Ambulance Deployment under Demand Uncertainty

Sean Shao Wei Lam
Health Services Research and Biostatistics Unit, Singapore General Hospital, Singapore
Email Address: lam.shao.wei@sgh.com.sg

Yee Sian Ng
Department of Industrial and Systems Engineering, National University of Singapore, Singapore
Email Address: ngyeesian@gmail.com

Mohanavalli Rajagopal Lakshmanan
Health Services Research and Biostatistics Unit, Singapore General Hospital, Singapore
Email Address: rajagopal.mohanavalli@sgh.com.sg

Yih Yng Ng
Medical Department, Singapore Civil Defence Force, Singapore
Email Address: ng_yih_yng@scdf.gov.sg

Marcus Eng Hock Ong
Department of Emergency Medicine, Singapore General Hospital, Singapore
Duke-NUS Graduate Medical School, Singapore
Email Address: marcus.ong.e.h@sgh.com.sg

***Abstract*—In this study, we develop a robust model for the deployment of a fleet of ambulances under demand uncertainty. The proposed model remains computationally tractable in a full scale model, despite the explicit consideration of uncertain demands, for deriving plans that are robust against the worst case shortfalls in demand coverage. The explicit consideration of resource commitments overcomes the problem of cross coverage where each ambulance location has the ability to cover more than one demand point. An actual case study based on data gleaned from the Singapore's emergency medical services is described.**

***Index Terms*—ambulance deployment plans; demand uncertainty; response times**

I. Introduction

Emergency Medical Services (EMS) plays an important role in the provision of health care services under emergency conditions in the prehospital setting. The quality of EMS is determined by many factors, such as proficiency of the EMS personnel, ambulance efficiency, reliability and utilization rates, and other organizational, communications, command and control systems. Among the various performance indicators used to evaluate EMS performance, response time is an important EMS industry benchmark, amongst a basket of other indicators. The widespread adoption of response times as a key indicator of EMS performance can be attributed to its relationship to specific time sensitive conditions, such as out-of-hospital cardiac arrest[1]. A response times of within 4 minutes to deliver successful defibrillation can help increase the chances of survivals of such patients [2]. In reality, most of the critical emergency incidents served by EMS providers, such as, stroke and severe trauma cases [3], are also time-sensitive conditions.

There are various response time targets used in different countries. In some US states for example, the target is set to cover 95% of the emergency incidents within 10 minutes in the urban area and within 30 minutes in the rural area [4]. In Singapore, the national EMS provider uses a quality indicator of 11 minutes as a maximum response time threshold for 80% of all emergency incidents [5]. In essence, time is a vital factor in emergency situations. Therefore, it is critical that ambulances be effectively pre-positioned at the correct locations in anticipation of emergency calls, so that there can be adequate coverage within the target response time thresholds.

This research provides a systematic review of existing mathematical programming (MP) models for the deployment planning of ambulances, and proposes a new model for planning under demand uncertainty which overcomes some limitations of existing models. The proposed model remains computationally tractable despite the explicit consideration of uncertain demands in

Manuscript received Sept. 19, 2014; revised Dec. 29 2014.

doi: 10.12720/joams.4.3.187-194

order to derive plans that are robust against the worst case shortfalls in demand coverage.

II. Background

Resource allocation problems can be classified into strategic, tactical and operational decision-making problems according to the decision horizon [6]. Similarly, real-world EMS providers generally faced the same short, medium and long term decision problems in the planning and allocation of EMS resources. Some of the decision problems faced by these providers are listed in Table I.

Operational planning by EMS providers considers short-term decisions such as for ambulance dispatch and dynamic ambulance relocations. Tactical planning involves medium term decision horizons which typically establish baseline deployment plans (static or dynamic) and baseline manpower shift schedules for operational planning. Strategic planning involves longer term decision horizons and attempts to develop strategies that are robust against long-term scenario uncertainties. These strategies may be related to strategic policy changes (such as in the increase of private ambulance providers in a public-private EMS partnership) or long term infrastructural decisions (e.g., sizing of ambulance fleets and expansion of EMS manpower). In essence, baseline tactical plans to guide operational decisions should be robust against short term uncertainties, whereas strategic plans should be robust against longer-term uncertainties. Table 1 lists some example of decision problems faced by EMS providers for the respective decision horizons.

TABLE I. Decision Problems Faced by EMS Providers

Planning Problem	Planning Horizon	Decision Problems
Strategic	Long term (typically 5-10 years)	• EMS appliance fleet sizing[7] • Siting of EMS bases[8,9] • Public and private EMS operations policies[10] • EMS manpower planning[11] • Ambulance diversion policies[12]
Tactical	Medium term (can range from few months to a year)	• Static EMS appliance deployment planning[13] • Developing dynamic system status plans[5,14] • EMS appliance and manpower scheduling and shift planning[5,15]
Operational	Short-term (daily or weekly decisions)	• Real-time ambulance dispatch [16] • Dynamic ambulance relocation[17,18] • Tiered ambulance dispatch [19] • Ambulance routing and diversion[20]

A primary objective for tactical planning is to develop baseline resource allocation schedules that are robust against short term uncertainties in demand and supply, such as those related to EMS call volumes, travel times and fleet reliability. Resource allocation issues range from the staff/shift scheduling to the development of dynamic system status plans [21]. One of the most prevalent tactical problems is that related to the location of bases and allocation of the EMS appliances to bases. These appliances may include ambulances and fast response paramedics on motorbikes. Simulation and heuristical approaches have been proposed to deal with these problems [9, 16, 22, 23].

A more rigorous quantitative approach is to develop a mathematical programming (MP) model with the aim of deriving a mathematically optimal ambulance deployment strategy under resource constraints. There are several existing ambulance location and relocation models based on MP approaches. These models have evolved over time and undergone different phases of evolution that have brought about the development of both deterministic and stochastic models under different operational needs and regulatory requirements[24]. Some examples of deterministic and stochastic ambulance deployment models will be discussed in the next section. These models will motivate the development of a new shortfall-aware stochastic model that will handle cross-coverage and demand uncertainties without sacrificing computational efficiencies.

III. Ambulance Deployment Models

This section provides a condensed discussion of relevant MP models for the deployment planning of EMS appliances. The discussion on existing models will systematically motivate the development of a new deployment model which considers demand uncertainties. An adaptation to deal with the issue of cross coverage in existing deterministic models will also be presented. In order to facilitate the discussions of these models, we first introduce some necessary notations as follows:

Notations:

Sets

$\mathbb{R}$	Set of real numbers
$\mathbb{Z}$	Set of integers
$\mathbb{B}$	Set of binary numbers

Model Parameters

r	Coverage standard in time units (minutes)
t_{ij}	Travel time from node i to node j
x_i	Number of ambulances to deploy at location i
c_i	Cost of deploying ambulances to location i
a_{ij}	Adjacency index defined on G such that $a_{ij} = 1$ if $t_{ij} \leq r$ and $a_{ij} = 0$ otherwise
λ_j	Cost associated with shortfalls of coverage in node j
z_j	Number of ambulances committed to demand node j
p	Number of ambulances available for deployment
u	Budget available for ambulance deployment

Matrices and Vectors

$\mathbf{c} \in \mathbb{R}_+^{\|I\|}$	Vector of costs in deploying ambulances to location i, c_i
$\mathbf{x} \in \mathbb{R}_+^{\|I\|}$	Vector of number of ambulances deployed to ambulance location i, x_i
$\mathbf{z} \in \mathbb{R}^{\|J\|}$	Vector of number of ambulances committed to demand location j
$\mathbf{A} \in \mathbb{B}^{\|I\| \times \|J\|}$	Adjacency matrix with elements a_{ij}

Decision Variables

d_j	Demand for ambulances at location j
$d \in \mathbb{R}+\|J\|$	Vector of demands in node j, d_j

We define ambulance location models on directed graphs $G = (V,E)$ whose edges are defined as $E=\{(i,j) : i \in I, j \in J\}$ and vertices, $V = (I, J)$, that is partitioned into two disjoint sets: (1) set of potential ambulance locations, $I \in V$, and; (2) set of demand points, $J \in V$. Given a coverage standard r, we define coverage of a demand point j by location i only when $t_{ij} \leq r$. In addition, we define the following notations:

A. Deterministic Models

One of the earliest deterministic models introduced for ambulance deployment planning is the location set covering model (LSCM) [25]. The objective of this model is to determine the minimum number of ambulances required to ensure all demand points covered by explicitly modeling the costs of deployment. The LSCM is given as follows:

LSCM Primal

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \quad & A^T\mathbf{x} \geq \mathbf{d} \\ & \mathbf{x} \geq 0 \end{aligned}$$

LSCM Dual

$$\begin{aligned} \max_{\boldsymbol{\lambda}} \quad & \boldsymbol{\lambda}^T\mathbf{d} \\ \text{s.t.} \quad & \boldsymbol{\lambda}^T A^T \geq \mathbf{c} \\ & \boldsymbol{\lambda} \geq 0 \end{aligned}$$

The dual formulation of the LSCM focuses on the maximization of demand coverage. The dual has been introduced in several early research given its closer relevance to the operations of EMS as a public service. Following the LSCM, the maximal covering location problem (MCLP) was also proposed in an early research[26]. The MCLP attempts to maximize the total demand covered given a fixed ambulance fleet size, thereby introducing a hard constraint which may render the problem infeasible.

A widely adopted target in public EMS service providers is to ensure an acceptably high proportion of incidents can be served within a predetermined threshold. The travel time targets are typically in the range of 8-11 minutes for 80-90% percentiles of all cases. Given such a target-oriented objective, neither the primal nor the dual instructs us on the optimal policy. A more appropriate model may consider the minimization of the overall shortfall in coverage, instead of maximizing the demand coverage. The shortfall is directly related to the proportion of incidents that can be served within the predetermined threshold. In order to deal with such a problem, a Minimal Shortfall Location Problem (MSLP) can be formulated as follows:

MSLP

$$\begin{aligned} \min_{x,z} \quad & \prec^T(\mathrm{d}-\mathrm{z}) \\ s.t. \quad & \mathrm{A}^T\mathrm{X} \geq \mathrm{Z} \\ & \mathrm{C}^T\mathrm{X} \leq u \\ & \mathrm{X} \geq 0 \end{aligned}$$

In the MSLP, budget constraints have been included. These budget constraints can also be formulated in terms of the number of ambulances available for deployment. The MSLP is in fact a generalization of the MCLP.

The MSLP model will continue to minimize shortfalls in coverage against the target response time, irrespective of whether the number of ambulances committed is in excess or shortage over the demands across all nodes, subjected to budget/capacity constraints. In order to more accurately model the decision criteria based on shortfall minimization, we model the following objective instead: $\boldsymbol{\lambda}^T \max\{0, (\mathbf{d}-\mathbf{z})\}$. In the consideration of uncertainties in the MP model, such a shortfall-aware objective criterion has been shown to be a special case of the success probability criterion, but possesses computationally desirable qualities in comparison to the success probability criterion under general distributional assumptions for a model considering uncertainties, and distributional ambiguities [27], [28].

Both the LSCM and MCLP approaches are static models and do not consider the possibility that a particular ambulance will be busy to answer emergency calls when it is needed–problems of cross coverage (or multiple overlapping demands). In order to handle the issue of cross coverage, several deterministic and stochastic approaches have been proposed. These models include the backup coverage models (BACOP) [29] and the double standard model (DSM). Another deterministic model which explicitly considers the problem of backup coverage is the Double Coverage Model (DCM) [30]. Even though the BACOP, DSM or DCM provides for the cross coverage, thereby mitigating the probability of the nearest ambulances being unavailable to answer emergency calls, these models did not consider stochastic demands.

B. Extension of Deterministic Model

Even without considering stochastic demands, one of the limitations of the MSLP formulation is that only the resource commitment information at the demand nodes is explicitly considered. There are no detailed specifications of where the resources will be committed from. This will result in insufficient coverage if multiple demand nodes can be reached from a base location within acceptable travel time thresholds. As an example, assume that we have the following adjacency matrix, **A**, defining the connectivity between two bases and three demand locations within a pre-specified travel time threshold:

$$\mathrm{A} = \begin{bmatrix} 110 \\ 011 \end{bmatrix}$$

For a demand scenario of cover a demand scenario of $\mathbf{d}^T = [1, 2, 1]$, a deployment configuration of $\mathbf{x}^T = [1,1]$ will completely satisfy the constraint $A^T\mathbf{x} \geq \mathbf{z}$ resulting in zero shortfalls, hence, an optimal solution. However, the ambulances at the base locations will not be sufficient to cover all three demand nodes simultaneously. In order to deal with this problem, we introduce an additional decision variable to represent "commitments" at the base locations, y_i, and decompose the adjacency matrix into two node-arc sub-matrices. Before describing the model, we need to introduce the following additional notations:

Decision Variables

y_k	Number of ambulances deployed to serve edge k, $k \in E$
$\mathbf{y} \in \mathbb{R}_+^{\lvert E \rvert}$	Vector of number of ambulances deployed to serve edge k, y_k

Model Parameters

b_{ik}	Node-arc adjacency index where $b_{ik} = 1$ if node i is the base location serving edge k and $b_{ik} = 0$ otherwise for $i \in I$
h_{jk}	Node-arc adjacency index where $h_{jk} = 1$ if node j is the demand location served by edge k and $h_{jk} = 0$ otherwise for $j \in J$
$B \in \mathbb{B}^{\lvert I \rvert \times \lvert E \rvert}$	Node-arc adjacency matrix with elements b_{ik}
$H \in \mathbb{B}^{\lvert I \rvert \times \lvert E \rvert}$	Node-arc adjacency matrix with elements e_{jk}

Given the above notations, the MSLP model can be reformulated to overcome the problem of overlapping demands in a static model:

$$\begin{aligned} &\text{MSLP(S1)} \\ &\min_{x,y,s} \quad \prec^{T} S \\ &s.t. \quad S \geq \mathrm{d} - \mathrm{Hy} \\ &\qquad\; \mathrm{X} \geq \mathrm{By} \\ &\qquad\; \mathrm{C}^{T}\mathrm{X} \leq u \\ &\qquad\; \mathrm{X} \geq 0; S \geq 0 \end{aligned}$$

C. Stochastic Models

One of the first models to consider demand uncertainty is the Maximal Expected Coverage Location Problem (MEXCLP) [31]. MEXCLP introduces a "busy fraction" q as the probability that any given ambulance will be unavailable to respond to an incoming emergency call. The motivation for this model is similar in spirit to the deterministic models which consider contingent coverage, such as the DCM. Another similar stochastic formulation proposed the use of a reliability level in a model which sought to maximize the demand covered with specified probability [32].

In order to consider stochastic demands, a two-stage stochastic programming (SP) model can also be formulated from the MSLP(S1) problem with the assumption that demands are generated by a random vector $\mathbf{d}(\omega)$ for scenarios $\omega \in \Omega$ with state-space Ω equipped with a set (σ-algebra) $\Im$ of events and a probability measure, P. Representing the second stage decisions by $\tilde{\mathrm{s}} = [\mathrm{s}(\omega)]_{\omega\in\Omega}$, the SP model can be formulated as follows:

$$\begin{aligned} &\text{MSLP(SP)} \\ &\min \quad E[Q(X, \tilde{S})] \\ &s.t. \quad \mathrm{X} \geq \mathrm{By} \\ &\qquad\; \mathrm{C}^{T}\mathrm{X} \leq u \\ &\qquad\; \mathrm{X} \geq 0; Y \geq 0 \end{aligned}$$

where

$$\mathrm{Q(x, s)} = \min \quad \prec^{T} s \quad \text{St} \quad s \geq d(\omega) - \mathrm{Hy}, \quad s \geq 0$$

The MSLP(SP) can be solved using the Sample Average Approximation (SAA) approach. However, one problem with adopting an SP approach is related to combinatorial intractability when stochastic demand scenarios are considered across all demand nodes. Furthermore, the optimal SAA solutions may vary under different realizations of the scenarios.

D. Extension of Stochastic Model

An alternative to SP is to consider the idea of decision robustness. A robust decision model can be formulated to ensure that the constraints associated with the uncertain demands may only be violated up to a certain acceptable robustness level [33], [34]. This threshold level can be defined as the reliability level of the robust solution. Assuming that the probability of keeping to the optimal shortfall levels is at most κ , the uncertain inequalities in MSLP(S1) can be formulated as a joint chance constraint as follows, giving rise to the chance constrained model MSLP (CC):

$$P(\tilde{d} - Hy \leq s) \leq \kappa \tag{1}$$

κ can be interpreted as the reliability of the solution given optimal shortfall levels. Its complement *(1-κ)* is similar to the probability of unavailability given the optimal shortfalls under the resource constraints. The implication for this is that the estimated level of unavailability based on historical data becomes an exogenous decision variable, instead of an endogenous estimate from data (which can be difficult to estimate). Although this is not an empirical probability estimate, it can be viewed as a projected level of unavailability. Together with the optimal shortfalls, this probability of unavailability will be useful for considering resource expansion decisions.

Considering MSLP (CC) and assuming independent demands and using Markov inequality, the left-hand side of constraint(1) in MSLP(CC) for each demand node j has the following lower bounds:

$$P\left(\tilde{d}_j \geq s_j + \sum_{k\in E} h_{jk} \mathrm{y}_{\mathrm{K}}\right) \geq 1 - \frac{E(\tilde{d}_j)}{s_j + \sum_{k\in E} h_{jk} \mathrm{y}_{\mathrm{K}}} \tag{2}$$

Consequently, (1) and (2) can be reformulated to arrive at a deterministic robust formulation. Considering demands at each discretized demand node and time segment follows independent Poisson distributions with arrival rates λ_j for all $j \in J$, the deterministic robust counterpart(RC) for MSLP(CC) is given by replacing the chance constraints with a set of deterministic constraints based on the Poisson arrival rates, giving rise to the MSLP(RC) formulation.

IV. Case Study

The case study is based on the national EMS system of Singapore. Emergency ambulance services in Singapore are provided by the national EMS service provider - the Singapore Civil Defence Force (SCDF). In 2011, the SCDF operated 46 ambulances. A central command centre for the SCDF coordinates all emergency calls through a national "995" emergency hotline for civil emergencies. The ambulances are then physically deployed from bases that may be fire stations or fire posts.

Fire posts are satellite locations where ambulances and fast response fire fighting vehicles may be deployed.

Upon emergency activation for medical and trauma emergencies, the nearest available ambulance will be deployed. The nearest available ambulance may be on standby in the bases, at the hospitals or returning to base following the conveyance of patients to hospitals. Emergency medical treatment will be given to the patients on-scene and assigned a triage status following a Patient Acuity Category (PAC) scale, ranging from priority 1 (PAC1) to priority 4 (PAC4) in decreasing levels of patient severity. Upon conveyance to hospitals, patients will be handed over to the emergency department of the respective hospital. Following the hospital handover, the ambulance will then be made available for serving the next incoming demand [5, 22].

In this study, data of all the unique cases of emergencies from 1st January 2011 to 30th June 2011 was used as the source of estimating the arrival rates of emergency calls. Fig. 1 shows the distribution of call demands over four time periods, Monday, Sunday and the rest of the weekdays. The demand profile can be distinguished into 8 hourly time blocks: from 0900 hours to 1600 hours, 1700 hours to 0000 hours and 0100hours to 0800 hours. The total demands from Tuesdays to Saturday follow a similar pattern, as compared to Sundays and Mondays. Among the total demands across all the days, Mondays typically have the largest average demands. Fig. 2 shows a chloropeth plot of the geospatial distribution of demand volumes according to demand volume per square km.

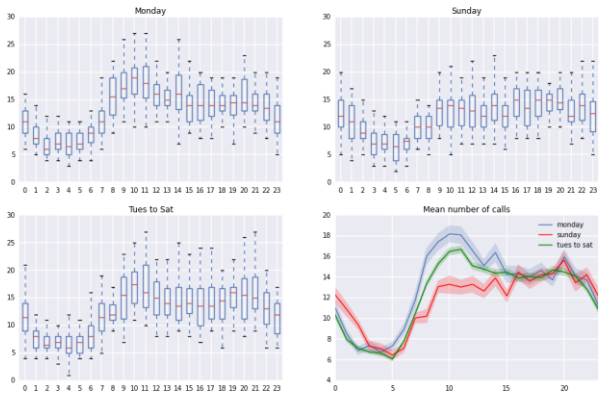

Figure 1. Temporal distribution of call demands across distinct weeks and time periods

For the MP modelling approach, the whole of Singapore is rasterized into 30 by 30 equally sized rectangular cells, and historical incidents that happened within each cell region were aggregated for each cell region within each hour. Only 782 cells were retained when the land area of Singapore's main island was considered. In the model, time was discretized in hourly segments. The distribution of turnaround times for all calls within the six monthly time period follows a lognormal distribution with a median of approximately 35 minutes. Fig. 3 shows the empirical turnaround and inter-arrival times between calls overlaid with lognormal and exponential distributions respectively.

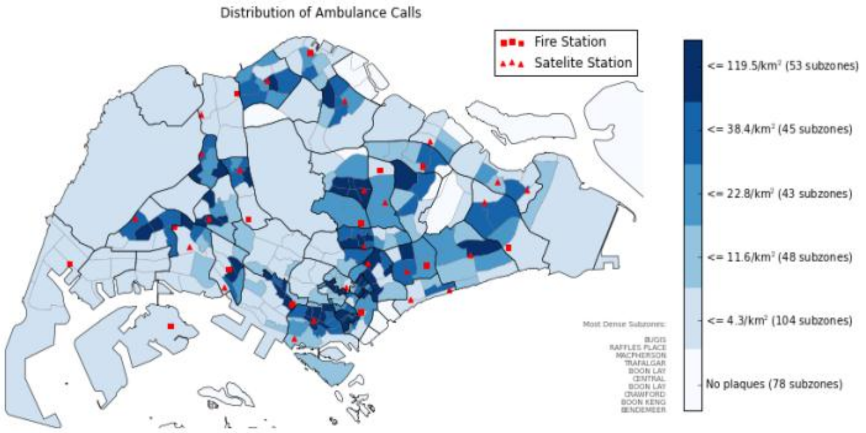

Figure 2. Geospatial distribution of demand volumes per square km

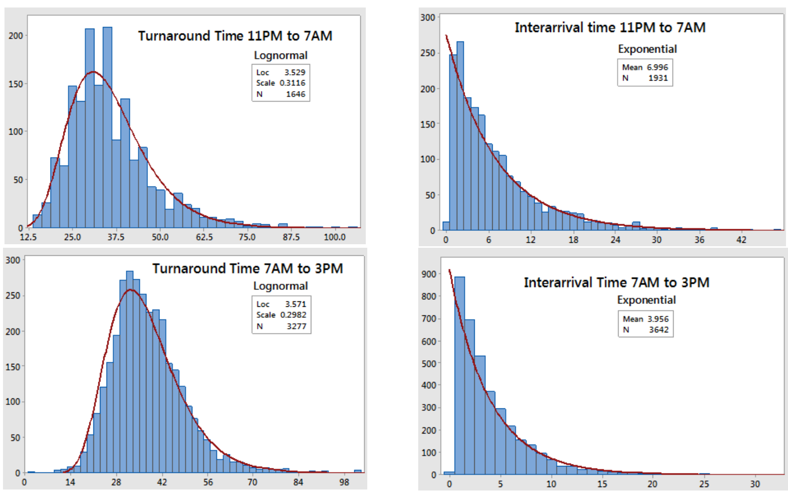

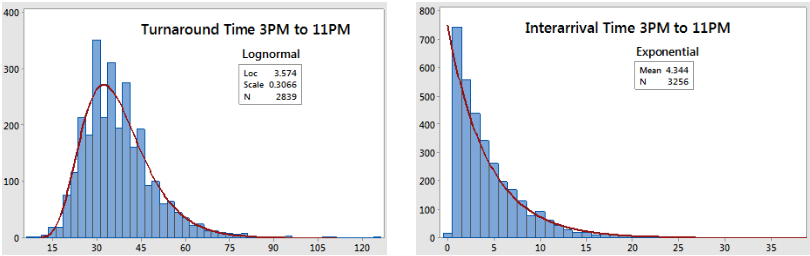

(a) (b)

Figure 3. Distributional assumptions: (a) turnaround times following lognormal distributions, and; (b) inter-arrival times following exponential distributions

V. Results and Discussions

Altogether 53, 300 incident calls were enrolled in this study. The breakdown of the characteristics of ambulance calls is shown in Table II.

TABLE II. Characteristics of Ambulance Calls in Singapore

Incident Type	**Total**	**(%)**
Medical	39640	(74%)
Trauma	12707	(24%)
False/Cancelled Calls/ Assistance not Required	953	(2%)
PAC+	**Total**	**(%)**
PAC0*	1161	(2%)
PAC1	6351	(12%)
PAC2	33332	(63%)
PAC3	11375	(21%)
PAC4	1059	(2%)
Age Group	**Total**	**(%)**
<10	1262	(2%)
10-19	1899	(4%)
20-29	5503	(10%)
30-39	5266	(10%)
40-49	5657	(11%)
50-59	7709	(14%)
60-69	7551	(14%)
70-79	7431	(14%)
>79	11022	(21%)
Day_Of_Week	**Total**	**(%)**
Sunday	7399	(14%)
Monday	8127	(15%)
Tuesday	7614	(14%)
Wednesday	7612	(14%)
Thursday	7698	(14%)
Friday	7336	(14%)
Saturday	7514	(14%)
Time of the Day	**Total**	**(%)**
11PM-7AM	12242	(23%)
7AM-3PM	21777	(41%)
3PM-11PM	19281	(36%)

+ 22 calls with unknown PAC status.
* Additional category of PAC0 refers to patients declared dead on scene

The MSLP(RC) model with an objective function which minimizes the worst case short falls was solved over three 8-hourly shifts for 782 rasterized demand cells using AIMMS software (AIMMS, Haarlem, The Netherlands) with CPLEX 12 solver (IBM Corporation, New York, US). A maximum of two ambulances can be deployed in each base. The optimal deployment plan for each of the shift is shown in Fig. 4.

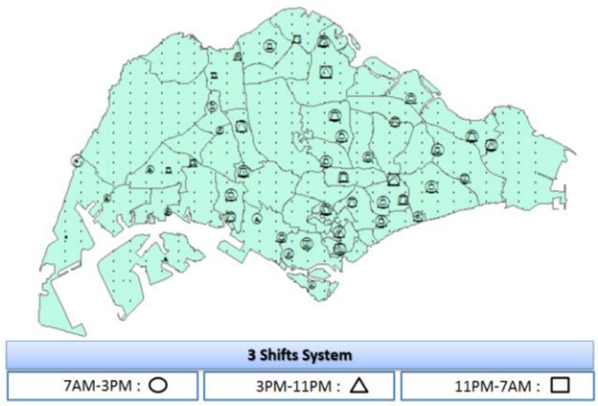

Figure 4. Deployment plans for a three shifts system with maximum capacity of two ambulances per base

Coverage proportions are the proportions of demand covered in each node assuming an acceptable travel time threshold. Fig. 5 shows that reliability level of 80% provides the best coverage proportions across all the demand nodes for the worst 30-50 demand nodes in terms of coverage proportions. The remaining demand nodes not shown in Fig. 5 essentially have 100% coverage under the recommended deployment plans). Consequently, a reliability level of 80% was chosen for the derivation of the optimal deployment plans.

The optimal solutions may not be the best deployment plan in consideration of the numerous practical complexities and uncertainties confronting decision makers that were not considered in the model. For example, travel time uncertainties have not been explicitly considered in the MSLP(RC) formulation. Given these limitations, the use of discrete events simulations (DES) [5,23] will help to provide more assurance in the quality of the deployment. The quality of the robust solution can also be effectively compared via a realistic DES model against other alternative strategies.

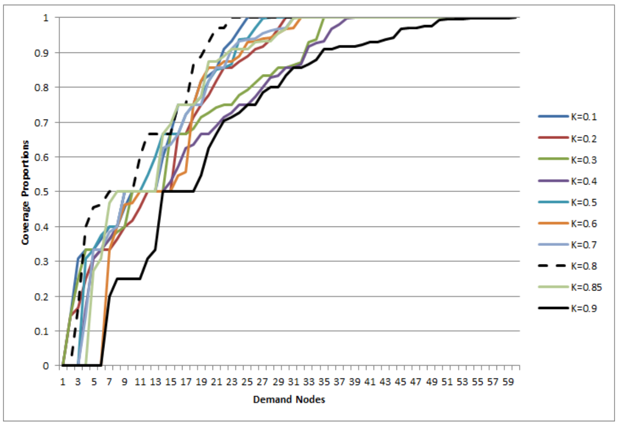

Figure 5. Coverage proportions across all demand nodes under different reliability levels

Apart from DES models, geospatial visualizations on how the coverage changes across multiple time thresholds over the entire demand region will enable decision makers to identify gaps, or areas of low coverage proportions, to focus on. Adaptive augmentation of stakeholder's perspectives can be incorporated into a iterative decision making framework that consist of the MSLP (RC) based optimal plan as the starting point for developing more practically realistic and convincing deployment strategies. Such a decision support framework is proposed in Fig. 6.

VI. Conclusions and Recommendations

This study proposed a robust MP model for the deployment of ambulances under demand uncertainty. The explicit consideration of resource commitments within the network overcomes the problem of cross coverage. A case study based on the Singapore's EMS system demonstrates that the deterministic reformulation of the robust model retains computational tractability for the deployment planning of ambulances in real systems.

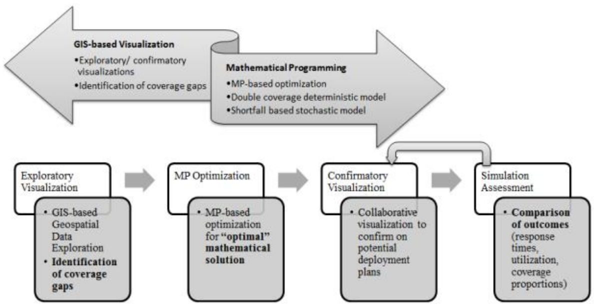

Figure 6. Decision support framework incorporating MP approach, discrete events simulations and geospatial visualizations

Acknowledgment

This work was supported in part by a grant (HSRNIG12Nov011) from the National Medical Research Council, Singapore.

References

[1] T. D. Valenzuela, D. J. Roe, G. Nichol, L. L. Clark, D. W. Spaite, and R.G. Hardman, "Outcomes of rapid defibrillation by security officers after cardiac arrest in casinos," *N Engl J Med*, vol. 343, no. 17, pp. 1206-1209, 2000.

[2] T. D. Valenzuela, D. W. Spaite, H. W. Meislin, L. L. Clark, A. L. Wright, and G. A. Ewy, "Emergency vehicle intervals versus collapse-to-CPR and collapse-to-defibrillation intervals: Monitoring emergency medical services system performance in sudden cardiac arrest," *Ann Emerg Med*, vol. 22, no. 11, pp. 1678-83, 1999.

[3] R. Sánchez-Mangas, A. García-Ferrrer, A. D. Juan, and A. M. Arroyo, "The probability of death in road traffic accidents. How important is a quick medical response?" *Accident Analysis & Prevention*, vol. 42, no. 4, pp. 1048-1056, 2010.

[4] M. O. Ball and F. L. Lin, "A reliability model applied to emergency service vehicle location," *Operations Research*, vol. 41, no. 1, pp. 18-36, 1993.

[5] S. S. W. Lam, Z. C. Zhang, H. C. Oh, Y. Y. Ng, W. Wah, and M. E. H. Ong, "Reducing ambulance response times using discrete event simulation," *Prehospital Emergency Care*, vol. 18, no. 2, pp. 207-216, 2013.

[6] S. Zeltyn, Y. N. Marmor, A. Mandelbaum, B. Carmeli, O. Greenshpan, Y. Mesika, S. Wasserkrug, P. Vortman, A. Shtub, T. Lauterman, D. Schwartz, K. Moskovitch, S. Tzafrir, and F. Basis, "Simulation-based models of emergency departments: Operational, tactical, and strategic staffing," *ACM Transactions on Modeling and Computer Simulation*, vol. 21, no. 4, 2011.

[7] A. J. Fischer, P. O'Halloran, P. Littlejohns, A. Kennedy, and G. Butson, "Ambulance economics," *Journal of Public Health Medicine*, vol. 22, no. 3, pp. 413-421, 2000.

[8] X. Li, Z. Zhao, X. Zhu, and T. Wyatt, "Covering models and optimization techniques for emergency response facility location and planning: A review," *Mathematical Methods of Operations Research*, 2011. vol. 74, no. 3, pp. 281-310.

[9] B. Adenso-Díaz and F. Rodríguez, "A simple search heuristic for the MCLP: Application to the location of ambulance bases in a rural region," *Omega*, vol. 25, no. 2, pp. 181-187, 1997.

[10] S. G. Reissman, "Privatization and emergency medical services," *Prehospital and Disaster Medicine*, vol. 12, no. 1, pp. 22-29, 1997.

[11] K. Kajino, T. Kitamura, T. Iwami, M. Daya, M. E. H. Ong, C. Nishiyama, T. Sakai, K. Tanigawa-Sugihara, S. Hayashida, T. Nishiuchi, Y. Hayashi, A. Hiraide, and T. Shimazu, "Impact of the number of on-scene emergency life-saving technicians and outcomes from out-of-hospital cardiac arrest in Osaka City," *Resuscitation*, vol. 85, no. 1, pp. 59-64, 2014.

[12] K. J. McConnell, C. F. Richards, M. Daya, C. C. Weathers, and R. A. Lowe, "*Ambulance diversion and lost hospital revenues*," *Annals of Emergency Medicine*, vol. 48, no. 6, pp. 702-710, 2006.

[13] L. C. Nogueira Jr, L. R. Pinto, and P. M. S. Silva, "Reducing emergency medical service response time via the reallocation of ambulance bases," *Health Care Management Science*, 2014.

[14] M. S. Maxwell, M. Restrepo, S. G. Henderson, and H. Topaloglu, "Approximate dynamic programming for ambulance redeployment," *INFORMS Journal on Computing*, vol. 22, no. 2, pp. 266-281, 2010.

[15] D. Cabrera, J. L. Wiswell, V. D. Smith, A. Boggust, and A. T. Sadosty, "A novel automatic staffing allocation tool based on workload and cognitive load intensity," *American Journal of Emergency Medicine*, vol. 32, no. 5, pp. 467-468, 2014.

[16] V. E. SilvaSouza and J. Mylopoulos, "Designing an adaptive computer-aided ambulance dispatch system with Zanshin: An experience report," *Software-Practice and Experience*, 2013.

[17] 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, pp. 216-231 2013.

[18] S. Dean, "The origins of system status management," *Emergency Medical Services*, vol. 33, no. 6, pp. 116-118, 2004.

[19] J. Stout, P. E. Pepe, and V. N. Mosesso Jr, "All-advanced life support vs tiered-response ambulance systems," *Prehospital Emergency Care*, vol. 4, no. 1, pp. 1-6, 2000.

[20] J. C. Pham, R. Patel, M. G. Millin, T. D. Kirsch, and A. Chanmugam, "The effects of ambulance diversion: A comprehensive review," *Academic Emergency Medicine*, vol. 13, no. 11, pp. 1220-1227 2006.

[21] J. Overton, J. Stout, and A. Kuehl, "System design, prehospital systems and medical oversight," *3rd ed. National Association of EMS Physicians.* Kendall/ Hunt Publishing, 2002.

[22] M. E. H. Ong, T. F. Chiam, F. S. P. Ng, P. Sultana, S. H. Lim, B. S. H. Leong, V. Y. K. Ong, E. C. Ching Tan, L. P. Tham, S. Yap, and V. Anantharaman, "Reducing ambulance response times using geospatial-time analysis of ambulance deployment," *Academic Emergency Medicine*, vol. 17 no. 9, pp. 951-957, 2010.
[23] C. H. Wu and K. P. Hwang, "Using a discrete-event simulation to balance ambulance availability and demand in static deployment systems," *Academic Emergency Medicine*, vol. 16, no. 12, pp. 1359-1366, 2009.
[24] G. Laporte, F. Louveaux, F. Semet, and A. Thirion, "Application of the double standard model for ambulance location," in *Innovations in Distribution Logistics*, J. A. E. E. Nunen, M. G. Speranza, and L. Bertazzi, Ed. Springer Berlin Heidelberg, 2009, pp. 235-249.
[25] C. Toregas, R. Swain, C. ReVelle, and L. Bergman, "The location of emergency service facilities," *Operations Research*, vol. 19, no. 6, pp. 1363-1373, 1971.
[26] R. Church and C. ReVelle, "The maximal covering location problem," *Papers of the Regional Science Association*, vol. 32, no. 1, pp. 101-118, 1974.
[27] D. B. Brown and M. Sim, "Satisficing measures for analysis of risky positions," *Management Science*, vol. 55, no. 1, pp. 71-84, 2009.
[28] S. W. Lam, T. S. Ng, M. Sim, and J. H. Song, "Multiple objectives satisficing under uncertainty," *Operations Research*, vol. 61, no. 1, pp. 214-227, 2013.
[29] K. Hogan, and C. ReVelle, "Concepts and applications of backup coverage," *Management Science*, vol. 32, no. 11, pp. 1434-1444, 1986.
[30] M. Gendreau, G. Laporte, and F. Semet, "Solving an ambulance location model by tabu search," *Location Science*, vol. 5, no. 2, pp. 75-88, 1997.
[31] M. S. Daskin, "A maximum expected covering location model: Formulation, properties and heuristic solution," *Transportation Science*, vol. 17, no. 1, pp. 48-70, 1983.
[32] C. ReVelle and K. Hogan, "The maximum availability location problem," *Transportation Science*, vol. 23, no. 3, pp. 192-200, 1989.
[33] S. L. Janak, X. Lin, and C. A. Floudas, "A new robust optimization approach for scheduling under uncertainty: II. Uncertainty with known probability distribution," *Computers & Chemical Engineering*, vol. 31, no. 3, pp. 171-195, 2007.
[34] X. Lin, S. L. Janak, and C. A. Floudas, "A new robust optimization approach for scheduling under uncertainty: I. Bounded uncertainty," *Computers & Chemical Engineering*, vol. 28, no. 6–7, pp. 1069-1085, 2004.

Sean Shao Wei Lam has a PhD and Masters in Industrial and Systems Engineering. He is currently the Manager of the Health Services Research Unit at Singapore General Hospital and Associate with the Centre for Quantitative Medicine. The primary focus of his research is on real-world healthcare operation research problems for the improvements of health services and healthcare operations. His publications and research have received several awards, including the Best Applications Paper Award from IIE and MTI Innovation Award.

Yee Sian Ng is a graduate student at the Operations Research Center, Massachussetts Institute of Technology. He previously worked as a Research Assistant at the National University of Singapore, before doing a Research Fellowship with the Future Urban Mobility group under the Singapore-MIT Alliance.

Ms. Mohanavalli Rajagopal is a Research Scientist at the Division of Research, Health Services Research in Singapore General Hospital. Ms. Mohana is actively involved in research projects in pre-hospital emergency care, strategies to improve patient waiting time at the Emergency Department and simulation modelling of emergency operations. She has presented her research in international and local conferences. Ms. Mohana holds a Master's degree in Logistics from Nanyang Technological University and awaiting PhD graduation from Singapore MIT Alliance.

COL(DR) Yih Yng Ng is currently the Chief Medical Officer and Director, Medical Department of the SCDF. He is a consultant to the Singapore Ministry of Health in the area of pre-hospital emergency care and works as an emergency physician at the Singapore General Hospital. He is also the vice-chairman of the Asian EMS Council.

Professor Marcus Eng Hock Ong is a Senior Consultant, Director of Research, and Clinician Scientist at the Department of Emergency Medicine in Singapore General Hospital. He is also the Director for the Health Services Research and Biostatistics Unit, Division of Research, in SGH. He serves as a Senior Consultant at the Hospital Service Division and Director for the Unit for Pre-hospital Emergency Care in Ministry of Health, Singapore. He is also Associate Professor at Duke-National University of Singapore Graduate Medical School. His research studies focus predominantly on pre-hospital emergency care, medical devices, and health services research. He has published 100 papers in international and local peer-reviewed journals.