is considered *robust* if it performs well under a wide range of changes in any stochastic parameters associated with the system. Baron et al. [2] consider robust facility location when demand for the product produced by the facility varies considerably over multiple time periods. Cui et al. [7] treat the case of facility disruption (e.g., a factory producing parts for a car unexpectedly goes out of commission) and develop robust solutions for this problem. Expanding on previous investigations, this paper finds optimal solutions to engine company location selection under *spatial uncertainty in demand* (i.e. it is not known precisely where fires will break out and engine companies must be located to accommodate a wide range of scenarios for fire locations).

This paper addresses the following research questions:

1. What is the minimum number of fire engine companies needed to provide satisfactory coverage of fires in a given city?
2. Given a required number of engine companies, where should these engine companies be located?
3. What resources should be made available at each fire engine company? In addition to the standard pumper-type truck, fire companies utilize a more functional and expensive ladder truck. How many *pumper* and *ladder* trucks should be placed at each location?

To address these research questions, we propose a **decision support system** (DSS) to determine engine company locations as outlined below and in Figure 1.

1. This paper proposes a new metric for measuring robustness in the context of fire engine company location. Robustness is defined in terms of a tuple (β, p), where at least β percent of fires that break out in a fixed time period are covered with threshold probability p (the probability of coverage p decreases with response distance (time) traveled). More specifically, the research answers the following question: what is the *minimum number of fire engine companies* needed within a city so that β % of fire incidents can be covered with a *threshold coverage probability* p?
2. Fire scenarios are generated via a Spatial Poisson process; the minimum number of engine companies and their locations are identified via the integer program **EC-PSCP** (*Equity Constrained Probabilistic Set Covering Problem*). This is Phase I in Figure 1.
3. The engine company locations identified for each scenario are combined via an **Ensemble Learning** algorithm, using the notion of "voting" for the most effective facilities (Phase II in Figure 1).

4. The current paper also considers ***resource allocation*** at each location. A variety of resource configurations may be deemed feasible at this stage, in the spirit of "*solution plurality*" [9]. Resources (e.g., pumper vs. ladder trucks) are allocated to each location by solving a **Constraint Satisfaction Problem** using **Genetic Algorithms**.

Section 2 of this paper describes the data used and provides details of the solution algorithm. *Section 3* presents a discussion of computational results and *Section 4* concludes the paper with relevant public-policy recommendations and offers future research directions.

2. Solution Approach

2.1.A. Method Overview

Figure 1 provides an overview of the procedure (algorithm) coded in Python, which is illustrated using data obtained for the city of Philadelphia.

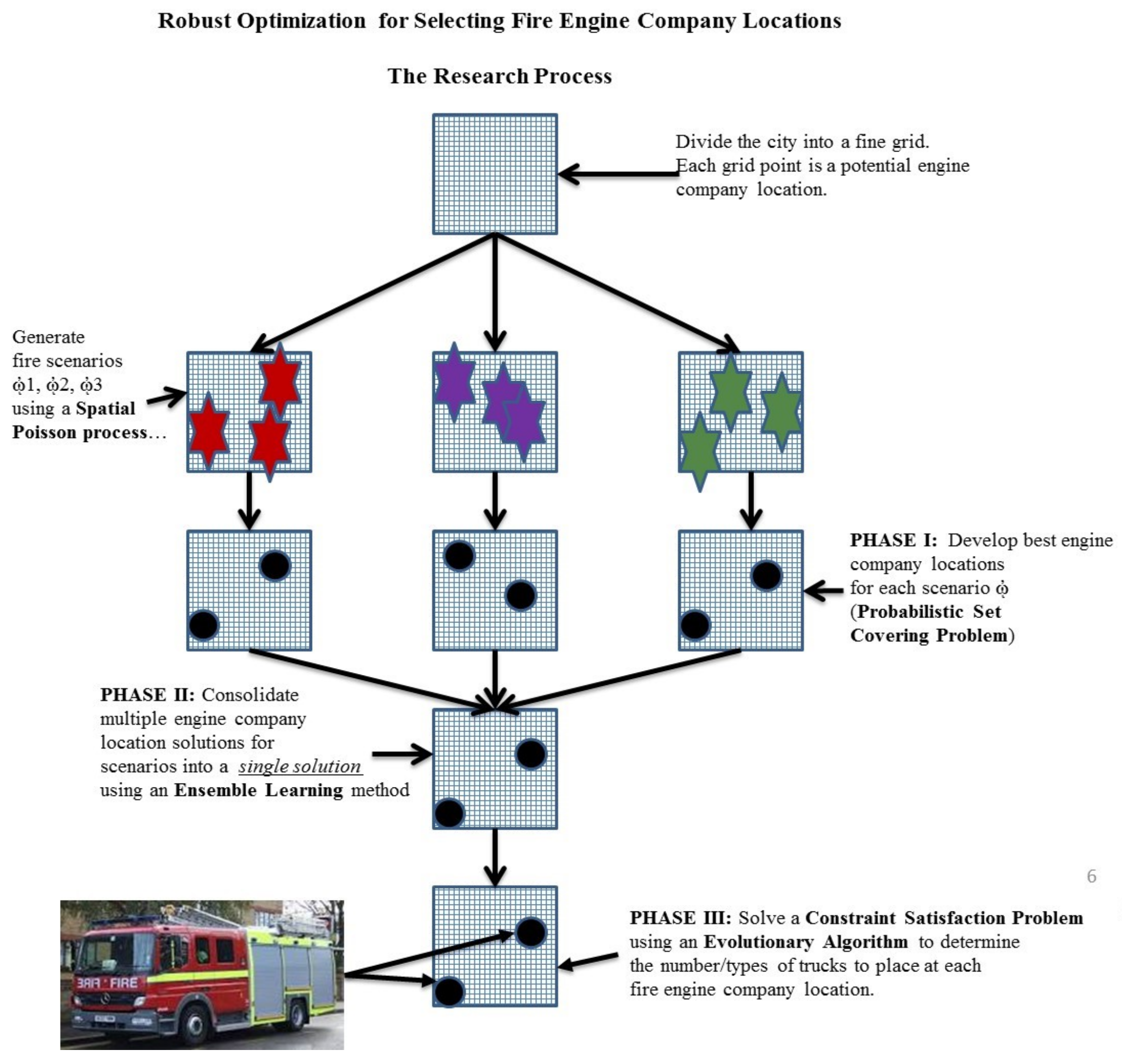

Figure 1: Robust Optimization Process for Selecting Fire Engine Company Locations

The algorithm (solution methodology) requires *user-provided inputs* and creates certain *outputs*:

1. **User provided inputs:** The user must provide a tuple (β, *p*), guiding the DSS to provide solutions where **β** % of fires can be covered with probability *p*. The coverage probability itself depends upon the travel speed of fire engines and the distance between the fire and the responding engine company. An expert user must provide a model for how the coverage probability varies with travel time (see Figure 2). In the computations presented in this paper, a linear decay in the coverage probability function is assumed. If a fire engine reaches a fire after *t* minutes, the coverage probability is stated below ($t_{\min}$ and $t_{\max}$ are user-specified parameters, see Figure 2):

$$p(t) = \begin{cases} 1, & t \le t_{\min} \\ -\frac{t - t_{\min}}{t_{\max} - t_{\min}} + 1, & t_{\min} < t \le t_{\max} \\ 0, & t > t_{\max} \end{cases}$$

Response times less than $t_{\min}$ are deemed satisfactory and response times exceeding $t_{\max}$ are unacceptable.

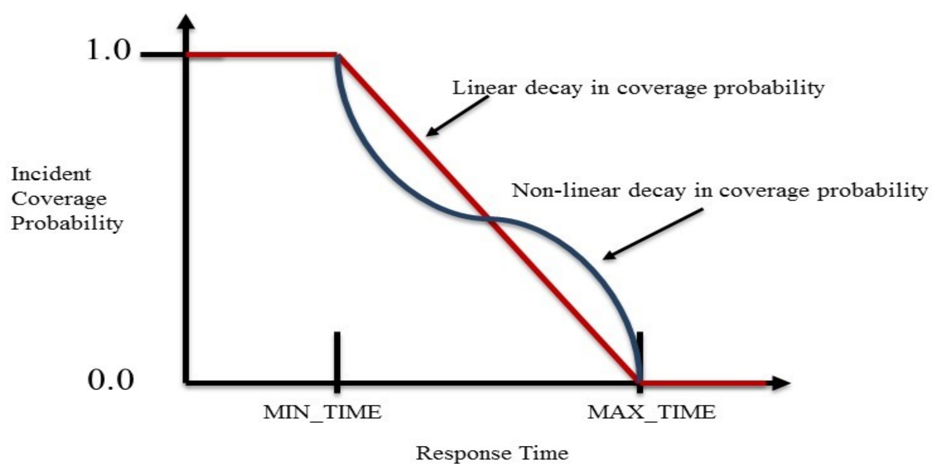

Figure 2: Illustration of Two Coverage Probability Functions

Outputs created by the solution methodology: *Number of engine companies* needed, their *arcGIS coordinates* and the *number and types of fire engines* (pumper or ladder trucks) at each location.

This **solution method**, discussed below, consists of the five parts (refer to Figure 1).

2.1.B. Data

As a case study to illustrate the algorithm, we consider the city of Philadelphia. The city is home to about 1.5 million residents and can be geographically partitioned into 155 neighborhoods (Figure 3). For each of these neighborhoods, a public-domain arcGIS data set was obtained (see Figure 3) identifying the *neighborhood population* and the *neighborhood centers* (i.e., the x and y coordinates for the centers in the arcGIS frame of reference). The precise spatial locations of fire incidents (for a 1-year period) were not available for this project, but from public-domain information, it is known that a total of 54485 fire incidents occurred during 2015 [18].

2.2 Methods

2.2.1 Generating Candidate Engine Company Locations

The city of Philadelphia consists of 155 neighborhoods (as shown in Figure 3 below) and for each neighborhood, a square grid is drawn, with the center of the square aligning with the center of the neighborhood, thereby masking the whole city (see top box in Figure 1). Each corner point of the grid is assumed to be a potential engine company location and the mesh size for the grid is a parameter that can be controlled by the user. The finer the mesh size, the closer the location problem is to a "continuous location" problem.

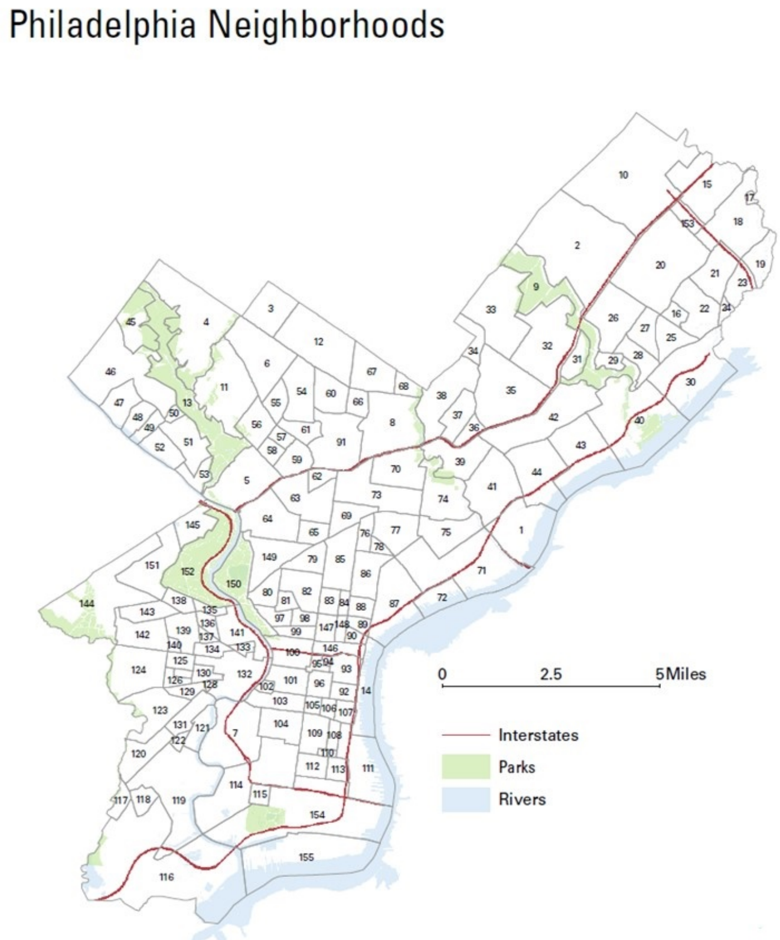

Figure 3: Map of Philadelphia Neighborhoods [16]

2.2.2 Generating Fire Location Scenarios

A scenario is simply a set of locations for fires that occur during a fixed period. If a data set were available with the exact spatial locations of fires for a year, the continuous spatial distribution to approximate the discrete set of real world fire events (the observed data) can be fit using Kernel Density Estimation (KDE), which is a non-parametric approach to density estimation from a set of discrete points [6, 13, 4]. The smoothing parameter (bandwidth) determines the extent of spread of each point. Once the spatial distribution is known, fire scenarios can be generated by sampling from this spatial distribution. However, fitting a spatial distribution for fire incidents is not the main objective of this paper. Moreover, fire incident location data were unavailable for this project, although an aggregate number for the total number of fires in the city of Philadelphia was available. The total number of fires was allocated to each of the 155 neighborhoods in proportion to the neighborhood's population. Fires were then randomly generated for each neighborhood. One complete fire scenario consists of a set of spatial locations for fires in each of the 155 neighborhoods. The scenario generation method repeatedly creates fire scenarios for the city of Philadelphia. The next step is to solve for the best engine company locations for each scenario using the method outlined in the next sub-section (2.3).

2.3 Solving the Equity Constrained Probabilistic Set Covering Problem (EC-PSCP) for each Scenario

The Equity Constrained Probabilistic Set Covering Problem (EC-PSCP) represents the core of the methodology presented in this paper. An EC-PCSP is solved once for each generated fire scenario (and the solutions later aggregated via an Ensemble Algorithm (see section 2.4)). The EC-PCSP is an integer program that is a variant of the classical set covering problem. As a preamble to solving the EC-PCSP, a set covering matrix $\boldsymbol{A}$ must be developed. The matrix $\boldsymbol{A}$ has one row for each fire incident and one column for each potential engine company location. A matrix entry a_{ij} is 1 if an engine company at location j can cover a fire at i. The following steps enable the computation of matrix $\boldsymbol{A}$.

Step 1: The Euclidean distance between each *candidate facility* j (node of the grid) and each fire incident location i in the scenario is computed.

Step 2: The Euclidean distance is converted to a travel distance using a **DETOUR INDEX** (equal to the ratio of actual travel to Euclidean distance) of 1.42. Prior research supports the use of such an index [9, 3].

Step 3: Using the user-specified travel speed for fire engines, the travel time matrix $\boldsymbol{T}$ is computed. $\boldsymbol{T}$ has (i, j)th entry = $t(i,j)$ = (travel distance between i and j)/(engine speed).

Step 4: Using the travel time, compute the *coverage probability* p (as in Figure 3) and determine if a facility at j can cover a fire at i. Develop a *set covering matrix* $\boldsymbol{A}$, with (i,j)th entry $a_{ij} = 1$ if the coverage probability $\geq$ user-specified threshold probability and 0 otherwise [Note: The user-specified threshold probability is the second field in the tuple (β, p)].

Step 5: Use the set covering matrix to develop a mathematical programming formulation for the fire engine company location problem. The set covering formulation is solved optimally by using the Gurobi optimization solver [15].

An example set covering matrix and the formulation for the Equity Constrained Probabilistic Set Covering Problem **EC-PSCP** is explained below.

In this example scenario, five fires must be covered by a set of four engine companies. The set covering matrix for this scenario is presented in Table 1.

Table 1: The set covering matrix for example scenario, where columns denote EC = Engine Company, and rows are fire incidents

	(EC 1)	(EC 2)	(EC 3)	(EC 4)
fire 1	1	1	1	0
fire 2	1	1	1	0
fire 3	1	1	1	1
fire 4	0	0	1	0
fire 5	0	0	1	1

The set covering matrix $\boldsymbol{A}$ has entry $a_{ij} = 1$ if an engine company at location (column) j can reach a fire at location (row) i with required coverage probability p. The constraint set is illustrated below for this example.

The Equity Constrained Probabilistic Set Covering Problem (EC-PSCP):

1. The mathematical program has a **variable** $\boldsymbol{y_j}$ for each node (*potential* EC facility).
2. $\boldsymbol{y_j}$ is 1 if an engine company is located at node j, and 0 otherwise (a switch variable).
3. The math program has a binary variable X_{ij} which is set equal to 1 if fire i is covered by an engine company located at j.
4. The math program has a switch variable S_i for each fire i. If $S_i = 1$ in the solution, fire i is NOT covered within the stipulated response time. The variable S_i provides a "pass" that allows the math program to pass on covering fire i (In this model,only β % of fires are covered within the stipulated response time).
5. A *scenario* is defined by a set of locations for fire incidents, drawn from a spatial distribution.
6. The mathematical program has five sets of constraints (explained below) in the *scenario* ω.
7. Every new *scenario* ω will lead to a new mathematical program because the locations of the fires will change.

OBJECTIVE: Minimize $\sum_{j=1}^{|J|} y_j$

Illustration for Example: OBJECTIVE: Minimize $\sum_{j=1}^{4} y_j = y_1 + y_2 + y_3 + y_4$

SUBJECT TO FIVE SETS OF CONSTRAINTS:

Constraint set I (cannot set X_{ij} to 1 unless y_j is also 1):

$$\sum_i a_{ij} X_{ij} \leq |I| y_j \text{ for all engine companies } j \in J$$

Illustration of Constraint Set I for Example:

$|I| = 5$, since there are 5 fires in the example matrix.

$$X_{11} + X_{21} + X_{31} \leq 5y_1$$
$$X_{12} + X_{22} + X_{32} \leq 5y_2$$
$$X_{13} + X_{23} + X_{33} + X_{43} + X_{53} \leq 5y_3$$
$$X_{34} + X_{54} \leq 5y_4$$

Constraint set II (either a fire i must be covered by an engine company y_j, OR the "pass" variable S_i must be set equal to 1):

$$\sum_j a_{ij} y_j \geq [1 - S_i] \text{ for all fires } i \in I$$

$$y_1 + y_2 + y_3 \geq [1 - S_1] \ (i.e., cover\ fire\ 1\ if\ S_1 = 0)$$
$$y_1 + y_2 + y_3 \geq [1 - S_2] \ (i.e., cover\ fire\ 2\ if\ S_2 = 0)$$
$$y_1 + y_2 + y_3 + y_4 \geq [1 - S_3] \ (i.e., cover\ fire\ 3\ if\ S_3 = 0)$$
$$y_3 \geq [1 - S_4] (i.e., cover\ fire\ 4\ if\ S_4 = 0)$$
$$y_3 + y_4 \geq [1 - S_5] \ (i.e., cover\ fire\ 5\ if\ S_5 = 0)$$

Constraint set III: Enforce the Equity Spread Constraint

$$\sum_i a_{ij} X_{ij} \leq Max \text{ for all engine companies } j$$

$$\sum_i a_{ij} X_{ij} \geq Min \text{ for all engine companies } j$$

$$Max - Min \leq EquitySpread\ [this\ is\ a\ parameter\ provided\ by\ user]$$
$$Illustration\ of\ Equity\ Spread\ constraints\ for\ example:$$
$$X_{11} + X_{21} + X_{31} \leq Max$$

$$X_{12} + X_{22} + X_{32} \leq Max$$

$$X_{13} + X_{23} + X_{33} + X_{43} + X_{53} \leq Max$$

$$X_{34} + X_{54} \leq Max$$

$$X_{11} + X_{21} + X_{31} \geq Min$$

$$X_{12} + X_{22} + X_{32} \geq Min$$

$$X_{13} + X_{23} + X_{33} + X_{43} + X_{53} \geq Min$$

$$X_{34} + X_{54} \geq Min$$

$$\boldsymbol{Max - Min \leq EquitySpread}$$

Constraint set IV: At least β % of fires must be covered.

$$\boldsymbol{\sum_{i=1}^{|I|} S_i \leq [1-\beta] \times |I|}$$

$\boldsymbol{\{equivalently,\ \sum_{i=1}^{|I|} \sum_{j=1}^{|J|} X_{ij} \geq \beta \times |I|\}}$

Illustration of constraint set IV for example, where β = 0.8 is assumed.

$\boldsymbol{S_1 + S_2 + S_3 + S_4 + S_5 \leq \{(1-0.8) \times 5\} = 1}$ **(This ensures 4 out of 5 fires covered)**

Constraint Set V:

Declare all variables to be BINARY (0 or 1)

$$\boldsymbol{X_{ij}, y_j, S_i \in \{0, 1\}}$$

In the example:

y_1, y_2, y_3, y_4 are all binary (0 or 1) variables

$S_1..S_5$ are ALSO binary (0 or 1) variables

$\boldsymbol{X_{ij}}$ is binary for $\boldsymbol{i = 1, \ldots, 5}$ and $\boldsymbol{j = 1,..4}$

2.4 Combining Multiple Solutions with an Ensemble Algorithm

The Ensemble Algorithm combines the location solutions for each generated scenario into a single location solution. Each candidate facility j (in *any* of the scenario solutions) is given a cumulative weight that depends upon the distance of facility j to either a) all fire scenarios *or* b) to all other selected engine company locations. Finally, all the engine company locations in *any* scenario solution are rank-ordered based upon this weight. Engine companies are successively selected from this rank-ordered list in a greedy fashion (i.e., the next best facility is the one that covers the most number of uncovered fires) until the robustness criterion of β % of fires covered is achieved.

2.5 Solving a Constraint Satisfaction Problem using a Genetic Algorithm

When the Ensemble Algorithm is completed, the number of engine companies selected is known along with their precise arcGIS locations. The purpose of Phase III (in Figure 1) is to place *resources* at these engine companies. Two kinds of resources are considered, pumper engines (that cost about $300,000) and ladder engines (that are more expensive, upwards of $900,000). The goal of the Genetic Algorithm is to determine the *type* and *number* of engines at each location. A Genetic Algorithm approach was selected as the search procedure as it offered a natural way to encode the solution space. At least one engine must be placed at each location. In addition, two kinds of constraints must be satisfied by a solution to the Constraint Satisfaction Problem: a) a constraint on the overall cost of engines for the entire city (a budget constraint) and b) every fire must be responded to within 8 minutes by some type of engine (pumper or ladder) and every fire must be responded to with a ladder truck within 18 minutes (these parameters are used for illustrative purposes only and can be changed by the user of the DSS). The *fitness function* for the GA is a two-component vector; one component is total cost, and the other is percentage of fires not covered as per constraint b above. A simple algorithm combining crossover and mutation of genes (where a gene is a resource profile for *all* the selected EC locations) has been designed for this phase. Figures 4 and 5 below provide schematic representations of the Genetic Algorithm steps.

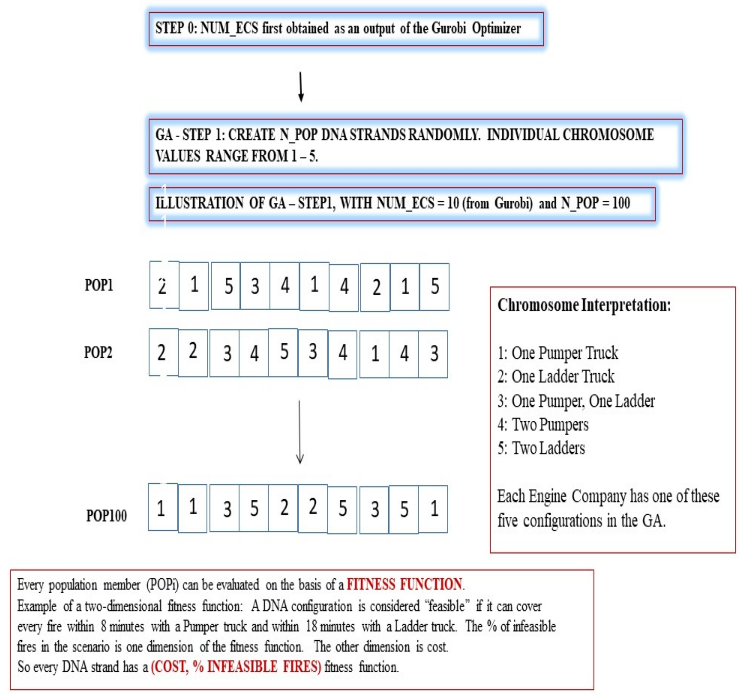

Figure 4: Schematic of Genetic Algorithm.

Figure 4 provides a schematic representation of the Genetic Algorithm step. The Gurobi Optimizer [15] is used to solve the probabilistic set covering problem in Phase I. The optimizer determines the location of the Engine Companies. The Genetic Algorithm determines how many pumper and ladder trucks to place in each location.

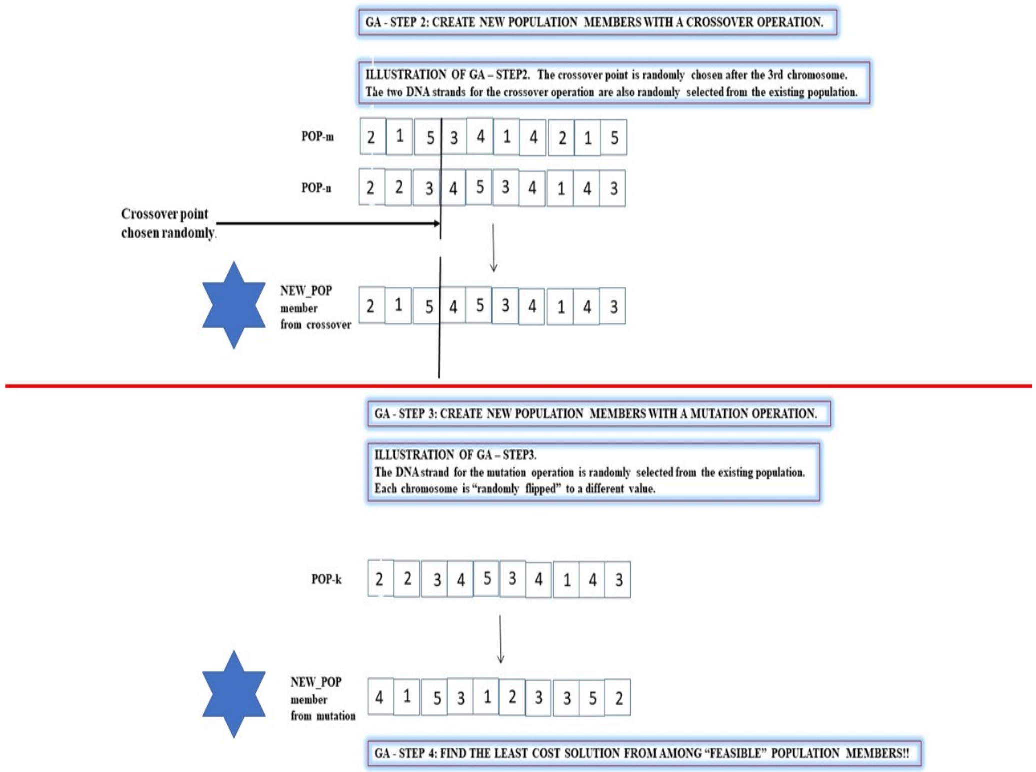

Figure 5: Schematic for Cross Over and Mutation Operations using Genetic Algorithms

Figure 5 provides a schematic representation for cross over and mutation operations in the Genetic Algorithm. There were 100 cross over operations followed by 100 mutation operations for the computational results presented in this paper. However, the user has the ability to set parameters for the number of crossover and mutation operations performed. While in principle, there could be genes that code for a larger number of resources at each Engine Company, in practice, budget limitations would make gene configurations other than the ones discussed in this paper unlikely.

3. Discussion of Computational Results

The methodology presented in the previous section was coded in the Python programming language. The Equity Constrained Probabilistic Set Covering Problem (EC-PSCP) was solved using the Gurobi optimization package [15]. All computations were performed on a 64-bit Lenovo laptop. In conformance with standard practice in the machine learning literature, five training scenarios were created to find engine company solutions and these solutions were validated using five test scenarios (results reported in this section are generally averages for the five test scenarios). Please note, the user can also select the number of training and test scenarios. The discussion of computational results is further organized into the following sub-sections:

1. Impact of the robustness parameter β
2. Impact of engine speed
3. Solution quality in terms of engine company equity
4. Impact of coverage probability function assumptions
5. Illustrative performance of the genetic algorithm.

3.1 Impact of the Robustness Parameter β

The parameter β, which is the proportion of fires covered effectively in a solution is the fundamental measure of "robustness" of any engine company solution. The higher the β, the more engine companies will be required. The following charts demonstrate the impact of β on the number of engine companies required and the maximum response time.

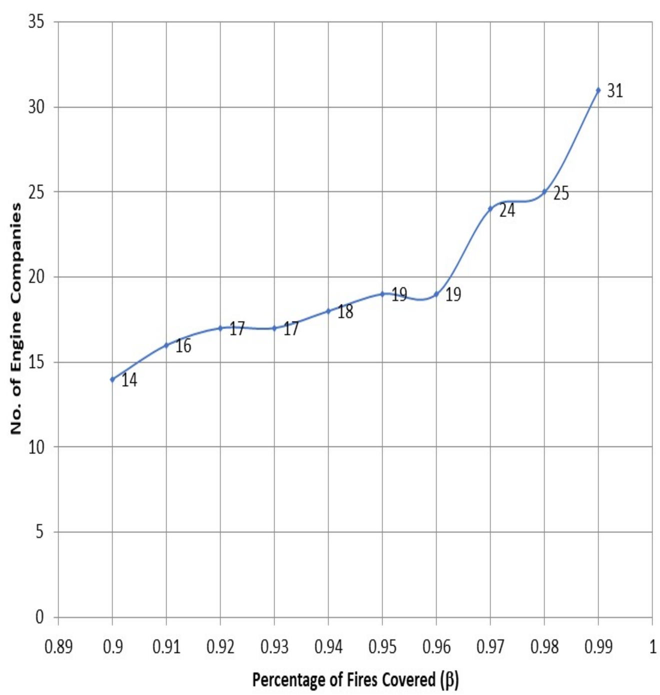

Chart 1: Impact of β on the Number of Engine Companies Required

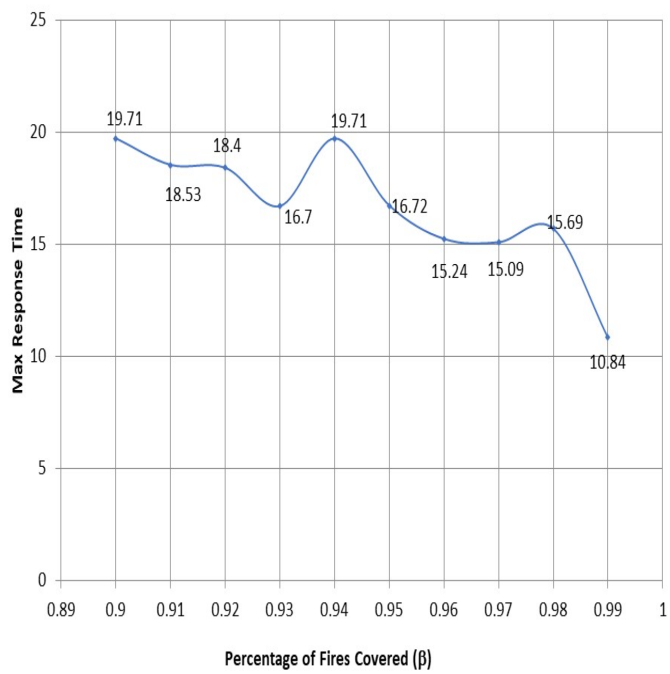

Chart 2: Impact of β on Maximum Response Times

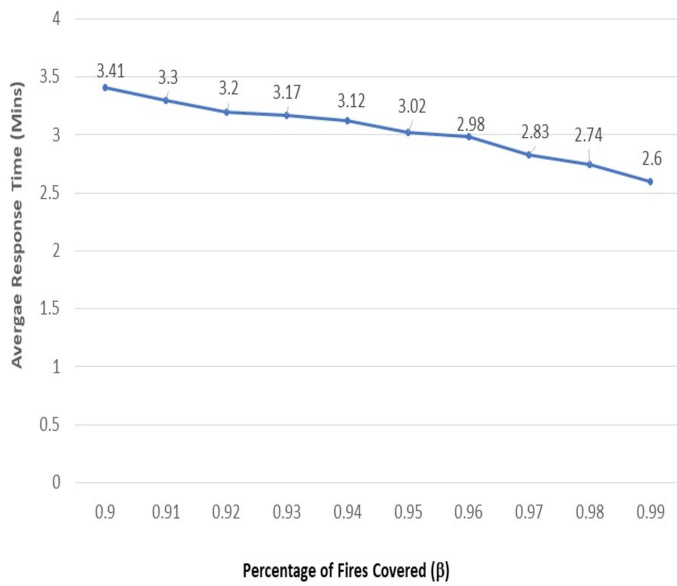

Chart 3: Impact of β on Average Response Times

As β increases from 0.9 to 0.99, the average response time decreased from 3.41 to 2.6 minutes. However Charts 1, 2, and 3 together document the price of robustness (# engine companies required more than doubles as β increases from 0.9 to 0.99) and the benefit of robustness (the maximum response time for *any* fire almost drops by 10 minutes). Policy makers must thoughtfully choose an operating point from these results.

3.2 Impact of Fire Engine Travel Speed

The speed of travel for any incident may depend upon extraneous factors such as time of day and traffic conditions. However, most cities have an inherently latent capacity for allowing fire trucks to move about the city with a certain speed. Chart 4, below, shows a steep increase in the number of engine companies is required for travel speeds less than 30 mph. In this research, rather than view the fire engine travel speed as a given of the environment, it is assumed that policy makers have a limited ability to influence the speed of travel for emergency response (e.g., by levying extreme fines for obstructors).

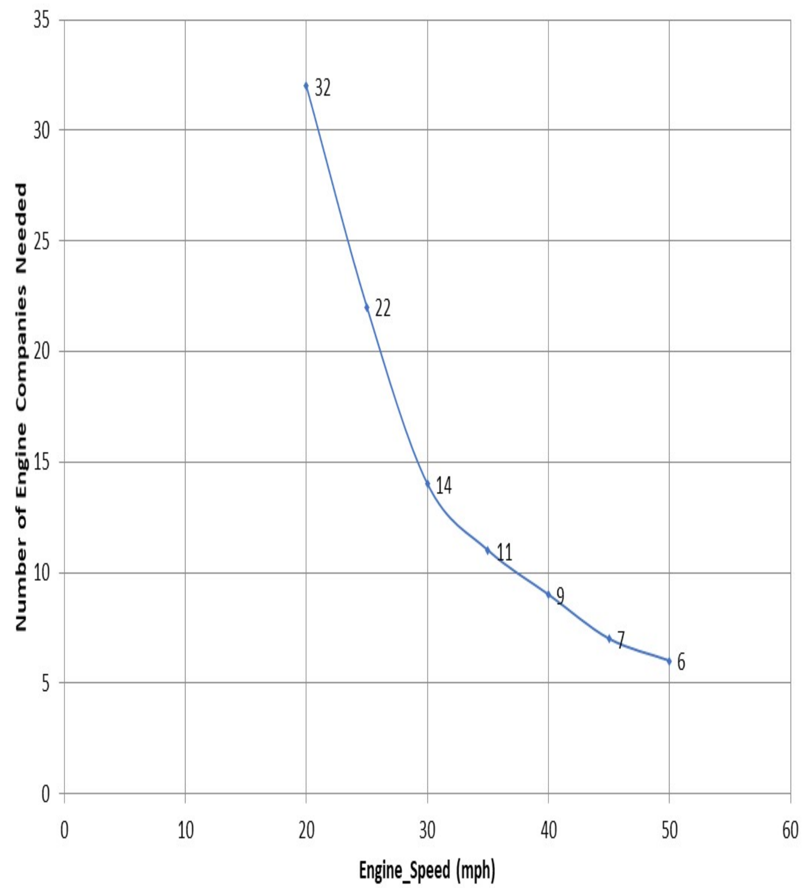

Chart 4: Impact of Engine Speed on Number of Engine Companies Required

3.3 Equity for Engine Company Workloads

A good solution for the engine company location problem must pay attention to the workloads allocated to various fire engine companies (some companies cannot remain idle too much of the time). Equity considerations can be explicitly added to the EC-PSCP, but this also makes the problem more difficult to solve. Future research should also explore improving the equity with Phase III (local search for better solutions). Chart 5 below indicates a discrepancy of about 10% (between the busiest and most idle engine companies) when EC-PSCP is solved *without* any constraints on equitable workloads. Moreover, equity is harder to achieve for higher values of β.

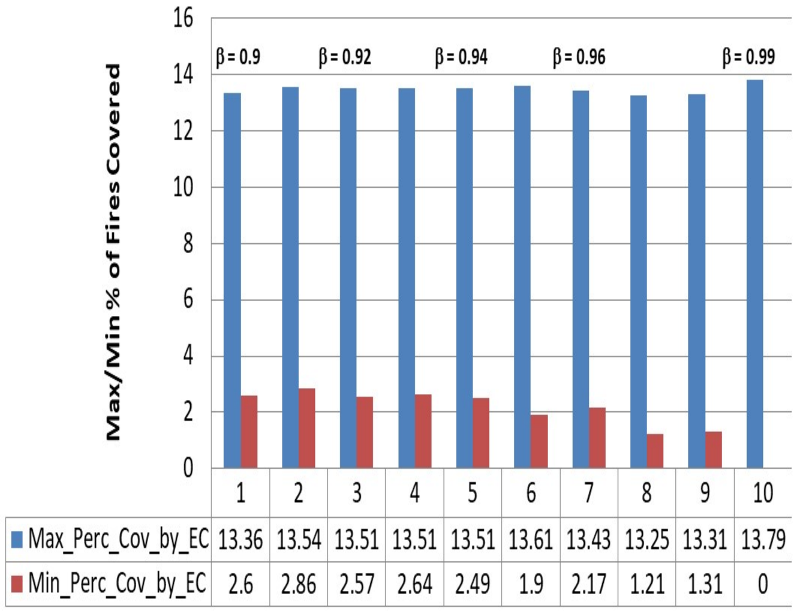

Chart 5: Impact of Equity Spread on Engine Company Workloads as β Increases

However, EC-PSCP can also be solved parametrically by varying the "Equity Spread" limit in the formulation. Clearly, the number of engine companies needed, as well as the total system cost, can be expected to increase as the Equity Spread limit decreases. The behavior of the system with decreasing Equity Spread is characterized in Charts 6 and 7 below.

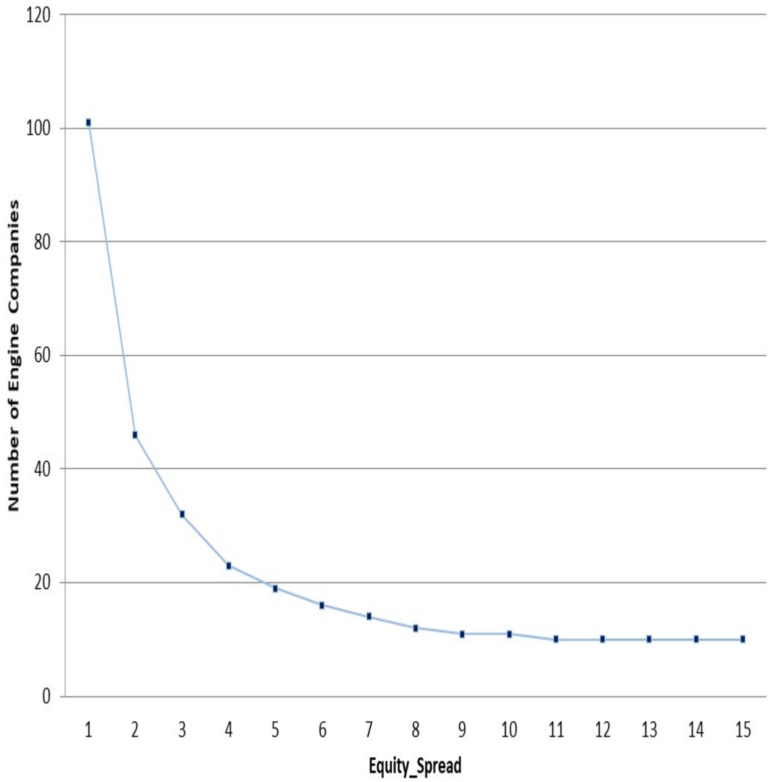

Chart 6: Impact of Equity Spread on Number of Engine Companies Required

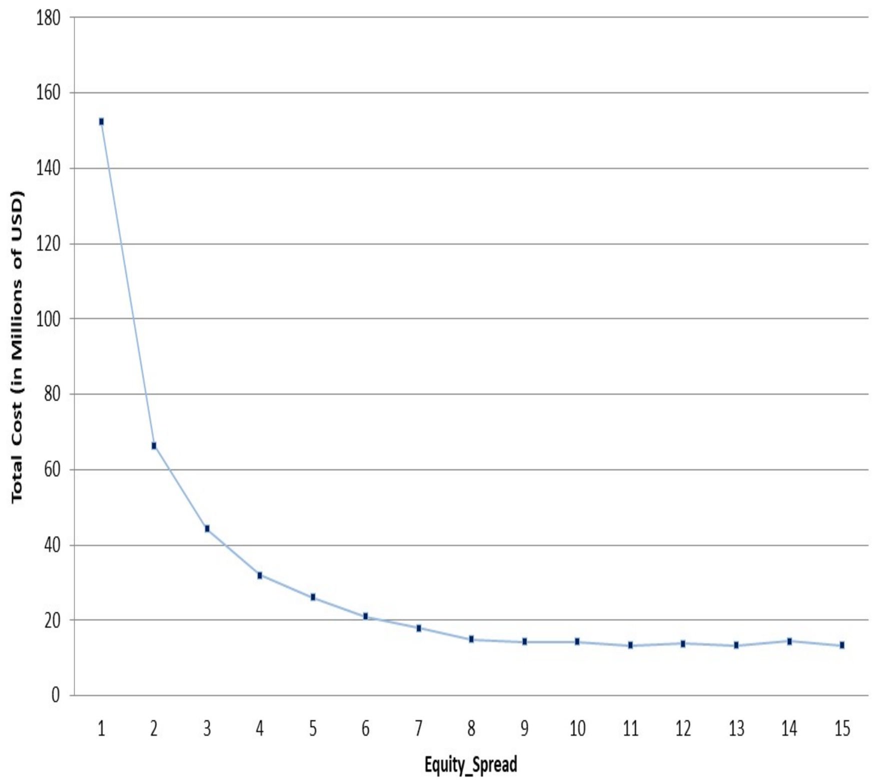

Chart 7: Impact of Equity Spread on Total Cost

For low levels of Equity Spreads, the system cost increases by as much as $1M for each percentage point reduction in the Equity Spread. From these plots, it can be observed that the increase in system cost depends upon the current level of the equity spread. So a decrease in Equity Spread from 15 to 14 has a much smaller impact on system cost as compared to a decrease from 2 to 1.

3.4 Impact of Coverage Probability Function Assumptions

The user-specified coverage probability function of Chart 1 is a key determinant of the solution quality. In particular, t_{min} and t_{max} in Chart 1 influence the form of the coverage probability function and in turn the number of engine companies needed. In some cases, the coverage probability function may be the subjective opinion of an expert user of the DSS. For this reason, the impact of changing the form of the coverage probability function is studied below in Chart 8. The National Fire Protection Association (NFPA) standard calls for the first due fire engines to arrive on scene within 5 minutes and 20 seconds after being dispatched for 90 percent of their runs. As expected, **response times lower than about 6 minutes impose an enormous cost on the system.** [18]

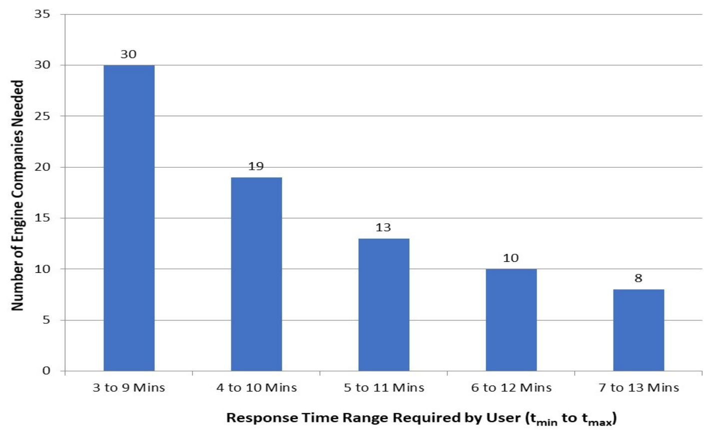

Chart 8: Impact of Response Time Standards on Number of Engine Companies

3.5 Illustrative Performance of the Genetic Algorithm (GA)

Phase III of the solution procedure employs a GA to place resources (pumper and ladder trucks) at each chosen engine company location. Only resource allocations are changed in this phase and engine company locations themselves are fixed. The GA uses a two-dimensional fitness function. The first dimension is the *budget* used. Given that there must be at least one engine at each location, a lower bound for the cost is the number of engine companies multiplied by the cost of a pumper (less expensive) truck. All resource configuration costs are expressed as an index with respect to this base cost (i.e., a cost of 150 means the resource configuration is 50% more expensive than the cost of placing one pumper truck at all stations). Likewise, the second component of the fitness function represents the *percentage of fires* that do not satisfy the constraint "first response (by any type of truck) within 8 minutes and a ladder truck available within 18 minutes". The GA searched for a resource configuration with a cost index less than 200 and the percentage of infeasible fires less than 10% (For illustration only: the user can tune feasibility parameters for genes). Chart 9 below illustrates how these two GA fitness attributes trade off in the set of feasible solutions found. The computational experiments indicate that solutions feasible to these two constraints are hard to find (less than 10% of the population members generated by GA were "feasible").

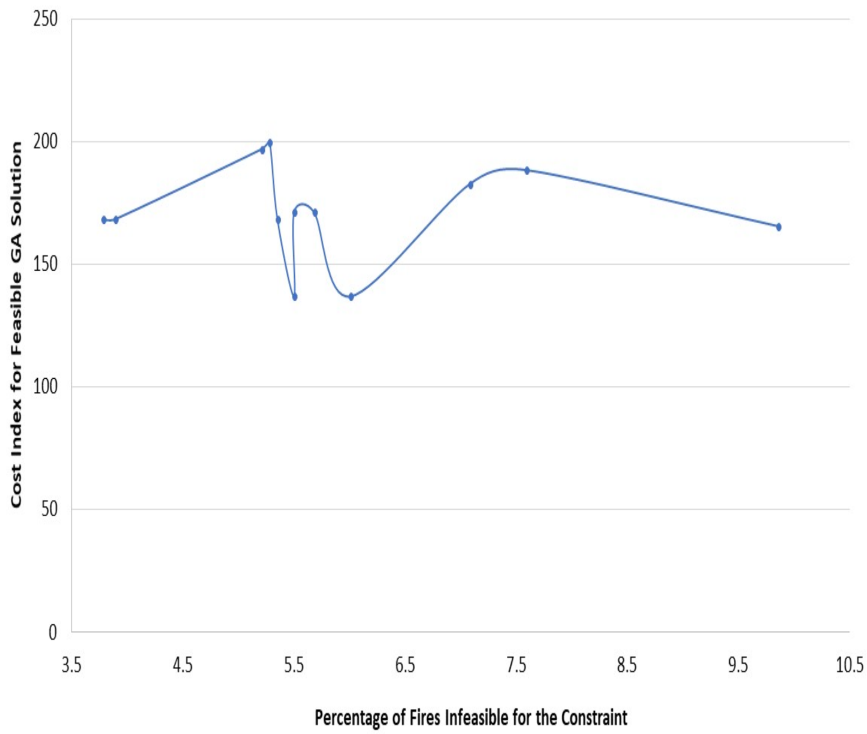

Chart 9: GA Performance: System Cost vs. Percentage of Fires Not Covered

Finally, a check was made for overfitting. If the β calculated for the test data sets is significantly lower than the β stipulated for the training data sets, overfitting has occurred. It was validated that the model was not overfitting the data. There were some "negative" results in the computations however. For instance, in the Ensemble Algorithm, it did not matter whether facilities were combined based upon distance to fires in scenarios, or distance to other facilities chosen in other solutions. Moreover, the number of candidate engine company locations that the DSS started with did not matter beyond a point. Most of the reported computations started with 625 candidate engine company locations.

4. Conclusions, Public Policy Implications, and Future Research Opportunities

4.1 Conclusions

This paper develops a robust optimization approach for locating fire engine companies. The main dimension of robustness addressed is the spatial uncertainty of fire incident locations. For this problem, the current paper provides a solution algorithm to find the minimum number of engine companies needed so that β % of fires can be covered with probability p. The main contribution is the development of a Probabilistic Set Covering Model formulation for this problem. In addition to the EC-PSCP, the paper also develops a method to combine solutions from different scenarios (the Ensemble algorithm) and a local search procedure for placing resources at chosen fire engine company locations.

4.2 Public Policy Implications

The solution algorithm and associated computational results raise several important public policy issues:

1. One of the responsibilities of city government is to choose an appropriate level of coverage β. **In fact, the required number of engine companies more than doubles as β increases from 90 to 99%.** Policy makers should carefully evaluate the trade-off between the increased cost of opening additional locations and the real benefits from covering the last few percentiles.
2. As engine speed increases, the number of engine companies required decreases (as expected). The data also seem to indicate that EMS responders must try to achieve an engine speed of at least 30 mph. **For engine speeds < 30 mph, there is again a steep increase in the number of engine companies required.** Given that this parameter has the greatest impact on the number of engine companies needed, policy makers should consider policy options that can increase fire engine travel speeds. Options such as special lanes for fire engine travel (like what is done for mass transit buses already) and higher fines on the roads for obstructing fire engine travel must be considered.
3. Special attention must be given to **engine company workload equity**. The computational results indicate that there might be as much as a 10% difference in the proportion of fires tackled by the busiest engine company and the engine company with the smallest workload. Territory design for engine companies (to attain equitable workloads) is an important extension of this project, as disparity may cause difficulties with labor unions or contract workers.
4. Finally, achieving response times of less than 5 minutes is extraordinarily difficult. **There is a steep increase in the number of engine companies required for achieving average response times less than 5 minutes.** Policy makers should consider options wherein a sufficient amount of education and equipment is provided locally (e.g., fire extinguishers) at building sites, so that local residents can contain the impact of the fire for about 6-8 minutes. If this 6-8-minute time threshold can be managed locally, the city can also drastically lower the costs of opening more fire engine companies.

4.3 Future Research Opportunities

The current research also has some modeling limitations. Travel times are computed ignoring conditions like traffic or time of day. Some fires require multiple response units and it may not be sufficient to dispatch just the closest unit. Some fires may also require specific equipment such as ladder companies that may not be available at the closest facility (this issue is partly addressed by the Genetic Algorithm). Fire engines may also be unavailable due to external circumstances such as maintenance and maintenance plans must be factored into developing engine company locations. In terms of modeling enhancements, this exercise can be repeated with GIS mapping tools to formulate the same model at a more fine-grained level (e.g., include details of one-way streets, traffic lights/intersections). Finally, the model developed herein has a rich set of other applications such as ambulance shelter location, police patrol improvement and logistics for emergency response (e.g., for hurricanes such as Katrina), where spatial demand uncertainty considerations are required.

5. References

[1] R. ALANIS, A. INGOLFSSON, and B. KOLFAL, A Markov Chain Model for an EMS System with Repositioning, *Production and Operations Management*, Vol. 22, No. 1, (**2013**), pp. 216–231.

[2] O. BARON, J. MILNER, and H. NASERALDIN, Facility location: a robust optimization approach, *Production and Operations Management*, Vol. 20, Issue 5, (**2011**), pp. 772-785.

[3] F.P. BOSCOE, K.A. HENRY, and M.S. ZDEB, A Nationwide Comparison of Driving Distance Versus Straight-Line Distance to Hospitals, *Prof Geogr.* , 64(2), Apr 1 (**2012**).

[4] F.P. BOSCOE, K.A. HENRY, and M.S. ZDEB, A Nationwide Comparison of Driving Distance Versus Straight-Line Distance to Hospitals, *Prof Geogr.* , 64(2), Apr 1 (**2012**).

[5] S. BUDGE, A. INGOLFSSON, and D. ZEROM, Empirical Analysis of Ambulance Travel Times: The Case of Calgary Emergency Medical Services, *Management Science*, Vol. 56, Issue 4, (**2010**) pp. 716 – 723.

[6] T.C.Y. CHAN, D. DEMIRTAS, and R. KWON, Optimizing the Deployment of Public Access Defibrillators, published as Articles in Advance, *Management Science*, January (**2016**).

[7] T.CUI, Y.OUYANG, and Z.M. SHEN, Reliable facility location design under the risk of disruptions, Operations Research, Vol. 58, Issue 4, (**2010**).

[8] L.V. GREEN and P.J. KOLESAR, Improving emergency response with management science, *Management Science*, Vol. 50, No. 8, (**2004**), pp. 1001 – 1014.

[9] S.O. KIMBROUGH, A. KUO, L.H. CHUIN, F.H. MURPHY, and D.H. WOOD, Solution pluralism and metaheuristics, The *IX Metaheuristics International Conference*, Udine, Italy, July 25-28, (**2011**).

[10] P. KOLESAR and E.H. BLUM, Square root laws for fire engine response distances, *Management Science*, 19(12), (**1973**), pp. 1368-1378.

[11] P. KOLESAR and W.E. WALKER, An algorithm for the dynamic relocation of fire companies, *Operations Research*, 22(2), (**1974**), pp. 249-274.

[12] R.C. LARSON and T.F. RICH., Travel time analysis of New York city police cars, *Interfaces*, 17(2), (**1987**), pp. 15-20.

[13] D.W. SCOTT, *Multivariate Density Estimation: Theory, Practice and Visualization*, Wiley Series in Probability and Statistics, Vol. 383 (John Wiley Sons, Hoboken, NJ) (**2009**).

[14] Four children killed in 3 alarm SW Philly fire, www.phillyfirenews.com/2014/07/05/four-4-kids-killed-3-alarm-sw-philly-fire/, downloaded on September 16, (**2016**).

[15] Gurobi Optimizer, http://www.gurobi.com/, downloaded on September 16, (**2016**).

[16] Map of Philadelphia Neighborhoods, http://www.arcgis.com/home/webmap/viewer.html?webmap=d21eb07ec54a4dcb9f80d057fe7c3cd6, downloaded on September 16, (**2016**).

[17] Three alarm fire at apartment building in west Philadelphia, Philadelphia.cbslocal.com/2016/06/12/three-alarm-fire-at-apartment-building-in-west-philadelphia, downloaded on September 16, (**2016**).

[18] PHILLY REPORT CRITICAL OF RESPONSE TIMES, http://www.philadelphiacontroller.org/publications/FireDepartment_ResponseTime_February2016.pdf, downloaded on July 18, (**2016**).