# Approximating Vehicle Dispatch Probabilities for Emergency Service Systems with Location-Specific Service Times and Multiple Units per Location 

Susan Budge (sbudge@ualberta.ca) ${ }^{2}$<br>Armann Ingolfsson (armann.ingolfsson@ualberta.ca) ${ }^{1,2}$<br>Erhan Erkut (erkut@bilkent.edu.tr) ${ }^{3}$<br>${ }^{1}$ Corresponding author<br>${ }^{2}$ School of Business<br>University of Alberta<br>Edmonton, Alberta, T6G 2R6<br>${ }^{3}$ Faculty of Business Administration<br>Bilkent University<br>Bilkent, Ankara, Turkey

31 May 2007


#### Abstract

To calculate many of the important performance measures for an emergency response system one requires knowledge of the probability that a particular server will respond to an incoming call at a particular location. Estimating these "dispatch probabilities" is complicated by four important characteristics of emergency service systems. We discuss these characteristics and extend previous approximation methods for calculating dispatch probabilities, to account for the possibilities of workload variation by station, multiple vehicles per station, call and station dependent service times, and server cooperation and dependence.


## 1. Introduction

The primary performance measures for an emergency service system are typically cost and some measure of speed of response, for example, average response time or fraction of calls responded to within some time standard (i.e., coverage). Therefore, when evaluating a strategic or operational change, it is important to accurately estimate the impact on speed of response. An essential input for calculating coverage and other response time performance measures is the probability that an incoming call at a particular location is served by a particular server. When these dispatch probabilities are known, many system-wide performance measures can be calculated easily, by conditioning on the location of the call and the location of the server, and then using the law of total probability.

Broadly speaking, there are three alternatives for evaluating the performance of the system we are interested in: (1) an exact approach, such as Larson's (1974) exact hypercube model, (2) discrete event simulation, or (3) approximations, such as Larson's (1975) approximate hypercube (AH) model. Compared to simulation and exact approaches, the advantages of the various versions of the AH model are that they are fast, with computation times that are relatively insensitive to system characteristics, and they are sufficiently accurate for many practical purposes. In many cases, we believe it is appropriate to use an approximation to facilitate comparison of alternatives (for example, as part of an optimization heuristic for station location, vehicle allocation, or shift scheduling) and then to use simulation or an exact approach to further evaluate the most promising alternatives. Approximate approaches generally make simplifying assumptions regarding one or more of the following four system characteristics:

1) Number of vehicles per station: It is common to assume only one vehicle per station. This is a restrictive assumption because neighborhoods with high demand density, fixed costs of building a station, or limitations on the number of available station sites may make it economical or necessary to have multiple vehicles at the same station.
2) Average workload: Some models assume that all vehicles have the same average utilization irrespective of the vehicle's home station. This is unrealistic because spatial variation in demand and transport network characteristics will tend to create imbalances in workload.
3) Average service time: Some models assume that average service time (the time a vehicle is unavailable to respond to new calls while responding to a specific call) is either independent of the location of the vehicle's home station, or independent of the location of the call, or both. The service time depends on these two locations due to the travel time between the two, but components other than the travel time might depend on the location of the call or responding station as well.
4) Server cooperation: Two extremes are to either assume that each station operates as an independent subsystem or assume that any call is equally likely to be responded to by any available vehicle. These extremes simplify modeling, but reality is somewhere in between. Assumptions of no server cooperation and a single vehicle per station imply that the status (busy or idle) of an ambulance is statistically independent of the status of other vehicles. Thus, the "no cooperation" assumption is related (but not equivalent) to the common "server independence" assumption.

All of these assumptions are violated to a significant degree in real systems. Larson's (1975) AH model allows workload to vary between servers and models server cooperation, but it assumes one server at each station and average service times that are independent of both the server and call location. Jarvis (1985) extended the AH model to allow average service times to depend on the server and call location.

Our main contribution is a version of the AH model that computes station-specific (rather than server-specific) busy fractions and dispatch probabilities, and allows multiple vehicles at a station. We find it more natural and convenient to organize the model around stations rather than servers for the systems that we have worked with. We give a theoretical convergence guarantee for a restricted version of the procedure (see Goldberg and Szidarovszky (1991) for related results). Burwell, Jarvis, and McKnew (1993) extended the AH model to account for ties in preferences for servers. Although there may be reasons for preference ties other than multiple vehicles at a station, in our experience these reasons are not common, and by focusing on multiple vehicles at a station rather than the general case of preference ties, we obtain a simpler procedure. The main difference between our work and that of Burwell, Jarvis, and McKnew (1993) is that they use the original correction factors developed by Larson (1975), assuming random sampling of vehicles, whereas we introduce a new set of correction factors, based on random sampling of stations.

Our procedure makes none of the simplifying assumptions about the four characteristics discussed above. The procedure is fast and converges in all cases that we have tried. It generally provides estimates of busy fractions with $2 \%$ relative error or less when compared with simulation, which is considerably more accurate than a procedure that ignores server dependence.

The remainder of the paper is structured as follows. Section 2 defines notation, derives stationspecific correction factors, and presents our algorithm and a convergence guarantee. Section 3 summarizes computational testing of the procedure. An online companion contains a more extensive literature review, a comparison of the AH model to simulation and the exact hypercube model, a derivation of the station-specific correction factors, a proof of the convergence result provided in
the paper, a discussion of sensitivity to the service time distribution, and additional computational results.

## 2 The Approximation Procedure

We assume that $s$ vehicles are distributed among $I$ stations, with $s_{i}$ vehicles at station $i$, with $s_{i} \geq 1$ for all i. Calls for service arrive according to independent Poisson processes from $J$ demand nodes, with arrival rate $\lambda_{j}$ from node $j$, and a total arrival rate of $\lambda=\sum_{j=1}^{J} \lambda_{j}$. The average service time for calls originating at node $j$ served by an ambulance from station $i$ is $\tau_{i j}$. This includes the average travel time between station $i$ and node $j$, the average time spent on scene, the average transport time to a hospital, and the average time spent at the hospital. We assume a fixed dispatch policy, where the preference of station $i$ in the dispatch order for node $j$ is given by $a_{i j}$ (for example $a_{i j}=3$ means that station $i$ is the $3^{\text {rd }}$ most preferred for responding to a call from node $j$ ). We assume that calls that arrive when all vehicles are busy do not queue but are handled through other means, for example by supervisor vehicles. This "no queue" assumption is common in the literature (for example see Jarvis, 1985) and it is consistent with the operation of the emergency service systems that we are familiar with. We refer to calls that arrive when all vehicles are busy as "lost," and define $P_{s}$ to be the probability that a call is lost. This loss probability is the same for all demand nodes, assuming that the node arrival rates are not influenced by which vehicles are busy. As Larson (1975), we use an $M / M / s / s$ system to approximate certain aspects of the system of interest. According to a well-known insensitivity result, steady occupancy probabilities for the $M / M / s / s$ model are insensitive to the shape of the service time distribution beyond the mean (Gross and Harris, 1998). Computational experiments by Jarvis (1985) suggest that the system we are modeling is also relatively insensitive to the shape of the service time distributions. This is important because service time distributions for emergency services will often be far from exponential.

We define the dispatch probabilities as

$$
\begin{equation*}
f_{i j}=\operatorname{Pr}\{\text { vehicle from station } i \text { responds } \mid \text { call from node } j\} \tag{1}
\end{equation*}
$$

with $\sum_{i=1}^{I} f_{i j}+P_{s}=1$ for all $j$. We use $r_{i}$ for the utilization of servers from station $i, \rho$ for the system-wide average offered load per server, and $r=\rho\left(1-P_{s}\right)$ for the system-wide average server utilization. If we know the distribution for the response time $R_{i j}$ for all station-node pairs, then we can calculate
performance measures such as coverage or average response time (for calls that are not lost) using

$$
\begin{align*}
& \text { coverage }=\sum_{j=1}^{J} \frac{\lambda_{j}}{\lambda} \sum_{i=1}^{I} f_{i j} \operatorname{Pr}\left\{R_{i j} \leq \text { time standard }\right\}  \tag{2a}\\
& \text { average response time }=\sum_{j=1}^{J} \frac{\lambda_{j}}{\lambda} \sum_{i=1}^{I} \frac{f_{i j}}{1-P_{s}} \mathrm{E}\left[R_{i j}\right] \tag{2b}
\end{align*}
$$

We can also compute coverage or average response times for specific demand nodes, with straightforward modifications to equations (2).

The procedure we present generalizes Larson's (1975) AH model. As in Jarvis (1985), we start by applying Little's law to the $s_{i}$ servers at station $i$. The dispatch rate for this station equals $\sum_{j=1}^{J} \lambda_{j} f_{i j}$. The average service time for calls that station $i$ responds to is $\sum_{j=1}^{J} \lambda_{j} f_{i j} \tau_{i j} / \sum_{j=1}^{J} \lambda_{j} f_{i j}$. Little's law implies that the average number of busy servers at station $i, s_{i} r_{i}$, equals the total dispatch rate to the station multiplied by the average service time for calls that the station responds to, which implies

$$
\begin{equation*}
r_{i}=\frac{1}{s_{i}} \sum_{j=1}^{J} \lambda_{j} f_{i j} \tau_{i j} \tag{3}
\end{equation*}
$$

The only unknown quantities on the right-hand-side of $(3)$ are the dispatch probabilities $f_{i j}$. If we could approximate these probabilities as a function of known quantities and the busy fractions $r_{i}$, then we would have the ingredients for an iterative procedure for estimating the busy fractions and dispatch probabilities.

Like Larson (1975) and Jarvis (1985), we begin with the "server independence" assumption. When $s_{i}=1$ for all $i$ and station $i$ is the $k^{\text {th }}$ preferred for node $j\left(a_{i j}=k\right)$, this assumption leads to approximating the dispatch probability, $f_{i j}$, with the product of the probabilities that ambulances at all more preferred stations are busy, multiplied with the probability that station $i$ has a free ambulance. To improve this approximation, Larson and Jarvis multiplied this product with a factor $Q$ to approximately correct for the erroneous assumption of server independence, i.e.,

$$
\begin{equation*}
f_{i j} \approx Q(s, \rho, k) \prod_{l=1}^{k-1} r_{l() j}\left(1-r_{i}\right), \tag{4}
\end{equation*}
$$

where $r_{(l) j}$ is the busy fraction for the $l^{\text {th }}$ preferred station for node $j$ and $\rho$ is an estimate of the overall system load per server (we discuss how to estimate $\rho$ at the end of the section). The correction factor involves occupancy probabilities for an $M / M / s / s$ loss system. Denoting the steady state probability that the loss system has $n$ customers by $P_{n}$, the correction factor in (4) can be expressed as

$$
\begin{equation*}
Q(s, \rho, k)=\frac{P_{0}}{s!\left(1-\rho\left(1-P_{s}\right)\right)} \cdot \frac{(s-k)!}{\left(1-P_{s}\right)^{k-1}} \cdot\left(\sum_{u=k-1}^{s-1} \frac{(s-u) \cdot s^{u} \cdot \rho^{u-k+1}}{(u-k+1)!}\right) \tag{5}
\end{equation*}
$$

Combining equations (3), (4), and (5) leads to an iterative procedure for approximating the busy fractions and dispatch probabilities, for a situation where each station has a single server.

To allow for more than one ambulance at some stations, we generalize equations (4) and (5). The multi-vehicle counterpart to equation (4) is

$$
\begin{equation*}
f_{i j} \approx Q_{j}\left(\left\{s_{(k) j}\right\}, \rho, k\right) \prod_{l=1}^{k-1} r_{(l) j}^{s_{l(j)}}\left(1-r_{i}^{s_{i}}\right), \tag{6}
\end{equation*}
$$

where $s_{(k) j}$ is the number of ambulances at the $k^{\text {th }}$ preferred station for node $j$ and we continue to assume that $a_{i j}=k$. In this equation, the correction factor depends not only on $s, \rho$, and $k$, but also on how the $s$ ambulances are distributed between stations and on the node $j$. This is because the probability that an ambulance from station $i$ is dispatched to a call from node $j$ depends not only on the number of stations that are preferred (by node $j$ ) to station $i$ but also on the number of ambulances at those stations. If the problem is constrained to allow only one ambulance per station, then this distinction between ambulances and stations disappears.

The station specific correction factors can be expressed as follows:

$$
\begin{equation*}
Q_{j}\left(\left\{s_{(k) j}\right\}, \rho, k\right)=\frac{P_{0} \sum_{n=z_{(k-1) j}}^{s-1} \frac{(\rho s)^{n}}{n!}\left[\prod_{u=0}^{z_{(k-1)} j^{-1}} \frac{n-u}{s-u}-\prod_{u=0}^{z_{(k) j}-1} \frac{n-u}{s-u}\right]}{r^{z_{(k-1) j}}\left(1-r^{s_{(k) j}}\right)} \tag{7}
\end{equation*}
$$

Equations (3), (6), and (7) provide the building blocks that we use in our algorithm to estimate stationspecific busy fractions and dispatch probabilities, as we describe next.

The algorithm begins by calculating the following;

$$
\begin{aligned}
b_{k j} & =k^{\text {th }} \text { preferred station for node } j \\
s_{(k) j} & =s_{b_{k j}}=\text { number of vehicles at } k^{\text {th }} \text { preferred station for node } j \\
z_{(k) j} & =s_{(1) j}+s_{(2) j}+\ldots+s_{(k) j} \\
\tau_{(k) j} & =\tau_{b_{k j j} j}
\end{aligned}
$$

Next, initialize the busy fractions and the system-wide average service time, by assuming that the system operates as an $M / M / s / s$ queue (superscripts are used as iteration counters):

$$
\begin{aligned}
& \tau^{0}=\frac{1}{\lambda s} \sum_{i=1}^{I} s_{i} \sum_{j=1}^{J} \lambda_{j} \tau_{i j} \\
& r_{i}^{0}=r^{0}=\lambda \tau^{0}\left(1-P_{s}^{0}\right) / s
\end{aligned}
$$

where $P_{s}^{0}$ is calculated using Erlang's loss formula. Set the iteration counter, $h$, to 1 and enter the iterative process. Each iteration consists of the following steps:

Step 1: Use $\tau^{h-1}, \lambda$, and $s$ to calculate $P_{0}^{h}$ and $P_{s}^{h}$.

Step 2: Calculate $V_{i}^{h}$ for all $i$, using (13) and the following:

$$
\begin{equation*}
V_{i}^{h}=\sum_{j=1}^{J} \lambda_{j} \tau_{i j} Q_{j}\left(\left\{s_{(k) j}\right\}, \rho^{h-1}, a_{i j}\right) \prod_{l=1}^{a_{j}-1}\left(r_{(l) j}^{h-1 \text { or } h}\right)^{s_{(l) j}} \tag{8}
\end{equation*}
$$

where $r_{(l) j}{ }^{h-1 \text { or } h}$ means $r_{(l) j}{ }^{h}$ if it has been computed, and otherwise $r_{(l) j}{ }^{h-1}$. In other words, we use a Seidel process, i.e., we always use the most recently computed station utilization. Then, update the station-specific busy fractions using (9) if $r^{h-1} \leq 0.5$ and using (10) otherwise.

$$
\begin{gather*}
r_{i}^{h}=\frac{V_{i}^{h}}{s_{i}+\left(r_{i}^{h-1}\right)^{s_{i}-1} V_{i}^{h}}  \tag{9}\\
r_{i}^{h}=\left(\frac{V_{i}^{h}}{V_{i}^{h}+s_{i} /\left(r_{i}^{h-1}\right)^{s_{i}-1}}\right)^{\frac{1}{s_{i}}} \tag{10}
\end{gather*}
$$

Step 3: Calculate $f_{i j}{ }^{h}$ using (6), normalize these probabilities using $f_{i j}^{h} \leftarrow f_{i j}^{h}\left(1-P_{s}^{h}\right) / \sum_{i=1}^{I} f_{i j}^{h}$ and calculate $\tau^{h}, \rho^{h}$, and $r^{h}$ using

$$
\tau^{h}=\frac{1}{\lambda\left(1-P_{s}\right)} \sum_{j=1}^{J} \lambda_{j} \sum_{i=1}^{I} f_{i j}^{h} \tau_{i j}, \quad \rho^{h}=\lambda \tau^{h} / s, \quad r^{h}=\frac{1}{s} \sum_{i=1}^{I} s_{i} r_{i}^{h}
$$

Step 4: If $\left|r_{i}^{h}-r_{i}^{h-1}\right|<\varepsilon$ for all $i$ then stop. Otherwise, set $h=h+1$.

Some observations about the algorithm are in order. First, there is no assurance that the dispatch probabilities $f_{i j}$, calculated using (6), will satisfy the normalization condition $\sum_{i=1}^{l} f_{i j}+P_{s}=1$. We found that enforcing this normalization in Step 3 tended to improve accuracy of the station-specific utilizations slightly. Larson (1975) discusses this and other more complex methods for normalizing dispatch probabilities. Second, equations (9) and (10) are both obtained by rearranging equation (3) after substituting equations (6) and (7). We found that equation (9) led to faster convergence for low system loads while equation (10), which guarantees that $0<r_{i}^{h}<1$, led to more reliable convergence for high system loads; hence our suggestion to use (9) when the initial system load estimate $r^{0}$ is less than or equal to $50 \%$ and to use (10) otherwise. Finally, although the algorithm above has converged in all of our experiments, we have no theoretical guarantee that it will always do so. However, a restricted version of the algorithm is guaranteed to converge (Goldberg and Szidarovszky, 1991, and our online supplement). Theorem 1: Suppose that the average service time $\tau$, the system load $\rho$, the loss probability $P_{s}$, and the correction factors $Q_{j}\left(\left\{s_{(k) j}\right\}, \rho, k\right)$ are calculated only once, at the beginning of the algorithm, equation (16) is always used to update the busy fractions, and a Gauss (rather than Seidel) iterative process is used. If the busy fractions are initialized to $r_{i}^{0}=1$ for all $i$, then the algorithm is guaranteed to converge.

## 3. Computational Results

We implemented the algorithm in VBA for MS Office. The computational times given below are based on this implementation, running on a 2.4 GHz PC .

We used two datasets to investigate the convergence of the algorithm. Using information from Burwell (1986) for Greenville County, South Carolina, with 99 demand nodes and 5 stations, we generated 3,124 cases. The procedure converged in all cases, requiring at most 10 iterations and 0.09 seconds of computation. The second dataset, from Edmonton, Alberta (Ingolfsson, Erkut, and Budge, 2003), consists of 180 demand nodes and 10 stations with specified capacities ranging from one to four, resulting in 55,403 cases. For this larger problem, the algorithm again converged in all cases, but required up to 32 iterations and 0.5 seconds to converge. The algorithm has been used as part of a decision support system for the Calgary (Alberta) Emergency Medical Services system, with over 1,000 demand nodes and 20-40 ambulances, depending on time of day. In this setting, each run of the algorithm takes about 15 seconds. We see that despite the speed of the method, computation time will limit the
number of scenarios that a user may wish to explore. If simulation of an exact solution method were used instead, the computation times would be considerably larger.

We compared the results of our procedure to the results of a discrete event simulation model, for 12 ambulance allocations using the Edmonton dataset, in order to evaluate the accuracy of the estimation procedure. We simulated each allocation for system loads ( $\rho=\lambda \tau / s$ ) ranging from 0.1 to 0.9 by varying the total arrival rate of calls to the system. We compared simulated and estimated vehicle utilizations for 108 cases ( 12 allocations $\times 9$ system utilizations). For ease of reading, for the remainder of the section we use "average error" to mean the average (across stations) of the absolute errors and "average relative error" to mean the average (across stations) of the relative errors. Table 1 shows the results, with typical average relative errors below $2 \%$. The average relative errors tend to be highest for system loads in the middle of the range (and lower for very low or very high system loads). This contrasts with the results under the independence assumption (the shaded cells in Table 1), with average relative errors often exceeding $2 \%$ and in general showing a pattern of increasing errors with increasing system loads. Table 2 shows an example of the results using our procedure, for a realistic allocation with 14 ambulances and system load of 0.3 . The agreement between the simulated and approximated busy fractions is rather good, with most of the relative errors below $2 \%$.

Figure 1 shows the impact of modeling the server unavailability on coverage (calculated using (2a)), and specifically includes the model with a constant system-wide busy fraction as a comparison point. First, comparing the estimated coverage when the vehicles are assumed to always be available (i.e., using a system-wide busy fraction of zero) to that when assuming a constant system-wide busy fraction (using the average of the estimated station-specific busy fractions), the overall coverage of the system is seriously overestimated (a difference of over 6\%). Next, dispatch probabilities based on a system-wide busy fraction assumption in turn overestimate the coverage of the system compared to the more realistic dispatch probabilities calculated using station-specific busy fractions and correction factors for dependence. Although the difference may seem small at $1.9 \%$, in previous work we found that such a difference was actually very significant, in that it would require fairly major changes to the system (for example adding two ambulances around the clock) in order to attain such a difference when the system coverage was in the vicinity of $90 \%$. Additionally, if the coverage values for sections of the city are considered, the differences can be much larger. Consider the two extreme cases; the neighborhoods around the stations with the highest and lowest estimated busy fractions respectively. In the first case, the ambulances are busier than average and so the coverage for this area is overestimated (by about 5.4\%) when assuming a constant system-wide busy fraction. In the latter case, the ambulances are less
busy than average and so the coverage for this area is underestimated (by almost $8 \%$ ) when assuming a constant system-wide busy fraction. This is significant because system designers may be concerned with equity in coverage between different communities, in addition to the overall system-wide coverage.

Other performance measures, such as average response times and frequency of interdistrict responses, are easily calculated using the estimated dispatch probabilities but will not be considered in detail here.

Table 2: Results for a sample ambulance allocation by station and averaged across stations (including estimated and simulated utilizations along with absolute and relative errors).

| Station, $i$ | $\mathbf{1}$ | $\mathbf{2}$ | $\mathbf{3}$ | $\mathbf{4}$ | $\mathbf{5}$ | $\mathbf{6}$ | $\mathbf{7}$ | $\mathbf{8}$ | $\mathbf{9}$ | $\mathbf{1 0}$ | Average |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Servers, $s_{i}$ | 1 | 2 | 1 | 1 | 1 | 1 | 2 | 2 | 1 | 2 | 1.4 |
| Estimated $r_{i}$ | 0.32 | 0.15 | 0.17 | 0.26 | 0.33 | 0.34 | 0.28 | 0.32 | 0.37 | 0.47 | 0.30 |
| Simulation $r_{i}$ | 0.33 | 0.16 | 0.17 | 0.26 | 0.33 | 0.35 | 0.28 | 0.32 | 0.37 | 0.46 | 0.30 |
| Absolute error | 0.006 | 0.006 | 0.003 | 0.000 | 0.002 | 0.008 | 0.001 | 0.003 | 0.001 | 0.008 | 0.004 |
| Relative error | $1.9 \%$ | $3.9 \%$ | $1.9 \%$ | $0.0 \%$ | $0.6 \%$ | $2.4 \%$ | $0.2 \%$ | $1.0 \%$ | $0.2 \%$ | $1.7 \%$ | $1.4 \%$ |



Figure 1: Comparison of estimated coverage for the same system using a system-wide busy fraction versus station-specific busy fractions.

## Acknowledgements

This research was partially funded by NSERC. We thank Stevanus Tjandra, Trevor Riehl, Dan Haight, Ramon Alanis, and Jonathan Cheng for useful suggestions and for programming assistance. The comments of the referees and Associate Editor led to substantial improvements to the paper.

## References

Burwell, T. (1986). A Spatially Distributed Queueing Model for Ambulance Systems. Ph.D. Dissertation. Departments of Management and Mathematical Sciences, Clemson University, Clemson.

Burwell, T. , J. Jarvis, M. McKnew (1993). Modeling Co-located Servers and Dispatch Ties in the Hypercube Model. Computers and Operations Research 20 113-119.

Church, R., C. ReVelle (1974). The Maximal Covering Location Problem. Papers of the Regional Science Association 32 101-118.

Goldberg, J., F. Szidarovszky (1991). Methods for Solving Nonlinear Equations Used in Evaluating Emergency Vehicle Busy Probabilities. Operations Research 39 903-916.

Gross, D., C. M. Harris (1998). Fundamentals of Queueing Theory, Third Edition. Wiley, New York.
Ingolfsson, A., E. Erkut, and S. Budge (2003). Simulating a Single Start Station for Edmonton EMS. Journal of the Operational Research Society 54 736-746.

Jarvis, J. (1985). Approximating the Equilibrium Behavior of Multi-Server Loss Systems. Management Science 31 235-239.

Larson, R. C. (1974). A Hypercube Queueing Model for Facility Location and Redistricting in Urban Emergency Services. Computers and Operations Research 1 67-95.

Larson, R. C. (1975). Approximating the Performance of Urban Emergency Service Systems. Operations Research 23 845-868.

Table 1: Average relative errors (in percent) for 12 ambulance allocations and 9 utilization levels (Edmonton dataset) resulting from our procedure and


