IMPROVING THE EMERGENCY SERVICE DELIVERY IN ST. ALBERT

ERHAN ERKUT[1], RUSSELL FENSKE[2]
STEVEN KABANUK[3], QUENTIN GARDINER[4], JERRY DAVIS[5]

[1]*School of Business, University of Alberta*
[2]*Intuit Canada, Edmonton, Alberta*
[3]*Centre for Operations Excellence, University of British Columbia*
[4]*Stout & Company Chartered Accountants, Edmonton, Alberta*
[5]*G. P. Davis & Associates Limited, St. Albert, Alberta*

ABSTRACT

This paper summarizes two studies conducted to improve the emergency service delivery in St. Albert, a small city near Edmonton, Alberta. The first study dealt with selecting the location of a new fire station. The goals of the study included an assessment of the performance of the current system, an identification of the area(s) with poor coverage, a selection of a site among a set of given candidate locations, and an assessment of the improvement in the system performance upon the addition of the new fire station. A geographical information system was used for storing and displaying the spatial data, computing service areas for given travel times, and for communicating the results of the study. The second study considered an evaluation of the resources available for emergency service. A probabilistic model was used to evaluate labor costs for different platoon sizes and a simulation model was used to evaluate the adequacy of the current staff and fleet sizes. Results were presented to EMS officers, city staff, as well as the City Council.

RÉSUMÉ

Cet article résume deux études réalisées pour améliorer la prestation de services d'urgence à St. Albert, une petite ville près d'Edmonton, Alberta. La première étude visait le choix d'un emplacement pour un nouveau poste du service des incendies. Les objectifs de l'étude incluaient une évaluation de la performance du système actuel, l'identification des zones avec une faible couverture de service, la sélection d'un site parmi un ensemble de candidats possibles et l'évaluation de l'amélioration du service suite à l'ajout du nouveau poste. Un système d'information à référence spatiale a été utilisé pour enregistrer et afficher les données spatiales, calculer des zones de service pour des temps d'intervention donnés, et pour présenter les résultats de l'étude. La deuxième étude portait sur une évaluation des ressources disponibles pour les services d'urgence. Un modèle probabiliste a été employé pour évaluer les coûts de main-d'œuvre pour différentes tailles d'équipe et un modèle de simulation a été employé pour évaluer l'adéquation des tailles présentes du personnel et de la flotte. Les résultats ont été présentés aux responsables des services médicaux EMS, au personnel de ville, ainsi qu'au conseil municipal.

1. INTRODUCTION

This paper describes two projects conducted for the Fire Service department of St. Albert. The first one is a location study and the second one is a resource evaluation study. This section serves as an introduction to both major parts of the paper.

St. Albert is a small city with a population of about 50,000 immediately Northwest of Edmonton, Alberta. Although the first settlement in St. Albert dates back to 1863, the population growth in the area is quite recent and St. Albert became a city in 1981. St. Albert is a planned community where 85% of the developed area contains single-family dwellings and the rest consists of low-rise multi-unit dwellings and institutional developments. The St. Albert Fire Department is a full-time, career organization that provides emergency services around the clock. Currently there are two fire stations in St. Albert, and one of these stations (Fire Hall #1) doubles as an ambulance station.

The Fire Service provides an integrated Fire and Emergency Medical Service (EMS) response

Recd. May 2001, Rev. Aug. 2001, Acc. Sept. 2001

to St. Albert, and (under contract) to defined areas in Sturgeon County and Edmonton Garrison. As well, mutual aid agreements exist with surrounding municipalities to provide back-up services to their forces, as needed. Services within St. Albert include fire prevention, fire suppression, rescue (e.g. traffic accidents, ice/water rescue), hazardous materials response, EMS response, emergency medical transfers, and disaster services. Service is provided by a staff of 57 composed of a Fire Chief, a Deputy Chief, a Fire Marshal, an EMS Coordinator, an Administrative Resource Coordinator, four Call Room Operators, and 48 Firefighter/EMS Personnel. A Medical Director is contracted, as required by legislation.

The quality of the emergency medical services provided by the City at the present time is acceptable. The EMS department can respond to over 90% of all calls within 9 minutes—a widely accepted standard in this sector. However, the City officials are concerned about deterioration in the quality of the service in the near future for the following reasons:

1. The city's population is growing and (perhaps more importantly) aging. This increases the call rate.
2. Recent changes to the health care system resulted in patients being discharged earlier from hospitals and persons with a higher level of acuity being taken care of at home or in non-hospital care environments. This results in increased call rates for ambulances.
3. The city is planning an expansion. Three new developments have been approved. This increases the total area that has to be covered (as well as the population).
4. The current Firefighter/EMS staff is aging which decreases the amount of time they are available (vacation entitlement and sick leave usage increase with length of service and age respectively), and age may increase risk of injury. In addition, as the current staff retires, there will be significant loss of experience.

Increased call volumes and longer travel distances, coupled with more shifts employing minimum staffing, impose an increased risk of not having sufficient resources available to answer calls. The City is concerned about the level of emergency service it provides and is interested in finding ways to maintain an acceptable service level. Options considered are the addition of a new fire hall, the addition of staff to each platoon, and the addition of new vehicles (fire trucks or ambulances).

The remainder of this paper is organized as follows. Section 2 describes a fire station location study, and Section 3 summarizes the study of the current resource levels. Conclusions follow in Section 4.

2. FIRE HALL LOCATION STUDY

2.1 Introduction

Currently there are two fire stations in St. Albert, and the city is considering adding a third station. (Figure 2.1 shows a map of the city with the two fire stations.) Two of the three fire stations would double as ambulance stations. The goal of the city is to respond to as many structures as possible within 5 minutes of an emergency call—a rather ambitious target. The third fire station location is to be chosen so as to maximize the number of homes and businesses that are within 5-minute travel of a fire station. The goals of the study are:

- Evaluate the coverage provided by the existing two stations
- Evaluate locations for a third station
- Recommend locations for ambulances

This study is not concerned with the number of emergency vehicles or staff.

2.2 Data

The decision problem has three main components: demand points, fire station locations, and travel distances.

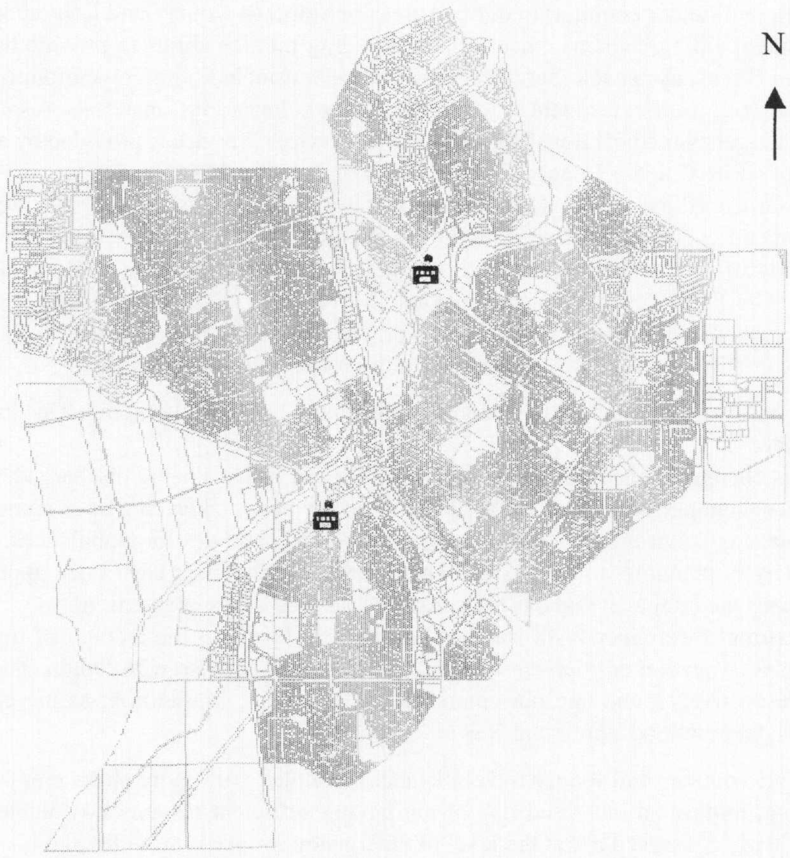

Figure 2.1: Map of St. Albert. The two fire stations (Churchill in the South and Boudreau in the North) are marked with station icons. The new developments are the less dense areas in the Northwest, North, and Northeast.

Demand points: The city provided us with a GIS database that contains all 16,000 existing dwellings as well as areas identified for future growth. We added a few hundred points to represent the new subdivisions and used this database for analysis with ArcView. We also aggregated the demand points to 1,000 points (using a clustering method) for spreadsheet analysis to reduce the computational burden. While any aggregation results in perturbations, we believe that the errors introduced by aggregating 16,000 points into 1,000 centroids are insignificant.

Existing and potential fire station locations: The locations of the two existing stations are known. In the first phase of the project, we identified the areas of the town that were not covered by the existing stations within the desired time limit. In the second phase, the city identified seven potential locations for the third station and we evaluated these candidate sites in terms of the additional coverage they provided.

Travel distances: We are interested in measuring the travel distances between demand points and service facilities. While we considered representing collections of dwellings as points on the plane, storing their coordinates, and using a distance formula to estimate travel distances, we opted for an explicit model of the travel network to compute actual travel distances for increased accuracy.

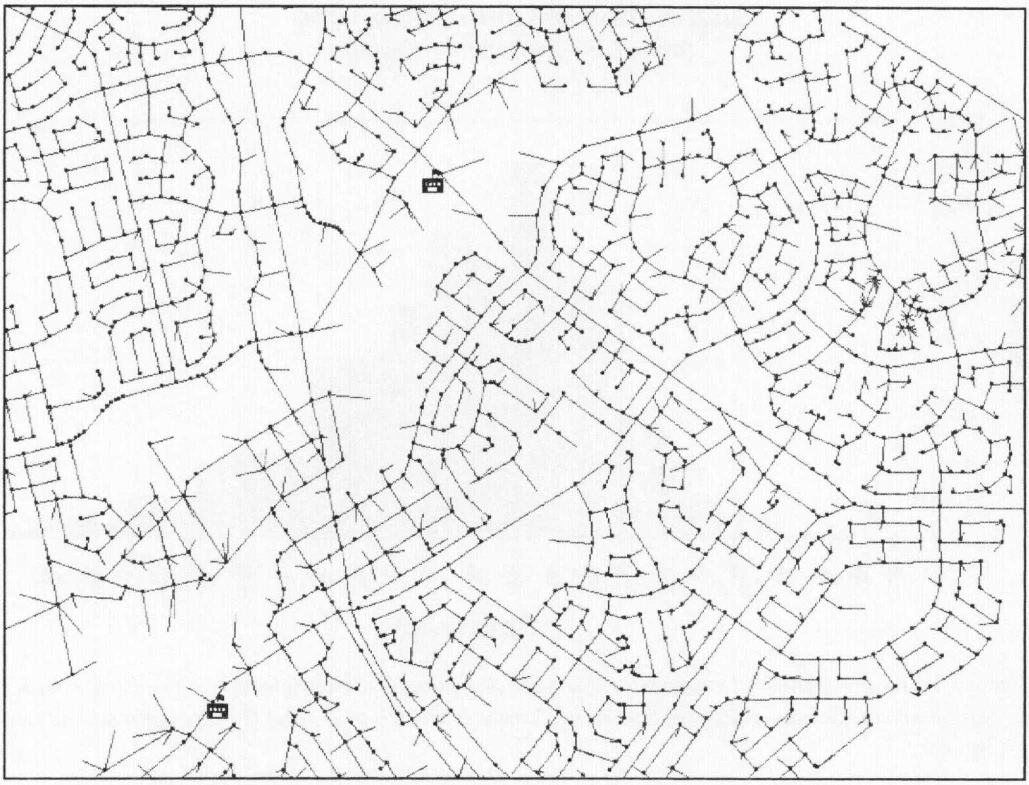

Figure 2.2: Central portion of the improved street network of St. Albert (improvements described in Section 2.3.1). The two fire stations are marked with station icons.

2.3 Preparation for Analysis

2.3.1. Improving the road network Although the city gave us a GIS database, the road network in this database was not directly useful for the modeling of travel times. We improved the road network in the GIS database in the following ways to prepare it for analysis:

- Broke up long links into shorter ones by adding new nodes to the network (to model road intersections and curvatures).
- Patched some areas of the network where links were not connected.
- Added the 1,000 demand aggregation points to the network by connecting each one to the closest node.
- Added artificial nodes to link the fire stations to the transport network.
- Added demand points to represent the demand in the new subdivisions and connected the new subdivisions to the existing road network.

While some of the changes were relatively simple and could be completed within ArcView using its tools, others required the export of data out of ArcView and involved some programming in C or in VBA (Excel macros). Figure 2.2 displays part of the network used for subsequent analysis.

2.3.2. Conversion of distances to times. This element is crucial since response time is the measure of service quality in this problem. The conversion of travel distances to travel times is not trivial since the average travel speed is not constant. We classified the road segments in three

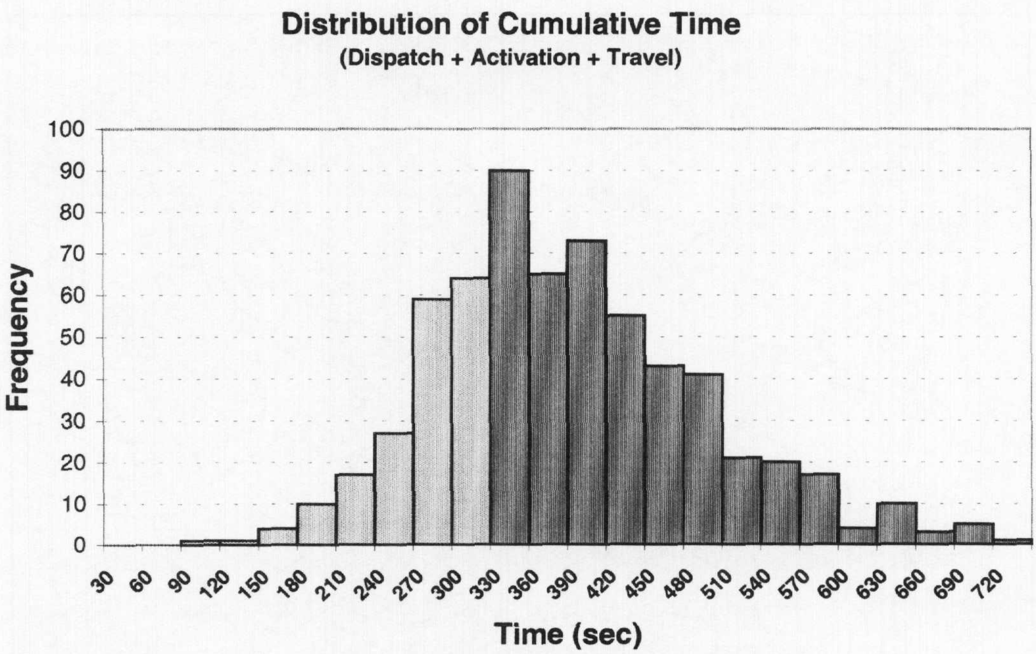

Figure 2.3: The response-time histogram for past data. Times over 5 minutes (the desired response time) are associated with darker bars. Response time is over 5 minutes for 70% of all calls. The average response time is 6 minutes.

groups: St. Albert Trail (a freeway that goes through the city), main roads, and residential roads. In addition, we considered the impact of the time of day (rush hour vs. non rush hour) and season (summer vs. winter) on travel speeds. We used several inputs for our travel speed estimates:

- Historical call data (We used a shortest path algorithm coded in VBA to compute the length of the shortest path between the responding station and the call location. This allowed us to compute an average speed for 750 past calls. The average speed was about 50 km/hr.)
- Interviews with four drivers
- Road test in a fire truck
- Road test in a car

We decided to use a total of 10 different average speeds, ranging from a low of 23 km/hr (residential streets in the winter) to a high of 67 km/hr (St. Albert Trail during summer non-rush hour). Our conversion of distances to times takes into account the season (winter vs. summer), time of day (rush vs. non-rush), and the road type (St. Albert Trail, major arteries, and residential roads). We exclude a number of other factors such as right vs. left turns, weekday vs. weekend, fire truck vs. ambulance, aggressiveness of the driver (and others on the road). After the completion of this phase, we asked EMS drivers to validate a few of the routes found by ArcView. As we had expected, the drivers were more likely to prefer the fastest routes to the shortest paths. As well, there seemed to be agreement that the fastest routes generated would be those taken by the drivers.

2.3.3. Analysis of historical response times. We used past response data to benchmark the status quo, to establish average travel speeds, and to validate our results. The city provided us with a database that contained every EMS call (750 calls) made during January–June 1999. Most importantly this database contained the location of the call and the response time. An analysis of this database showed that the system met the widely accepted EMS standard of responding to

90% of calls within 9 minutes. However, it was far from being able to respond to every call within 5 minutes. In fact we found that the response time was under 5 minutes for about 30% of the calls and the average response time was 6 minutes.

In our preliminary analysis of the response time data, we discovered that our travel time estimates and the actual response times differed by several minutes consistently. Hence, we focused on the definition of "response" in this phase, and noticed that a response time consists of three parts: dispatch time (the time that passes between the arrival of the emergency call and the dispatch of the vehicles), activation time (the time that passes between the dispatch and the rolling out of the vehicle) and travel time. The median dispatch time was 69 seconds, and the median activation time was 90 seconds, implying that approximately 2.5 minutes was spent before travel started. It seemed to us that this time could be cut back by at least one minute if EMS crews were to be activated as soon as a call arrived. While this would inconvenience the crews to some extent (crews in both stations would have to start preparing without knowing the nature or the address of the emergency), we thought that the reduction in response time might be worth the inconvenience. We decided to keep the dispatch and activation times separate from the travel times, and used three different scenarios for dispatch plus activation: 1.5 minutes (optimistic), 2 minutes, 2.5 minutes (pessimistic). Note that pessimistic does not mean worst case here (about 40% of the calls required more than 2.5 minutes of setup time).

2.4 Preliminary ArcView and Excel Analysis

Using the network analysis tool of ArcView we found service areas for the two fire stations for the following travel times: 2, 3, 4, and 5 minutes using two different sets of average travel speeds: summer non-rush (highest average travel speeds) and winter rush (lowest average travel speeds). (Note that a 5-minute travel time translates into a 9-minute response time if the setup takes 4 minutes – a pessimistic scenario.) We displayed the covered portion of the network using different colors in ArcView. Figure 2.4 shows the coverage areas for the winter rush scenario (in black-and-white). This figure indicates that, while a number of small pockets in the peripheral areas are not covered within 5 minutes, by far the most exposed area was the new development in the Northwest.

To complement the GIS-based analysis, we conducted some analysis on Excel as well. We selected 120 candidate points at random in the city and sought an answer to the question "Given the two existing locations, where should the third facility be located to maximize the coverage?" using different travel speed and setup time scenarios. There were no surprises; in every run a point in the new subdivision in Northwest was selected as the best location. This finding was consistent with the expectations of the fire department and we focused in this area for the final location.

2.5 Selecting the Final Location

We used nine different scenarios on a spreadsheet to compare the seven candidate sites provided to us by the city: travel times changing from 1.5 minutes to 3.5 minutes in increments of 0.25 minutes. (We focused on low travel times since we were most interested in the impact of the new station in its immediate vicinity.) For each given travel time we compared the seven sites using the summer non-rush and winter rush travel speeds. For every one of the 18 scenarios considered, we evaluated the additional coverage provided by each of the seven candidates for the third station using a spreadsheet. We normalized the additional coverage figures so that the maximum was always 100%. Hence, the coverage percentage for a candidate location indicated what percent of the maximum addition coverage the site was offering. Although the seven sites were in the same subdivision, there were some significant differences between the additional coverage provided. For example, the two extreme locations in the North and the South only provided about half of the additional coverage provided by the more central locations (Figure 2.5). We recommended the selection of one of the two central locations.

To help the decision-makers assess the incremental value of the third station, it was important

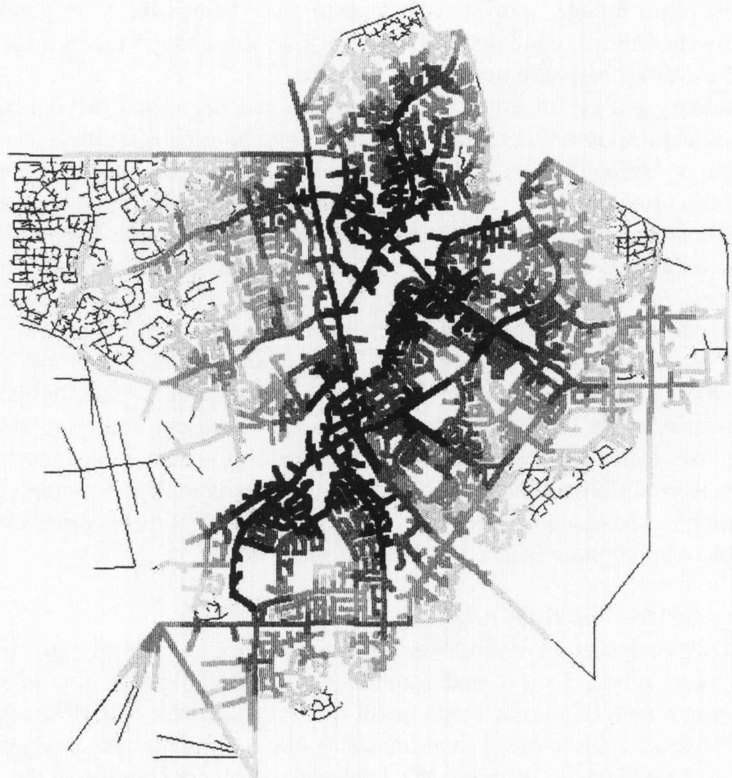

Figure 2.4: Service areas of the existing fire stations during winter rush hour. The street segments that can be reached within 2 minutes (travel time) are represented by thick dark lines. Segments that can be reached within 3, 4, and 5 minutes are shown in thick lines with increasingly lighter shades of gray. The peripheral parts of the network that are thin and black are the parts that cannot be reached within 5 minutes.

to quantify the difference between the coverage provided by two and three stations. We generated a cumulative coverage chart using our spreadsheet. As Figure 2.6 shows, the third station (if located at one of the two suggested locations) increases the winter-rush coverage for 4 minutes of travel from 48% to 68% of the aggregation points in the city. Likewise, the summer non-rush coverage for 3.5 minutes of travel would increase from 61% to 80%.

2.6 Locating the Second Ambulance

After finalizing the location of the third fire station in the Northwest (Site 2 or 3), we moved on to the last phase of the location project: the location of the second ambulance. We considered all possible pairs of the three fire stations and found that the best combination would be to locate the ambulances at the existing two stations. As it turns out, the activation of the second ambulance will reduce the ambulance response times considerably. Figure 2.7 displays the cumulative coverage of a single ambulance located at Fire Hall #1 (Churchill) and the cumulative coverage of two ambulances located at Churchill and Boudreau stations (the optimal configuration for two ambulances).

2.7 Summary

Our conclusions based on an analysis of the status quo were:

- Current performance is within industry standards.
- Only 30% of the city is covered within 5 minutes.

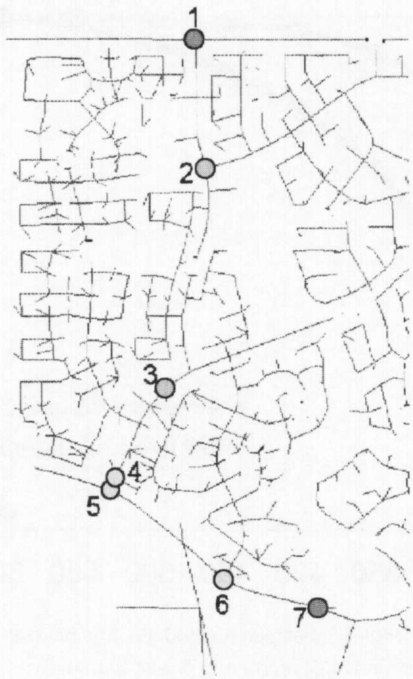

Candidate	Score
3	92%
2	90%
4	82%
6	80%
5	80%
1	53%
7	45%

Figure 2.5: The seven Northwest candidate locations evaluated with respect to the additional coverage provided. Sites 2 and 3, which are more central to the area that is not covered by the two existing station offer the best additional coverage.

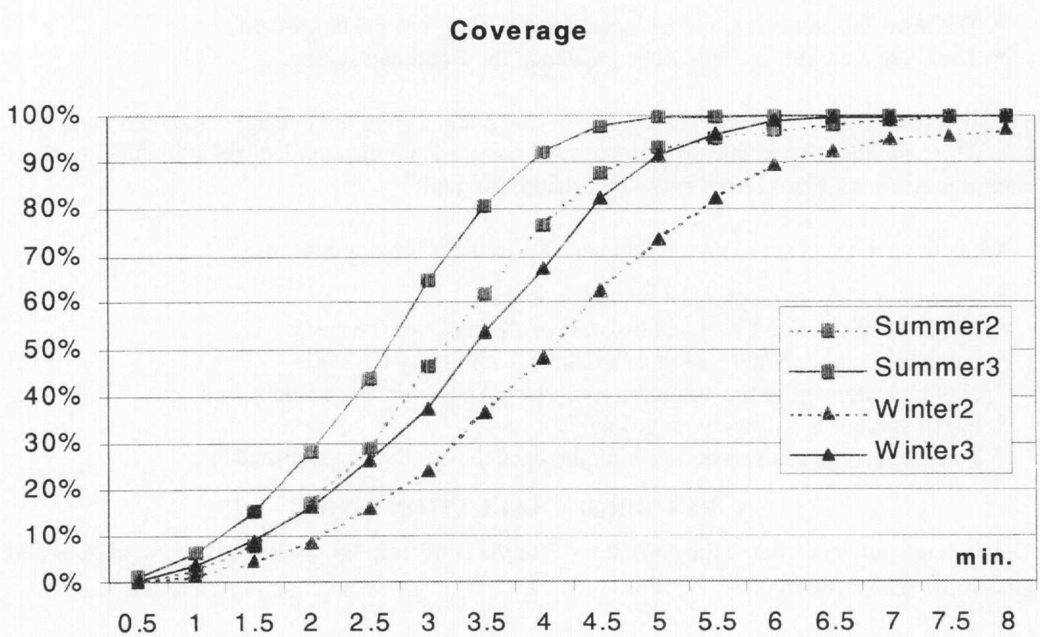

Figure 2.6: The percentage of the aggregation points covered with 2 vs. 3 stations using summer non-rush and winter rush hour travel speeds.

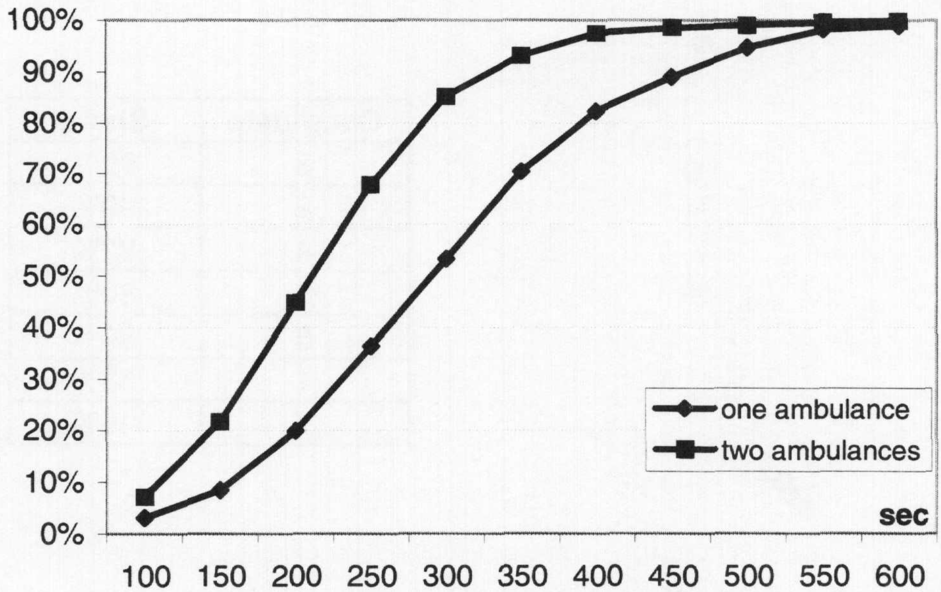

Figure 2.7: The cumulative coverage (travel time) of a single ambulance located at Churchill station (status quo) versus the cumulative coverage of two ambulances, one at Churchill and one at Boudreau.

- Performing dispatch and activation in parallel would reduce response times.
- Most exposed area: *Northwest*.

Our recommendations based on our location study were:

- The third fire station should be located at Site 2 or 3 in the Northwest.
- The second ambulance should be located at the Boudreau station.

The final results were presented to a group of city planners and the manager, and then to the city councilors and the mayor. The project was well received and the city decided to include the third station in its future plans as the new subdivisions are built.

We believe the success of the project was due to the following reasons:

- use of ArcView to display the results on maps,
- inclusion of the Fire Chief and his staff in all phases of the study,
- validation of the models using past data and vetting by the staff,
- explicit statement of the assumptions and weaknesses of the models used,
- use of spreadsheets for coverage analysis, and
- focusing on the analysis and leaving the decisions to the decision-makers.

3. RESOURCE EVALUATION STUDY

The second study considered the current level of the staff from two perspectives: economics and quality of service provided.

3.1 Economic Platoon Sizing

3.1.1 Introduction. The first phase of the study focused on the size of a platoon (i.e. the number of firefighters and EMS staff working a given shift). Staff works 10-hour day-shifts and 14-hour night-shifts, on a two days, two nights, 4 days-off rotation. On each shift there are 10 firefighters, assigned to apparatus located at two fire halls, and two EMS staff, assigned to an ambulance located at Fire Hall #1 (Churchill). This compliment of 12 may be reduced due to vacation, sick time, or other leave. If the staffing compliment for a shift is reduced by three or more, off-duty staff are called in on overtime to achieve a minimum of 10 personnel (4/fire truck, 2/ambulance). Staff can be moved from one fire hall to the other to ensure a minimum fire apparatus assignment of four persons. Firefighters can staff ambulances if EMS staff is absent.

The city was interested in quantifying the labor costs associated with different crew sizes in the existing emergency response platoons. The goal of the study was to measure the economic impact of potential staffing scenarios and to provide recommendations to the city based on our findings. The Fire Service is consistently faced with the issue of calling in off-duty staff to achieve the minimum compliment of 10. These off-duty staff is being paid overtime wages which are twice as high as their normal wages. The purpose of the analysis is to determine the most economic platoon size considering regular salaries and overtime costs. Further, the analysis was extended to include the impact of a seasonal pattern in the rate of staff absenteeism.

3.1.2 Data. There are several reasons a staff member might miss a shift. Sick leave, long term disability, vacation, and training accounted for more than 95% of the absences in 1999. In total this represented 8.19 Full-Time Equivalent person-years lost to absenteeism. Given that each FTE represents 182.5 shifts, this amounts to a total of just less than 1500 lost shifts. The total number of possible shifts is given as: 4 platoons * 12 people/platoon * 182.5 shifts/year = 8,760 shifts per year. The probability of absence can be computed by dividing the total number of absences by the total number of shifts: 1,495 missed shifts per year / 8,760 shifts per year = 0.171. Note that we lump all absences into one category.

3.1.3 Calculating the Expected Number of Absences. We used the binomial probability distribution to compute the probability that a given number of employees will be absent on any given shift.

$$P(x \text{ employees absent}) = \frac{12!}{x!(12-x)!} 0.171^x 0.829^{(12-x)}$$

The use of the binomial distribution assumes that all absences are random, absences are independent of each other (Person A's absence does not impact Person B), and probability of absence is independent of the day of the week, the month of the year, and the time of the shift (day vs. night). Figure 3.1 shows the probability distribution for the absences for a platoon of size 12, the status quo.

The probability that all members of a platoon will show up is 0.106. The most likely scenario entails two absent employees (0.296). The expected number of absentees is 2.05 per shift, and the expected number of overtime workers is 0.52 per shift. The probability of overtime is 0.337.

3.1.4 Calculating Labor Costs. The total labor cost is calculated by adding the regular labor costs to the overtime costs. The regular labor cost is computed using an hourly rate of $22.85 plus 19% benefits. The labor cost for one person is $22.85/hour * 1.19 * 12 hours/shift * 182.5 shifts/year = $59,550/year, and the total cost for a crew of 48 is $2,858,370/year.

Anticipated annual overtime payments are calculated as the average number of staff being paid overtime per shift, multiplied by the overtime wage per employee per shift, multiplied by the number of shifts in a year per platoon, multiplied by the number of platoons. Employees are paid $45.70/hour (double time) for overtime. For a platoon size of 12, the annual overtime expenses are: $45.70/person-hour * 0.52 OT persons/platoon * 12 hours/shift * 182.5 shifts/year * 4 platoons = $208,518/year. The sum of regular and overtime labor costs is estimated to be $3,066,889 for a crew of size 12.

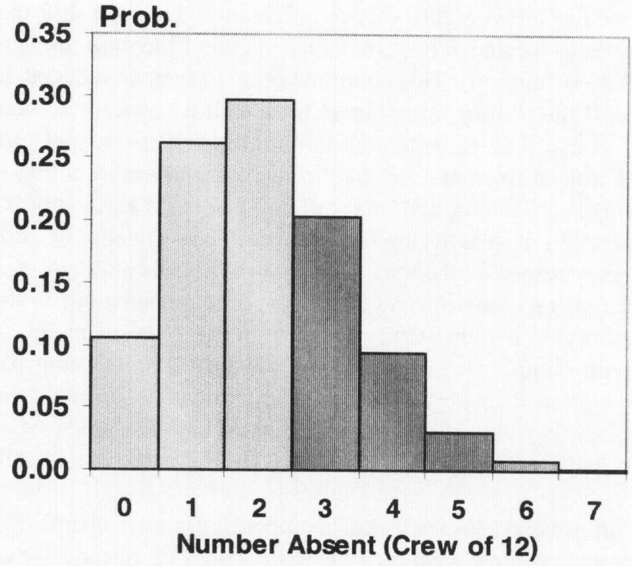

Figure 3.1: The probability of absences in a platoon of 12. Darker gray bars indicate the cases corresponding to overtime.

3.1.5 Incorporating Seasonality. We found out that most employees take their vacation time in the summer and that most training is scheduled in the winter. Since the bulk of the absences are vacation related this introduces seasonality into the problem and brings into question the assumption of uniform absence probability over the year. To deal with the seasonality of the absence probability, we assumed that the year was made up of two 6-month parts, one for summer, and the other for winter. We then assumed that all of the crew's vacation time was taken during the summer (total absenteeism = 5.3 Person-Years or 966 shifts), and that all of the training was taken during the winter (total absenteeism = 2.9 Person-Years or 528 shifts). We further assumed that the remaining causes of absenteeism were uniformly distributed throughout the year.

We conducted the analysis described above in two parallel streams, one for summer the other for winter. This resulted in an increase in the overtime expenses which is attributable to increased overtime payments in the summer. This expected annual overtime costs increased to $229,341, which is less than 5% off from the actual overtime expense in 1999. Although the model used is a very simplistic one, it estimates actual overtime payments relatively well.

3.1.6 Varying the Crew Size and the Rate of Absenteeism. We created a simple-to-use spreadsheet model of the labor costs which allowed the user to experiment with different crew sizes and absence probabilities. Using the estimated absence probability of 0.171, we found that the optimal platoon size is 11. It is possible to save approximately $55,000/year (1.8%) by going from a platoon of 12 to a platoon of 11.

The seniority-based system will allow senior employees more vacation entitlements in the near future. For example by 2006, it is expected that the probability of absence in the summer will increase from 22% to 28%. We conducted parametric analysis on labor costs by varying the probability of absence to gain a better appreciation of the labor costs.

For all reasonable absence probabilities, the lowest-cost platoon size is 11. However, the difference between platoon sizes of 11 and 12 disappears as the P(absence) increases.

3.1.7 Recommendation. Although the lowest cost platoon size was 11, we did not recommend a reduction in the platoon size for reasons we discuss next. With a platoon of size 11, the probabil-

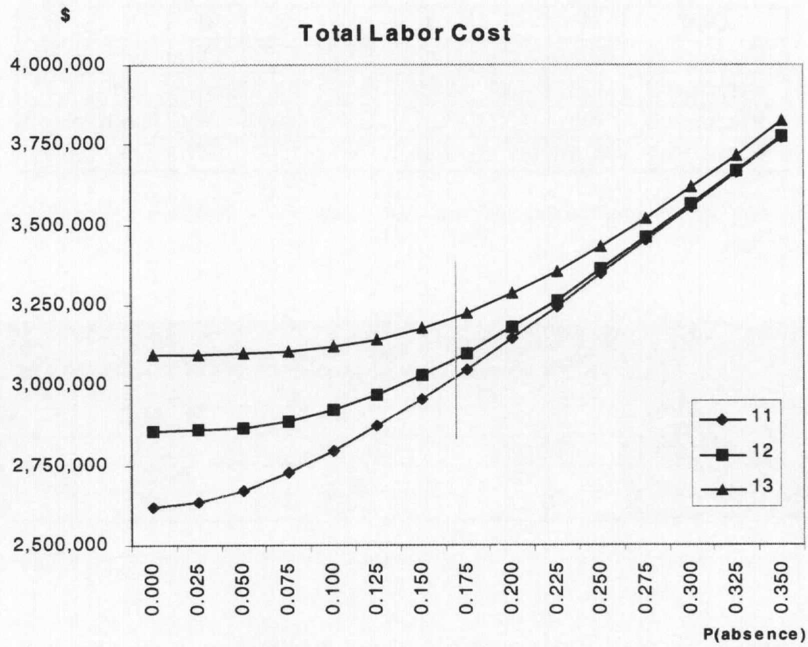

Figure 3.2: The labor cost for platoon sizes of 11, 12, and 13 as a function of P(absence). The vertical gray line indicates the current probability of 0.171.

ity of overtime goes to 0.583 per shift even if we ignore the seasonal effects. Such frequent overtime may not be sustainable for extended periods of time. Increased overtime frequency may impact the employee morale negatively and may even contribute to absenteeism.

Another problem is presented by the small pool of staff available for overtime. Consider the shift schedules in Table 3.1. In half of the 8 days only one of the two non-working platoons has rested for at least a full day. For example, if someone from Platoon 2 is absent on Day 3, it is not possible to draw a member from Platoon 4 for overtime since Platoon 4 completed the night shift the day before. The only platoon that can provide overtime is Platoon 3.

In the other half of the days, one of the two-non-working platoons is about to start work the next day and it would be difficult to recruit an overtime employee from that platoon. For example, on Day 4 the only platoon that can provide overtime (without going into a 5-day work sequence) is Platoon 4. Hence, we have one platoon providing overtime to two platoons on any given day.

The problem of chronic overtime is exasperated over the summer months when many staff members are scheduled for overtime in other platoons due to scheduled vacations. This reduces the number of available individuals to cover absences. During summer months, almost three quarters of all shifts would have at least one employee working on overtime (with 21% of the shifts requiring three or more employees on overtime), and this may be infeasible.

Finally, the Fire Chiefs suggested that the implied reduction in the platoon size from 12 to 11 is likely to impact the quality of the service provided negatively. Given these considerations, we did not think it was a good idea to go to a platoon size of 11.

The discussion above suggests that an increase in the platoon size to 13 may be desirable although the economic analysis indicates this would increase the total labor costs by \$128,400 (4.2%). We recommended consideration of a hybrid solution: hire a seasonal crewmember for each platoon during the summer when the majority of the staff takes their vacations. This would allow them to capitalize on the benefits of a 13-person crew in the summer when they need the

Day	1	2	3	4	5	6	7	8
Platoon 1	Day	Day	Night	Night	Off	Off	Off	Off
Platoon 2	Off	Off	Day	Day	Night	Night	Off	Off
Platoon 3	Off	Off	Off	Off	Day	Day	Night	Night
Platoon 4	Night	Night	Off	Off	Off	Off	Day	Day

Table 3.1: The shift schedule for the four platoons over 8 days. An annual schedule can be generated by cycling this schedule.

	Platoon Size			
Criteria	11	12	13	12.5
Total Labor Cost	$ 3,032,785	$ 3,087,711	$ 3,216,120	$ 3,126,700
P(OT | summer)	73.5%	51.3%	31.8%	31.8%
P(OT | winter)	39.0%	16.9%	6.2%	16.9%
Avg. OT/shift (summer)	1.49	0.92	0.52	0.52
Avg. OT/shift (winter)	0.57	0.23	0.08	0.23

Table 3.2: The performance characteristics of platoon sizes of 11, 12, 13, and 12.5, the hybrid alternative.

additional coverage the most. This alternative would represent an increase in cost of $39,000 only, even if the temporary staff members were paid the same average salary as permanent staff members. We are aware that there may be institutional and labor issues surrounding temporary staff, and the department may not wish to consider this option. Table 3.2 summarizes the performance characteristics for different platoon sizes.

3.2 Service Level Simulation

3.2.1 Introduction. The platoon sizing study summarized in Section 2.1 looked strictly at the economics of different platoon sizes. Perhaps the more important issue is the ability of the department to respond to calls, which is a function of the number of staff members (as well as the number of vehicles). This last phase of the study was intended to investigate the effects of manipulating the call volumes, as well as the distribution of call types and durations, on the available manpower and vehicles. While it is desirable that an emergency service system is always available to respond to calls, there will be times during which all resources of the system are deployed (Code Red) and it will not be possible to respond to an incoming call. It is useful to estimate the probability of such an incident as a function of resource levels under different scenarios. We used Monte-Carlo simulation to deal with the complexity of the EMS operation.

The calls to St. Albert Fire and EMS can be broken into two overlapping groups – fire related calls and medical related calls. All calls to the St. Albert Fire Service are increasing in numbers. At the same time the types of calls that are arriving are shifting with the demographics. As the average age of the people in St. Albert increases, so does the proportion of medical related calls. The goal of this phase is to develop a decision-support tool to evaluate the performance of the system under different scenarios.

3.2.2 Data. We received a database that contained all calls during an 18-month period (January 1999 – July 2000). Each call was documented, describing the call type, the duration, the date/time of the call, and the resources deployed to handle the situation. Calls are classified in 17 different categories. Some calls (such as grass fire or accidental injuries) require minimal resources (a single fire truck or an ambulance), while others (such as structure fires) require the deployment of all vehicles (three fire trucks and an ambulance).

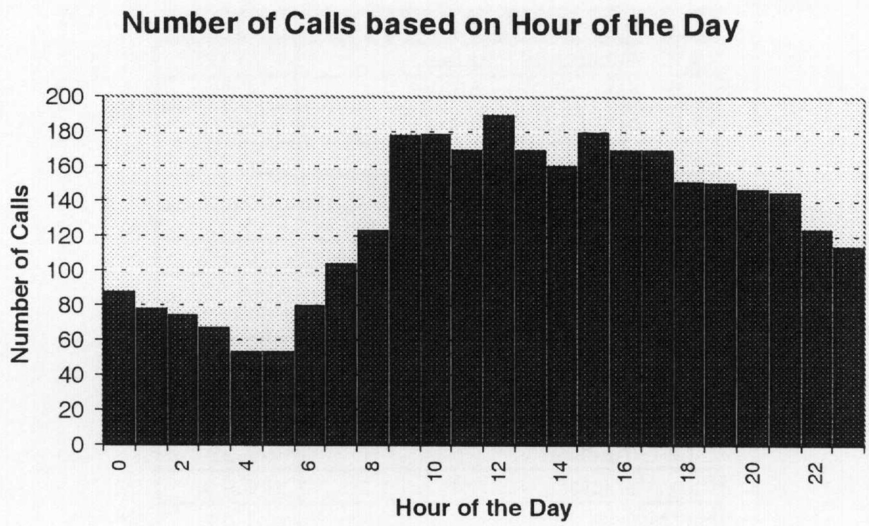

Figure 3.3: The distribution of calls over the hours of the day for the 3106 calls received in an 18-month period (January 1999 – July 2000).

To simulate calls, we need call arrival rates for each type of call for each hour of the day. Defining

n_i = total number of calls of type i ($i = 1, \ldots, 17$),
N = total number of days in the database (562)
t_j = total number of calls during time period j ($j = 0, \ldots, 23$),
M = total number of calls in the database (3106)

and assuming each call type follows the same time-of-day pattern, we can compute R_{ij}, the arrival rate for call type i during time period j, as follows:

$$R_{ij} = \frac{n_i}{N} \frac{t_j}{M}$$

Table 3.3 shows the breakdown of all calls to the 17 call types. Five of the 17 categories make up for over 70% of all calls. The average number of calls per day is 5.53.

We give an example of the arrival rate calculation. There were 342 calls classified as "EMS: Accidental Injuries," in 562 days for an average of 0.61 calls/day. A total of 178 of the 3106 calls (5.7%) arrived between 10:00 a.m. and 11:00 a.m. Hence, the estimated probability of arrival of an EMS Accidental Injury call from 10:00 to 11:00 a.m. is 0.61 * 0.057 = 0.035. The hourly arrival rates range from a low of 0.00006 (EMS: Critical Emergency Transfer between 4:00 and 5:00 a.m.) to a high of 0.088 (EMS: General Medical from noon to 1:00 p.m.).

3.2.3 Simulating Calls and Service. The Poisson distribution can be used to determine the probability that a finite number of calls come in over a time period. In this case the Poisson distribution is used to estimate the probability if the following alternatives for each call type for any given hour: no calls, one call, two or more calls. Since the probability of more than two calls of a given type happen in the same hour is extremely small, the distribution is truncated to exclude situations where more than two calls of a given type occur within one hour. The probability of 0, 1, or 2 incidents occurring is compared to a random number to determine how many incidents of each type occur in each simulated hour. The exact time (in minutes) of the incident is determined by choosing a random number between 1 and 60.

Call Type	% of calls
EMS: Accidental injuries	11.0%
EMS: Alcohol/Drug Abuse	0.2%
EMS: ALS/BLS Transfer	6.2%
EMS: Attempted Suicide	2.4%
EMS: Cardiac Arrest	0.3%
EMS: Collapse/Fainting	4.5%
EMS: Critical Emergency Transfer	0.1%
EMS: General Medical	26.1%
EMS: Heart/Respiratory	17.6%
EMS: Traffic	6.8%
Fire: Grass/Brush	2.5%
Fire: Rubbish	1.4%
Fire: Structural	8.3%
Fire: Vehicle	0.9%
HazMat	1.8%
Miscellaneous	2.0%
Service Call	7.7%

Table 3.3: The breakdown of all calls to the 17 call types.

Both, call duration and resource allocation, were selected by sampling from the actual historical distributions for each call type. Both durations and resources were determined using the same random number which is meant to replicate the severity of the call.

The simulation is repeated for each hour of each day for each incident type for one year. One of the outputs of the model is a database that closely resembles the one that is currently being used to track calls. The second output is a measure of the amount of time the three types of resources (ambulances, fire trucks, staff) spent in each state. For instance, it can be used to measure the amount of time that one fire truck was in use.

*3.2.4 Simulation Logic and Assumptions.*The following pseudo-code summarizes the logic of the simulation program, which was coded in VBA.

```
For each of the 1000 days
    For each of the 17 call types
        For each of the 24 hours of the day
            How many calls came in?
                When did the call come in?
                How long did the call service take?
                What was dispatched to the call?
                Adjust resource levels
            Repeat for all calls
        Repeat for all 24 hours
    Repeat for all 17 types
Repeat for all days
```

Here is a list of the assumptions of the simulation model along with a brief discussion.

- All vehicles and staff dispatched to an incident will remain blocked for the entire duration of the call. This means there is no triage. This is probably a restrictive assumption which may result in an overestimate of staff demand.

% of time spent in each state				
	Increase in Call Volume			
Staff	0%	25%	50%	100%
0	83.89%	80.00%	76.24%	69.54%
2	3.74%	4.75%	5.55%	6.14%
4	1.07%	1.32%	1.59%	1.91%
6	6.26%	6.89%	7.85%	9.36%
8	1.44%	2.00%	2.25%	3.12%
10	0.95%	1.07%	1.38%	2.03%
12	1.15%	2.01%	2.10%	3.00%
14	0.43%	0.68%	0.93%	1.27%
16	0.56%	0.51%	0.93%	1.07%
18	0.24%	0.32%	0.40%	0.80%
20	0.07%	0.14%	0.22%	0.53%
22 or more	0.19%	0.31%	0.57%	1.23%

Table 3.4: The percentage of time spent in each state (in terms of staff deployed on calls) during the 365-day simulation.

- Call duration and the number of vehicles dispatched to an incident are directly related. Longer calls will have more vehicles dispatched to them. This is a realistic assumption.
- In all circumstances, there are 2 crewmembers on an ambulance, and 4 on a fire truck. The department tries hard not to deviate from this.
- No seasonal variability (summer vs. winter) is taken into account. We do not have enough data to estimate seasonal frequencies.
- No more than two incidents of a given type can occur in the same hour. In the even that two incidents occur within the same hour, one will happen in the first half of the hour and the other will happen in the second half of the hour. It is extremely unlikely that more than two incidents of the same type will happen within one hour.
- Resources are assumed to the unlimited and there is no queueing of calls. If a call comes in, the program assumes that the appropriate crew will respond to it. It does not wait until a service is completed before deploying resources. This way we can generate the state probabilities for all crew sizes (going to infinity). This will allow us to estimate code yellow (all resources deployed) and code red (all resources deployed and a call comes in) probabilities.

3.2.5 Results. Table 3.4 shows the amount of time spent in each staff state based on a simulation of 365 days.

Under the current call rates everyone is idle about 84%. Yet there are some instances that require many more staff than is available. The maximum number of staff needed during the simulation with the current call volume was 36. On June 4, 2000, at 1:30 a.m. there was a commercial structural fire requiring 3 pumps, one ladder truck, 2 ambulances. The incident lasted 13 hours and 47 minutes. Since we are sampling service durations and vehicle assignments from actual data, this event has a small but positive probability of being selected whenever a structural fire is scheduled by the simulation code. If another serious event occurs during these 13 hours and 47 minutes, then the total staff needed can go up to 36. This example makes it clear that any reasonable platoon size (12 – 16) will result in a positive probability for Code Red.

The system requires more than 12 crewmembers (i.e. "guaranteed" Code Red) only around 1.5% of the time. This goes up to 2% with a 25% increase in the call volumes, and to 3% with a 50% increase in call volumes. However, recall that the minimum allowable platoon size is 10. If two or more individuals are absent (a likely event with a probability of 0.633), then the platoon is down to 10 persons. A reduced platoon of 10 is inadequate to deal with calls 2.7% of the time—almost twice as often as a full platoon. With a 25% increase in call volume, the probability of a platoon of 10 not being adequate goes to 0.04, a relatively high figure for EMS.

While it may seem somewhat surprising at first that the system performance does not drastically deteriorate under increased call volumes, we have to keep in mind that the system is designed with a lot of redundancy—everyone is idle 85% of the time under the current call volume. Hence, the current system can absorb significant increases in the call volume. However, note that in this sector the system performance is measured by focusing on the right tail of the distribution. A system that is in code red 4% of the time is unlikely to satisfy the 90%-of-calls-within-9-minutes standard. It seems that platoon sizes must be increased if the call volume increases by 25%.

Currently, one fire truck is used 8.0% of the time, two are used 1.8% of the time, and all three are used 0.8% of the time. One ambulance is needed 11.3% of the time, two are used 2.9% of the time, and three are needed 0.4% of the time. In severe incidents, additional staff is brought in so the department can mobilize all of its vehicles (3 fire trucks, one ladder truck, and 3 ambulances). Hence, it seems that staff is the bottleneck and not the vehicles.

The simulation study seems to lend further support to an increased platoon size. With a platoon of size 12, the probability of having all 12 staff members on a given shift is only 0.106. However with a platoon of size 13, the probability of having 12 or 13 on a given shift goes up to 0.323. The simulation model suggests that the percentage of time a staff of 12 is required is not trivial (ranging from 1.15% to 3.00% for different call volumes) and an increase in the probability of having 12 persons on a shift has an impact on the Code Red probability. In fact, P(Code Red | 12-person platoon) = 3.8% and P(Code Red | 13-person platoon) = 3.3% for the scenario with the 25% increase in call volume. It seems that the cost of adding the 13th member to a platoon does bring some quantifiable benefit in terms of the quality of the service provided.

The proposed addition of a new fire hall, equipped and staffed with fire apparatus, will increase the platoon size by 4 members. This will, in turn reduce the Code Red probability significantly. The new fire hall will probably be in place in 2006. Using the 25% increase in the call volume scenario, we evaluate the probability of Code Red as 0.77% if there are 16 staff members on the shift, and 1.28% if 14 or 15 staff members are on the shift. It is unrealistic to reduce this probability to zero, and every community must decide what is acceptable risk and how much they are willing to pay to get there. It seems that the probabilities we calculated here constitute acceptable risk for emergency service delivery in St. Albert, and they are willing to pay the price of a new fire hall ($2M), new fire truck ($0.4M), and 16 new staff members ($1M/year) to achieve this increase in service quality.

4. CONCLUSIONS

This paper summarizes two projects completed for the emergency services department of St. Albert. We found that the quality of the current service was within acceptable limits. However the planned growth of the city coupled with an increase in the per-capita call volume and an aging staff will force the city to spend more money to provide the same quality of service to the residents.

Our recommendations can be summarized as follows:

1) Locate a third fire hall in the Northwest to improve coverage
2) Staff a second ambulance (in Boudreau) to improve response times
3) Consider changing the dispatch and activation process to speed up responses
4) Go to a platoon of size 12.5 by using temporary staff over the summer
5) Consider a platoon size of 13 to reduce Code Red probability

The city has already included a new fire station in its capital plans for the Year 2004. However, the exact date of construction is not yet determined since it will depend on how quickly the new developments are filled with structures. The department is planning to phase-in the second ambulance by staffing it during the busier part of the day, and it may become fully staffed in the near future. A temporary staff was deemed to be unacceptable, and they will increase the platoon size to 13 soon.

We consider the study a success. We had an opportunity to apply what we teach in our programs to a real-world problem, and several students (Fenske, Kabanuk, and Gardiner) had interesting and useful experiences. We are currently using parts of the projects in teaching management science and operations management courses in the School of Business at the University of Alberta. We also plan on maintaining a professional relationship with the City of St. Albert and offer them similar services for other interesting urban operational research problems.

Erhan Erkut is a Professor of Management Science in the School of Business at the University of Alberta. His research areas include location of facilities and transportation of hazardous materials. He has published papers in journals such as *Operations Research, Transportation Science, European Journal of Operational Research, Networks, Interfaces, Computers & Operations Research, Naval Research Logistics, Journal of Business Logistics,* and *INFOR.* He teaches introduction to business, operations management, and management consulting courses, and he is involved in consulting projects with local organizations. He is the Editor-in-Chief of *INFORMS Transactions in Education.*

Russell Fenske is currently studying Management Science in the Master of Science program in the faculty of Business at the University of British Columbia. His received his undergraduate degree and CORS diploma at the University of Alberta in 1999.

Steven Kabanuk graduated from the School of Business at the University of Alberta in 1998. He received an M.Sc. in Management Science from the University of British Columbia in December of 2000. Currently, he works as a consultant for RP Options in Vancouver and as a project manager in the Centre for Operations Excellence at the University of British Columbia.

Quentin Gardiner graduated from the University of Alberta in April 2000. He is pursuing his CA designation and wrote the UFE in September 2001. He is currently employed with Stout & Company Chartered Accountants in Edmonton.